

Traffic Management on the QFX Series

Release
13.2X52



Published: 2014-07-15

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Traffic Management on the QFX Series

13.2X52

Copyright © 2014, Juniper Networks, Inc.

All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About the Documentation	xxiii
Documentation and Release Notes	xxiii
Supported Platforms	xxiii
Using the Examples in This Manual	xxiii
Merging a Full Example	xxiv
Merging a Snippet	xxiv
Documentation Conventions	xxv
Documentation Feedback	xxvii
Requesting Technical Support	xxvii
Self-Help Online Tools and Resources	xxvii
Opening a Case with JTAC	xxviii

Part 1

Chapter 1

Overview

Software Features Overview	3
Overview of Junos OS CoS for the QFX Series	4
CoS Standards	4
How Junos CoS Works	5
Default CoS Behavior	6
Overview of CoS Upgrade Requirements (Junos OS Release 11.1 or 11.2 to a Later Release)	7
Overview of CoS Upgrade Requirements to Junos OS Release 12.2	8
Overview of CoS Upgrade Requirements to Junos OS Release 12.3 (QFX3500 and QFX3600 Switches) or to Junos OS Release 13.1 (QFabric Systems) . . .	10
Support for Six Lossless Forwarding Classes	10
Scheduling on QFabric System Node Device Fabric (fte) Ports	12
Strict-High Priority Scheduling on QFabric System Node Device Fabric (fte) Ports	12
Overview of CoS Changes Introduced in Junos OS Release 11.3	13
CoS Default Value Changes	13
Default Forwarding Classes	13
Default IEEE 802.1p Unicast Classifiers	15
Default IEEE 802.1p Multidestination Classifiers	16
Default Scheduler	17
Queue Priority Configuration Changes	18
Minimum Guaranteed Bandwidth (Transmit Rate and Guaranteed Rate) Changes	19
Excess Rate Statement Disabled	19
Queue Scheduling (Low and Strict-High Priority Queues)	20
Multidestination Traffic Changes	20

	Overview of CoS Changes Introduced in Junos OS Release 12.2	21
	Lossless Forwarding Classes (fcoe and no-loss)	21
	Default MTU for Headroom Buffer Calculation for Lossless Forwarding Classes	22
	CoS for Layer 3 Physical Interfaces	22
	DSCP IPv6 Classifiers and Rewrite Rules	22
	Overview of Policers	23
	Policer Overview	23
	Policer Types	24
	Policer Actions	25
	Policer Colors	25
	Filter-Specific Policers	26
	Suggested Naming Convention for Policers	26
	Policer Counters	27
	Policer Algorithms	27
	How Many Policers are Supported?	27
	Policers can Limit Egress Firewall Filters	27
Chapter 2	CoS Overview	29
	Understanding Junos CoS Components	30
	Code-Point Aliases	30
	Policers	30
	Classifiers	30
	Forwarding Classes	31
	Forwarding Class Sets	31
	Flow Control (Ethernet PAUSE, PFC, and ECN)	31
	WRED Profiles	32
	Schedulers	33
	Rewrite Rules	33
	Understanding CoS Packet Flow	34
	CoS Inputs and Outputs Overview	36
	Understanding Default CoS Settings	37
	Default Forwarding Classes and Queue Mapping	37
	Default Forwarding Class Sets (Priority Groups)	38
	Default Code-Point Aliases	38
	Default Classifiers	40
	Default Rewrite Rules	43
	Default Drop Profile	43
	Default Schedulers	43
	Default Scheduler Maps	45
	Default Shared Buffer Configuration	45
	Understanding Host Inbound Traffic Classification	46
	Understanding Host Routing Engine Outbound Traffic Queues and Defaults	47
	Understanding CoS Code-Point Aliases	49
	Default Code-Point Aliases	49

Understanding CoS Classifiers	51
Interfaces and Output Queues	51
Behavior Aggregate Classifiers	52
Default Behavior Aggregate Classification	53
Importing a Classifier	54
Multidestination Classifiers	54
PFC Priorities	54
Fixed Classifiers on Ethernet Interfaces	55
Fixed Classifiers on Native Fibre Channel Interfaces (NP_Ports)	56
Multifield Classifiers	56
Packet Classification for Routed VLAN Interfaces (RVIs)	57
Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces	57
Supported Classifier and Rewrite Rule Types	58
Ethernet Interfaces Supported for Classifier and Rewrite Rule	
Configuration	59
Interface Types That Support Classifier and Rewrite Rule	
Configuration	59
Classifier and Rewrite Rule Physical and Logical Ethernet Interface	
Support	59
Routed VLAN Interfaces (RVIs) and Integrated Routing and Bridging	
(IRB) Interfaces	61
Default Classifiers	61
Default Rewrite Rules	61
Classifier Precedence	61
Unicast Classifier Precedence on Physical Ethernet Interfaces	62
Unicast Classifier Precedence on Logical Ethernet Interfaces	62
Classifier Behavior and Limitations	63
Rewrite Rule Precedence and Behavior	64
Classifier and Rewrite Rule Configuration Interaction with Ethernet Interface	
Configuration	64
Scenario 1: Applying a Classifier to an Ethernet Interface Causes a	
Commit Check Error	66
Scenario 2: Configuring the Ethernet Interface Family Causes a Syslog	
Error	66
Understanding CoS Forwarding Classes	68
Default Forwarding Classes	68
Forwarding Class Configuration Rules	70
Queue Assignment Rules	70
Scheduling Rules	70
Rewrite Rules	71
Lossless Transport Support	71
Understanding CoS Forwarding Class Sets (Priority Groups)	73
Understanding Default CoS Scheduling and Classification	74
Default Classification	75
Default Scheduling	78
Default DCBX Advertisement	80
Default Scheduling and Classification Summary	80

Understanding CoS Hierarchical Port Scheduling (ETS)	81
Hierarchical Scheduling Tiers	81
Hierarchical Scheduling and ETS	82
ETS Advertisement in DCBX	83
Hierarchical Scheduling Process	83
Strict-High Priority Queues and Hierarchical Scheduling	84
Default Hierarchical Scheduling	85
Understanding CoS Output Queue Schedulers	86
Output Queue Scheduling Components	86
Default Schedulers	87
Transmit Rate (Minimum Guaranteed Bandwidth)	90
Sharing Extra Bandwidth	90
Shaping Rate (Maximum Bandwidth)	91
Scheduling Priority	91
Scheduler Drop-Profile Maps	92
Buffer Size	92
Explicit Congestion Notification	93
Scheduler Maps	94
Understanding CoS Priority Group Scheduling	95
Priority Group Scheduling Components	95
Default Traffic Control Profile	96
Guaranteed Rate (Minimum Guaranteed Bandwidth)	96
Sharing Extra Bandwidth	96
Shaping Rate (Maximum Bandwidth)	97
Scheduler Maps	97
Understanding CoS Traffic Control Profiles	98
Understanding CoS Priority Group and Queue Guaranteed Rates (Minimum Bandwidth)	99
Guaranteeing Bandwidth Using Hierarchical Scheduling	99
Priority Group Guaranteed Rate (Minimum Bandwidth)	101
Queue Transmit Rate (Minimum Bandwidth)	101
Understanding CoS Priority Group Shaping and Queue Shaping (Maximum Bandwidth)	102
Priority Group Shaping	102
Queue Shaping	102
Shaping Maximum Bandwidth Using Hierarchical Scheduling	103
Understanding CoS Scheduling Behavior and Configuration Considerations	104
Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows	109
Lossless Transport Features Introduced in Junos OS Release 12.3	110
Default Lossless Priority Configuration	110
Configuring Lossless Priorities	113
Configuring Lossless Forwarding Classes (Packet Drop Attribute)	113
Congestion Notification Profiles (PFC Configuration)	115
Configuring DCBX (Application Protocol TLV Exchange)	121
Fate Sharing Among Traffic Classes	121
Transit Switch Configuration Versus FCoE-FC Gateway Configuration	123
Configuration Results and Commit Checks	123

Backward Compatibility with Junos OS Releases Earlier Than Release 12.3	125
Configuration Rules and Recommendations	126
Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway	128
Priority Remapping Configuration	128
Configuration Rules	129
Fate Sharing	130
Understanding CoS Buffer Configuration	131
Buffer Pools	132
Buffer Handling of Lossless Flows (PFC) Versus Ethernet PAUSE	133
Shared Buffer Pool and Partitions	134
Dedicated Port Buffer Pool and Buffer Allocation to Queues	135
Trade-off Between Shared Buffer Space and Dedicated Buffer Space	139
Order of Buffer Consumption	140
Default Buffer Pool Values	140
Total Buffer Pool Size	141
Shared Buffer Pool Default Values	141
Dedicated Buffer Pool Default Values	143
Shared Buffer Configuration Recommendations for Different Network Traffic Scenarios	143
Balanced Traffic (Default Configuration)	144
Best-Effort Unicast Traffic	144
Ethernet PAUSE Traffic	145
Best-Effort Multicast (Multidestination) Traffic	146
Lossless Traffic	146
Optimizing Buffer Configuration	147
General Buffer Configuration Rules and Considerations	148
Understanding CoS WRED Drop Profiles	149
Drop Profile Parameters	150
Default Drop Profile	151
Packet Drop Method	151
Drop Profile Maps	152
Congestion Prevention	152
Configuring a WRED Drop Profile and Applying it to an Output Queue	153
Understanding CoS Rewrite Rules	154
Understanding DCB Features and Requirements	156
Lossless Transport	156
PFC	157
Buffer Management	157
Physical Interfaces	157
ETS	157
DCBX	158
Understanding CoS Flow Control (Ethernet PAUSE and PFC)	160
Ethernet PAUSE	161
Symmetric Flow Control	162
Asymmetric Flow Control	162
PFC	165
Lossless Transport Support Summary	168

	Understanding DCBX	170
	DCBX Basics	170
	DCBX Modes and Support	171
	DCBX Modes (Versions)	171
	Autonegotiation	173
	CNA Support for DCBX Modes	174
	Interface Support for DCBX	174
	DCBX Attribute Types	174
	Asymmetric Attributes	174
	Symmetric Attributes	175
	DCBX Application Protocol TLV Exchange	175
	Application Protocol TLV Exchange	175
	FCoE Application Protocol TLV Exchange	175
	Disabling Application Protocol TLV Exchange	176
	DCBX and PFC	176
	DCBX and ETS	177
	Default DCBX ETS Advertisement	177
	ETS Advertisement and Peer Configuration	177
	ETS Recommendation TLV	178
	Understanding DCBX Application Protocol TLV Exchange	179
	Applications	180
	Application Maps	180
	Classifying and Prioritizing Application Traffic	181
	Enabling Interfaces to Exchange Application Protocol Information	182
	Disabling DCBX Application Protocol Exchange	182
Chapter 3	QFabric-Specific CoS Overview	185
	Understanding CoS Fabric Forwarding Class Sets	186
	Default Fabric Forwarding Class Sets	187
	Fabric Forwarding Class Set Configuration and Implementation	190
	Mapping Forwarding Classes to Fabric Forwarding Class Sets	191
	Fabric Forwarding Class Set Implementation	191
	Fabric Forwarding Class Set Scheduling (CoS)	192
	Class Groups for Fabric Forwarding Class Sets	192
	Class Group Scheduling	192
	QFabric System CoS	194
	Support for Flow Control and Lossless Transport Across the Fabric	194
	Viewing Fabric Forwarding Class Set Information	196
	Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences	198
	Understanding CoS Scheduling on QFabric System Node Device Fabric (fte) Ports	199
	Hierarchical Scheduling Architecture on QFabric System Node Devices	199
	Default Scheduling on Node Device Fabric Interfaces	200
	Configuring Scheduling on Node Device Fabric Interfaces	201
	Similarities Between Node Device Fabric Interface and Access Interface Scheduling	202
	Differences Between Node Device Fabric Interface and Access Interface Scheduling	202

	Understanding Default CoS Scheduling on QFabric System Interconnect Devices (Junos OS Release 13.1 and Later Releases)	203
	Hierarchical CoS Architecture Across a QFabric System Interconnect Device	203
	Fabric FC-Sets	204
	Class Groups for Fabric FC-Sets	204
	Default CoS on Interconnect Device Fabric Interfaces	205
	Default Class Group Scheduling	205
	Default Fabric FC-Set Scheduling	206
	Default Class Group and Fabric FC-Set Scheduling Example	209
	Default PFC and Lossless Transport Across the Interconnect Device . .	210
Chapter 4	Learn About Technology	215
	Data Center Technology Overview Videos	215
	Learn About Video: Why Do We Need an IP Fabric?	215
	Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?	215
	Learn About Video: Why Use an Overlay Network in a Data Center?	215
	Conceptual Documents That Contain Technology Overview Videos	216
Part 2	Configuration	
Chapter 5	Configuration Examples	219
	Example: Configuring CoS Hierarchical Port Scheduling (ETS)	220
	Example: Configuring CoS PFC for FCoE Traffic	241
	Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG . .	249
	Example: Configuring Unicast Classifiers	273
	Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers	277
	Example: Configuring WRED Drop Profiles	279
	Example: Configuring Drop Profile Maps	281
	Example: Configuring Forwarding Classes	283
	Example: Configuring Forwarding Class Sets	286
	Example: Configuring Queue Schedulers	289
	Example: Configuring Queue Scheduling Priority	295
	Example: Configuring Traffic Control Profiles (Priority Group Scheduling)	298
	Example: Configuring Minimum Guaranteed Output Bandwidth	300
	Example: Configuring Maximum Output Bandwidth	305
	Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)	309
	Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface	317
	Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces	326
	Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)	340
	Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway	356

	Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic	365
	Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled	371
	Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic	377
	Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic	383
	Example: Configuring DCBX Application Protocol TLV Exchange	389
Chapter 6	Configuration Tasks	401
	Configuring CoS	402
	Defining CoS Code-Point Aliases	404
	Defining CoS Unicast BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)	404
	Defining CoS Multidestination (Multicast, Broadcast, DLF) BA Classifiers	406
	Configuring CoS WRED Drop Profiles	407
	Configuring CoS Drop Profile Maps	408
	Defining CoS Forwarding Classes	408
	Defining CoS Forwarding Class Sets	410
	Defining CoS Queue Schedulers	410
	Defining CoS Queue Scheduling Priority	414
	Changing the Host Outbound Traffic Default Queue Mapping	415
	Defining CoS Traffic Control Profiles (Priority Group Scheduling)	416
	Configuring CoS PFC (Congestion Notification Profiles)	417
	Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control	420
	Configuring CoS Asymmetric Ethernet PAUSE Flow Control	421
	Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces (NP_Ports)	423
	Configuring Global Ingress and Egress Shared Buffers	425
	Defining CoS Rewrite Rules	427
	Assigning CoS Components to Interfaces	429
	Configuring the DCBX Mode	430
	Configuring DCBX Autonegotiation	431
	Disabling the ETS Recommendation TLV	434
	Defining an Application for DCBX Application Protocol TLV Exchange	434
	Configuring an Application Map for DCBX Application Protocol TLV Exchange	436
	Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange	437
Chapter 7	Configuration Statements	439
	application (Application Maps)	442
	application (Applications)	443
	application-map	444
	application-maps	445
	applications (Applications)	446
	applications (DCBX)	447
	buffer-partition (Egress)	448
	buffer-partition (Ingress)	450
	buffer-size	452

cable-length (Congestion Notification)	454
class-of-service	455
class (Forwarding Classes)	459
class (Forwarding Class Sets)	460
classifiers	461
code-point (Fibre Channel Interfaces)	463
code-point (Input Congestion Notification)	464
code-point (Output Congestion Notification)	465
code-point (Rewrite Rules)	466
code-point-aliases	466
code-points (Application Maps)	467
code-points (CoS)	467
configured-flow-control	468
congestion-notification-profile	469
dcbx	471
dcbx-version	472
destination-port (Applications)	473
disable (DCBX)	474
drop-probability	475
drop-profile	476
drop-profile-map	476
drop-profiles	477
dscp	478
dscp-ipv6	480
dscp-code-point	481
egress (Buffer Configuration)	482
enhanced-transmission-selection	483
ether-type	484
fill-level	485
flow-control	486
flow-control-queue (Output Congestion Notification)	487
forwarding-class	489
forwarding-class (Fibre Channel Interfaces)	491
forwarding-class (Host Outbound Traffic)	492
forwarding-class-set	493
forwarding-class-sets	493
forwarding-classes	494
guaranteed-rate	496
host-outbound-traffic	497
ieee-802.1	498
ieee-802.1 (Fibre Channel Interfaces)	500
ieee-802.1 (Input Congestion Notification)	501
ieee-802.1 (Output Congestion Notification)	502
import	503
ingress (Buffer Configuration)	504
input (Congestion Notification)	505
input (Fibre Channel Interfaces)	506
interface (DCBX)	507
interfaces (Class of Service)	508

	interpolate	509
	loss-priority (Classifiers)	510
	loss-priority (Drop Profiles)	511
	loss-priority (Rewrite Rules)	512
	multi-destination	513
	mru	514
	output (Congestion Notification)	515
	output-traffic-control-profile	516
	pf (Input Congestion Notification)	517
	policy-options	518
	priority (Schedulers)	519
	priority-flow-control	520
	protocol (Applications)	521
	protocol (Drop Profile Map)	522
	queue-num	523
	recommendation-tlv	524
	rewrite-rules	525
	rewrite-value (Fibre Channel Interfaces)	526
	rx-buffers	528
	scheduler	529
	scheduler-map	529
	scheduler-maps	530
	schedulers	531
	shaping-rate	532
	shared-buffer	534
	traceoptions (Class of Service)	536
	traffic-control-profiles	538
	transmit-rate	539
	tx-buffers	541
	unit	542
Part 3	Administration	
Chapter 8	Routine Monitoring	545
	Monitoring CoS Classifiers	545
	Monitoring CoS Forwarding Classes	546
	Monitoring Interfaces That Have CoS Components	547
	Monitoring CoS Rewrite Rules	548
	Monitoring CoS Scheduler Maps	549
	Monitoring CoS Value Aliases	550
Chapter 9	Operational Commands	553
	show class-of-service	554
	show class-of-service classifier	558
	show class-of-service code-point-aliases	560
	show class-of-service congestion-notification	562
	show class-of-service drop-profile	565
	show class-of-service forwarding-class	568
	show class-of-service forwarding-class-set	570
	show class-of-service forwarding-table	572

show class-of-service forwarding-table classifier	576
show class-of-service forwarding-table classifier mapping	578
show class-of-service forwarding-table drop-profile	580
show class-of-service forwarding-table rewrite-rule	582
show class-of-service forwarding-table rewrite-rule mapping	584
show class-of-service forwarding-table scheduler-map	585
show class-of-service interface	587
show class-of-service multi-destination	613
show class-of-service rewrite-rule	614
show class-of-service scheduler-map	616
show class-of-service shared-buffer	618
show class-of-service traffic-control-profile	620
show dcbx	624
show dcbx neighbors	625
show interfaces queue	647
show pfe next-hop	685
show pfe route	690
show pfe terse	696
show pfe version	698

Part 4

Chapter 10

Troubleshooting

Troubleshooting Procedures	701
Troubleshooting Dropped FCoE Traffic	701
Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth	704
Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth	705
Troubleshooting Egress Queue Bandwidth Impacted by Congestion	706
Troubleshooting an Unexpected Rewrite Value	707
Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic	708

List of Figures

Part 1	Overview	
Chapter 1	Software Features Overview	3
	Figure 1: Packet Flow Across the Network	6
	Figure 2: Flow of Tricolor Marking Policer Operation	24
Chapter 2	CoS Overview	29
	Figure 3: CoS Classifier, Queues, and Scheduler	35
	Figure 4: Packet Flow Through Configurable CoS Components	35
	Figure 5: Hierarchical Scheduling Tiers	82
	Figure 6: Hierarchical Scheduling Packet Flow	84
	Figure 7: Allocating Guaranteed Bandwidth Using Hierarchical Scheduling	100
	Figure 8: Setting Maximum Bandwidth Using Hierarchical Scheduling	104
	Figure 9: WRED-Drop Profile Packet Drop	150
Part 2	Configuration	
Chapter 5	Configuration Examples	219
	Figure 10: Hierarchical Port Scheduling Components Block Diagram	224
	Figure 11: Hierarchical Port Scheduling Packet Flow Block Diagram	224
	Figure 12: PFC for FCoE Traffic Configuration Components Block Diagram	244
	Figure 13: Supported Topology for an MC-LAG on an FCoE Transit Switch	252
	Figure 14: WRED Drop Profile Packet Drop Example	280
	Figure 15: Topology of the Two Lossless FCoE Priorities Example	328
	Figure 16: Topology of the Lossless FCoE and iSCSI Priorities Example	342
	Figure 17: Topology of the IEEE 802.1p Priority Remapping Example	358

List of Tables

	About the Documentation	xxiii
	Table 1: Notice Icons	xxv
	Table 2: Text and Syntax Conventions	xxv
Part 1	Overview	
Chapter 1	Software Features Overview	3
	Table 3: Junos OS Release 11.1 and 11.2 Default Forwarding Classes and Queue Mapping	13
	Table 4: Junos OS Release 11.3 Default Forwarding Classes and Queue Mapping	14
	Table 5: Junos OS Release 11.1 and 11.2 Default IEEE 802.1 Unicast Classifiers	15
	Table 6: Junos OS Release 11.3 Default IEEE 802.1 Unicast Classifiers	15
	Table 7: Junos OS Release 11.1 and 11.2 Default IEEE 802.1 Multidestination Classifiers	16
	Table 8: Junos OS Release 11.3 Default IEEE 802.1 Multidestination Classifiers	16
	Table 9: Junos OS Release 11.1 and 11.2 Default Schedulers	17
	Table 10: Default Schedulers	18
	Table 11: Policer Actions	25
Chapter 2	CoS Overview	29
	Table 12: CoS Mappings—Inputs and Outputs	36
	Table 13: Default Forwarding Classes and Queue Mapping	37
	Table 14: Default IEEE 802.1 Code-Point Aliases	38
	Table 15: Default DSCP and DCSP IPv6 Code-Point Aliases	39
	Table 16: Default IEEE 802.1 Unicast Classifiers for Ports in Trunk Mode or Tagged Access Mode (Trusted Classifier)	40
	Table 17: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier)	41
	Table 18: Default IEEE 802.1 Multidestination Classifiers	41
	Table 19: Default DSCP IP and IPv6 Unicast Classifiers	42
	Table 20: Default Drop Profile	43
	Table 21: Default Schedulers	43
	Table 22: Default Scheduler Maps	45
	Table 23: Default Ingress Shared Buffer Configuration	45
	Table 24: Default Egress Shared Buffer Configuration	45
	Table 25: Routing Engine Protocol Default Queue Mapping	47
	Table 26: Default IEEE 802.1 Code-Point Aliases	49
	Table 27: Default DSCP and DSCP IPv6 Code-Point Aliases	50
	Table 28: Default BA Classification	53

Table 29: Default IEEE 802.1p Code Point to PFC Priority, Output Queue, and Forwarding Class Mapping	55
Table 30: Supported Classifiers and Rewrite Rules	58
Table 31: Ethernet Interface Support for Classifier and Rewrite Rule Configuration	60
Table 32: Default Forwarding Classes for Unicast Packets	69
Table 33: Default Forwarding Classes for Multicast Packets	69
Table 34: Default IEEE 802.1 Unicast Classifiers for Ports in Trunk Mode or Tagged-Access Mode (Trusted Classifier)	75
Table 35: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier)	75
Table 36: Default IEEE 802.1 Multidestination Classifiers	76
Table 37: Default DSCP IP and IPv6 Unicast Classifiers	76
Table 38: Default Scheduler Configuration	78
Table 39: Hierarchical Scheduling Tiers	81
Table 40: Output Queue Scheduler Components	87
Table 41: Other Scheduling Components	87
Table 42: Default Schedulers	88
Table 43: Priority Group Scheduler Components	95
Table 44: Other Scheduling Components	96
Table 45: Mapping of Default Unicast Forwarding Class to Queue, IEEE 802.1p Priority, and Drop Attribute	112
Table 46: FCoE and No-Loss Forwarding Class Configuration in Junos OS Release 12.3	114
Table 47: Default Output Flow Control Profile	119
Table 48: User-Configured Output Flow Control Profile	120
Table 49: Results of Lossless Priority Configuration	124
Table 50: Default Dedicated Buffer Allocation to Egress Queues (Based on Default Scheduler)	136
Table 51: Egress Queue Dedicated Buffer Allocation (Example 1)	138
Table 52: Egress Queue Dedicated Buffer Allocation with Another Remainder Queue (Example 2)	139
Table 53: QFX5100 Switch Default Shared Ingress Buffer Values (KB)	141
Table 54: QFX3500 and QFX3600 Switch Default Shared Ingress Buffer Values (KB)	142
Table 55: Default Shared Ingress Buffer Values (Percentage)	142
Table 56: QFX5100 Switch Default Shared Egress Buffer Values (KB)	142
Table 57: QFX3500 and QFX3600 Switch Default Shared Egress Buffer Values (KB)	142
Table 58: Default Shared Egress Buffer Values (Percentage)	142
Table 59: Default Ingress Shared Buffer Configuration	144
Table 60: Default Egress Shared Buffer Configuration	144
Table 61: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Unicast Traffic	145
Table 62: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Unicast Traffic	145
Table 63: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Traffic and Ethernet PAUSE Enabled	145

	Table 64: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Traffic and Ethernet PAUSE Enabled	145
	Table 65: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Multicast Traffic	146
	Table 66: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Multicast Traffic	146
	Table 67: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Lossless Traffic	146
	Table 68: Recommended Egress Shared Buffer Configuration for Networks with Mostly Lossless Traffic	146
	Table 69: Asymmetric Ethernet PAUSE Flow Control Configuration	162
	Table 70: Flow Control State Advertised to the Connected Peer (Autonegotiation)	163
	Table 71: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces . .	164
	Table 72: Default PFC Priority to Queue and Forwarding Class Mapping	166
	Table 73: Summary of Differences Between IEEE DCBX and DCBX Version 1.01	172
Chapter 3	QFabric-Specific CoS Overview	185
	Table 74: Default Fabric Forwarding Class Sets	187
	Table 75: Default Forwarding Class to Fabric Forwarding Class Set Mapping . .	188
	Table 76: Class Group Scheduling Properties and Membership	193
	Table 77: Lossless Priority (Forwarding Class) Support for QFX3500 and QFX3600 Node Devices When Fewer than Six Lossless Priorities Are Supported	196
	Table 78: show class-of-service forwarding-class-set Command Output Fields	197
	Table 79: Summary of Differences Between Fabric fc-sets and Node Device fc-sets	198
	Table 80: Class Group Default Scheduling Properties and Membership on Node Device Fabric Interfaces	201
	Table 81: Hierarchical Scheduler Architecture on Node Devices and Interconnect Devices	203
	Table 82: Class Group Default Scheduling Properties and Membership	205
	Table 83: Default Fabric FC-Set Scheduler Configuration	206
	Table 84: Lossless Priority (Forwarding Class) Support for QFX3500 and QFX3600 Node Devices When Fewer than Six Lossless Priorities Are Supported	212
Part 2	Configuration	
Chapter 5	Configuration Examples	219
	Table 85: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology	222
	Table 86: Components of the PFC for FCoE Traffic Configuration Topology . . .	243
	Table 87: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology	252
	Table 88: ba-ucast-classifier Loss Priority Assignments	275
	Table 89: BA-mcast-classifier Loss Priority Assignments	278
	Table 90: Forwarding-Class-to-Queue Example Configuration	285

	Table 91: Components of the Forwarding Class Sets Configuration Example . . .	287
	Table 92: Components of the Queue Scheduler Configuration Example	292
	Table 93: Components of the Queue Scheduler Priority Configuration Example	296
	Table 94: Components of the Minimum Guaranteed Output Bandwidth Configuration Example	302
	Table 95: Components of the Maximum Output Bandwidth Configuration Example	306
	Table 96: Components of the Configuration Topology for FCoE Traffic That Does Not Use Priority 3	310
	Table 97: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology	318
	Table 98: Components of the Two Lossless FCoE Priorities Configuration Topology	328
	Table 99: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology	342
	Table 100: Components of the IEEE 802.1p Priority Remapping Configuration Topology	358
	Table 101: Components of the Recommended Shared Buffer Configuration for Best-Effort Unicast Network Topologies	367
	Table 102: Components of the Recommended Shared Buffer Configuration for Best-Effort Network Topologies with Links Enabled for Ethernet PAUSE . .	373
	Table 103: Components of the Recommended Shared Buffer Configuration for Multicast Network Topologies	379
	Table 104: Components of the Recommended Shared Buffer Configuration for Lossless Network Topologies	385
	Table 105: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier)	391
	Table 106: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier)	391
	Table 107: Components of DCBX Application Protocol Exchange Configuration Topology	392
Chapter 7	Configuration Statements	439
	Table 108: Default Egress Shared Buffer Partitioning	448
	Table 109: Default Ingress Shared Buffer Partitioning	450
	Table 110: Default Output Queue Buffer Sizes	453
Part 3	Administration	
Chapter 8	Routine Monitoring	545
	Table 111: Summary of Key CoS Classifier Output Fields	545
	Table 112: Summary of Key CoS Forwarding Class Output Fields	546
	Table 113: Summary of Key CoS Interfaces Output Fields	547
	Table 114: Summary of Key CoS Rewrite Rule Output Fields	548
	Table 115: Summary of Key CoS Scheduler Maps Output Fields	549
	Table 116: Summary of Key CoS Value Alias Output Fields	550
Chapter 9	Operational Commands	553
	Table 117: show class-of-service Output Fields	554

Table 118: show class-of-service classifier Output Fields	558
Table 119: show class-of-service code-point-aliases Output Fields	560
Table 120: show class-of-service congestion-notification Output Fields	562
Table 121: show class-of-service drop-profile Output Fields	565
Table 122: show class-of-service forwarding-class Output Fields	568
Table 123: show class-of-service forwarding-class-set Output Fields	570
Table 124: show class-of-service forwarding-table classifier Output Fields	576
Table 125: show class-of-service forwarding-table classifier mapping Output Fields	578
Table 126: show class-of-service forwarding-table drop-profile Output Fields	580
Table 127: show class-of-service forwarding-table rewrite-rule Output Fields	582
Table 128: show class-of-service forwarding-table rewrite-rule mapping Output Fields	584
Table 129: show class-of-service forwarding-table scheduler-map Output Fields	585
Table 130: show class-of-service interface Output Fields	588
Table 131: show class-of-service multi-destination Output Fields	613
Table 132: show class-of-service rewrite-rule Output Fields	614
Table 133: show class-of-service scheduler-map Output Fields	616
Table 134: show class-of-service shared-buffer Output Fields	618
Table 135: show class-of-service traffic-control-profile Output Fields	620
Table 136: show dcbx output fields	624
Table 137: show dcbx neighbors Output Fields	625
Table 138: Layer 2 Overhead, Transmitted Packets/Bytes	648
Table 139: show interfaces queue Output Fields	651
Table 140: Byte Count by Interface Hardware	654

Part 4

Chapter 10

Troubleshooting

Troubleshooting Procedures	701
Table 141: Components of the Rate Shaping Troubleshooting Example	709

About the Documentation

- Documentation and Release Notes on page xxiii
- Supported Platforms on page xxiii
- Using the Examples in This Manual on page xxiii
- Documentation Conventions on page xxv
- Documentation Feedback on page xxvii
- Requesting Technical Support on page xxvii

Documentation and Release Notes

To obtain the most current version of all Juniper Networks[®] technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at <http://www.juniper.net/books>.

Supported Platforms

For the features described in this document, the following platforms are supported:

- QFabric System

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

If the example configuration does not start at the top level of the hierarchy, the example is a *snippet*. In this case, use the **load merge relative** command. These procedures are described in the following sections.

Merging a Full Example

To merge a full example, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration example into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following configuration to a file and name the file **ex-script.conf**. Copy the **ex-script.conf** file to the **/var/tmp** directory on your routing platform.

```
system {
  scripts {
    commit {
      file ex-script.xml;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
      family inet {
        address 10.0.0.1/24;
      }
    }
  }
}
```

2. Merge the contents of the file into your routing platform configuration by issuing the **load merge** configuration mode command:

```
[edit]
user@host# load merge /var/tmp/ex-script.conf
load complete
```

Merging a Snippet

To merge a snippet, follow these steps:

1. From the HTML or PDF version of the manual, copy a configuration snippet into a text file, save the file with a name, and copy the file to a directory on your routing platform.

For example, copy the following snippet to a file and name the file **ex-script-snippet.conf**. Copy the **ex-script-snippet.conf** file to the **/var/tmp** directory on your routing platform.

```
commit {
  file ex-script-snippet.xml; }
```

2. Move to the hierarchy level that is relevant for this snippet by issuing the following configuration mode command:


```
[edit]
user@host# edit system scripts
[edit system scripts]
```

3. Merge the contents of the file into your routing platform configuration by issuing the **load merge relative** configuration mode command:

```
[edit system scripts]
user@host# load merge relative /var/tmp/ex-script-snippet.conf
load complete
```

For more information about the **load** command, see the *CLI User Guide*.

Documentation Conventions

Table 1 on page xxv defines notice icons used in this guide.

Table 1: Notice Icons

Icon	Meaning	Description
	Informational note	Indicates important features or instructions.
	Caution	Indicates a situation that might result in loss of data or hardware damage.
	Warning	Alerts you to the risk of personal injury or death.
	Laser warning	Alerts you to the risk of personal injury from a laser.
	Tip	Indicates helpful information.
	Best practice	Alerts you to a recommended use or implementation.

Table 2 on page xxv defines the text and syntax conventions used in this guide.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the configure command: user@host> configure

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none">Introduces or emphasizes important new terms.Identifies guide names.Identifies RFC and Internet draft titles.	<ul style="list-style-type: none">A policy <i>term</i> is a named structure that defines match conditions and actions.<i>Junos OS CLI User Guide</i>RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none">To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level.The console port is labeled CONSOLE.
< > (angle brackets)	Encloses optional keywords or variables.	stub <default-metric <i>metric</i>>;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Encloses a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identifies a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">In the Logical Interfaces box, select All Interfaces.To cancel the configuration, click Cancel.

Table 2: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select Protocols>Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page at the Juniper Networks Technical Documentation site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>

- Download the latest versions of software and review release notes:
<http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications:
<http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum:
<http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://tools.juniper.net/SerialNumberEntitlementSearch/>

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.

PART 1

Overview

- [Software Features Overview on page 3](#)
- [CoS Overview on page 29](#)
- [QFabric-Specific CoS Overview on page 185](#)
- [Learn About Technology on page 215](#)

CHAPTER 1

Software Features Overview

- Overview of Junos OS CoS for the QFX Series on page 4
- Overview of CoS Upgrade Requirements (Junos OS Release 11.1 or 11.2 to a Later Release) on page 7
- Overview of CoS Upgrade Requirements to Junos OS Release 12.2 on page 8
- Overview of CoS Upgrade Requirements to Junos OS Release 12.3 (QFX3500 and QFX3600 Switches) or to Junos OS Release 13.1 (QFabric Systems) on page 10
- Overview of CoS Changes Introduced in Junos OS Release 11.3 on page 13
- Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 21
- Overview of Policers on page 23

Overview of Junos OS CoS for the QFX Series

When a network experiences congestion and delay, some packets must be dropped. Junos OS class of service (CoS) enables you to divide traffic into classes and set various levels of throughput and packet loss when congestion occurs. You have greater control over packet loss because you can configure rules tailored to your needs.

You can configure CoS features to provide multiple classes of service for different applications. CoS also allows you to rewrite the Differentiated Services code point (DSCP) or IEEE 802.1p code-point bits of packets leaving an interface, thus allowing you to tailor packets for the network requirements of the remote peers.

CoS provides multiple classes of service for different applications. You can configure multiple forwarding classes for transmitting packets, define which packets are placed into each output queue, schedule the transmission service level for each queue, and manage congestion using a random early detection (RED) algorithm.

In designing CoS applications, you must carefully consider your service needs, and you must thoroughly plan and design your CoS configuration to ensure consistency and interoperability across all platforms in a CoS domain.

Because QFX Series implements CoS in hardware rather than in software, you can experiment with and deploy CoS features without affecting packet forwarding and switching performance.



NOTE: CoS policies can be enabled or disabled on each switch interface. Also, each physical and logical interface on the switch can have associated custom CoS rules.

When you change or when you deactivate and then reactivate the class-of-service configuration, the system experiences packet drops because the system momentarily blocks traffic to change the mapping of incoming traffic to input queues.

This topic describes:

- [CoS Standards on page 4](#)
- [How Junos CoS Works on page 5](#)
- [Default CoS Behavior on page 6](#)

CoS Standards

The following RFCs define the standards for the QFX Series CoS capabilities:

- RFC 2474, *Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers*
- RFC 2597, *Assured Forwarding PHB Group*
- RFC 2598, *An Expedited Forwarding PHB*

- RFC 2698, *A Two Rate Three Color Marker*
- RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*

QFX Series also supports the following data center bridging (DCB) standards to provide the CoS (and other characteristics) Fibre Channel requires for transmitting storage traffic over an Ethernet network:

- IEEE 802.1Qbb, priority-based flow control (PFC)
- IEEE 802.1Qaz, enhanced transmission selection (ETS)
- IEEE 802.1AB (LLDP) extension called Data Center Bridging Capability Exchange Protocol (DCBX)

How Junos CoS Works


Junos CoS works by examining traffic entering at the edge of your network. The switch classifies traffic into defined service groups to provide the special treatment of traffic across the network. For example, you can send voice traffic across certain links and data traffic across other links. In addition, the data traffic streams can be serviced differently along the network path to ensure that higher-paying customers receive better service. As the traffic leaves the network at the far edge, you can reclassify the traffic to meet the policies of the targeted peer by rewriting the DSCP or IEEE 802.1 code-point bits.

To support CoS, you must configure each switch in the network. Generally, each switch examines the packets that enter it to determine their CoS settings. These settings dictate which packets are transmitted first to the next downstream switch. Switches at the edges of the network might be required to alter the CoS settings of the packets that enter the network to classify the packets into the appropriate service groups.

In [Figure 1 on page 6](#), Switch A is receiving traffic. As each packet enters, Switch A examines the packet's current CoS settings and classifies the traffic into one of the groupings defined on the switch. This definition allows Switch A to prioritize its resources for servicing the traffic streams it receives. Switch A might alter the CoS settings (forwarding class and loss priority) of the packets to better match the defined traffic groups.

When Switch B receives the packets, it examines the CoS settings, determines the appropriate traffic groups, and processes the packet according to those settings. It then transmits the packets to Switch C, which performs the same actions. Switch D also examines the packets and determines the appropriate groups. Because Switch D sits at the far end of the network, it can reclassify (rewrite) the CoS code-point bits of the packets before transmitting them.

Figure 1: Packet Flow Across the Network



Default CoS Behavior

If you do not configure CoS settings, the software performs some CoS functions to ensure that the system forwards traffic and protocol packets with minimum delay when the network is experiencing congestion. Some CoS settings, such as classifiers, are automatically applied to each logical interface that you configure. Other settings, such as rewrite rules, are applied only if you explicitly associate them with an interface.

Related Documentation

- [Overview of Policers on page 23](#)
- [Understanding Junos CoS Components on page 30](#)
- [Understanding CoS Packet Flow on page 34](#)
- [Understanding Default CoS Settings on page 37](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)

Overview of CoS Upgrade Requirements (Junos OS Release 11.1 or 11.2 to a Later Release)

Before you upgrade to Junos OS Release 11.3, you must deactivate the CoS configuration if the CoS configuration includes any of the following features:

- **excess-rate** option
- **strict-high** or **high** priority queues
- Any of the Junos OS Release 11.1 or 11.2 default multidestination forwarding classes



CAUTION: If your CoS configuration contains any of the features listed above and you attempt to upgrade from Junos OS Release 11.1 or 11.2 to a later version without first editing the configuration, the Junos OS might not restart.

Junos OS Release 11.3 and later for QFX Series no longer supports the **excess-rate** statement, the **strict** priority option, or the default multidestination forwarding classes used in Junos OS Release 11.1 and 11.2. In addition, Junos OS Release 11.3 introduces new restrictions on how to configure and use **strict-high** priority queues.

This topic does not describe how to perform the software upgrade procedure. It describes how to deactivate your CoS configuration, edit your CoS configuration, and reactivate your CoS configuration at the appropriate times.

Use the following procedure to upgrade safely from Junos OS Release 11.1 or 11.2 to a later release:

1. Deactivate the CoS configuration *before* you upgrade the software:

```
user@switch# deactivate class-of-service
```
2. Follow the upgrade procedure to Junos OS Release 11.3 or later software.
3. Make the following changes to the CoS configuration while the CoS configuration is still deactivated:
 - Remove the **excess-rate** statement from the CoS configuration if you have used it at the **[edit class-of-service schedulers]** or **[edit class-of-service traffic-control-profiles]** hierarchy level.
 - Remove the **strict-high** and **strict** priority queue configurations if you have used them at the **[edit class-of-service schedulers]** hierarchy level.
 - Remove the default multidestination forwarding classes (**mcast-be**, **mcast-af**, **mcast-ef**, and **mcast-nc**) if you have used them at the **[edit class-of-service schedulers]**, **[edit class-of-service rewrite-rules]**, **[edit class-of-service classifiers]**, **[edit class-of-service scheduler-maps]**, or **[edit class-of-service forwarding-class-sets]** hierarchy level. Alternatively, you can change the mapping of the multidestination traffic to use the new default multidestination forwarding class (**mcast**).
4. If desired, configure **strict-high** priority queues in accordance with the Junos OS Release 11.3 or later configuration rules, and map multidestination traffic to the default multidestination forwarding class (**mcast**).

5. Activate the CoS configuration:

```
user@switch# activate class-of-service
```

6. Commit the CoS configuration:

```
user@switch# commit
```



NOTE: If you configured the `transmit-rate` option for any queues under the `[edit class-of-service schedulers]` hierarchy level, if the rate is configured as an exact rate in Mbps, we recommend that you reconfigure the `transmit-rate` option as a percentage. This is because the scheduler converts exact rates to percentages, and when the exact rate is below 1 Gbps, some granularity may be lost in the conversion. You can avoid this potential issue by specifying the `transmit-rate` option as a percentage.

Related Documentation

- [Upgrading Software on QFX3500, QFX3600, and QFX5100 Switches](#)
- [Understanding CoS Classifiers on page 51](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Understanding CoS Traffic Control Profiles on page 98](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.2 on page 8](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.3 \(QFX3500 and QFX3600 Switches\) or to Junos OS Release 13.1 \(QFabric Systems\) on page 10](#)
- [Example: Configuring Unicast Classifiers on page 273](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)

Overview of CoS Upgrade Requirements to Junos OS Release 12.2

Before you upgrade to Junos OS Release 12.2, you might need to edit the class-of-service (CoS) configuration, because the way the QFX Series handles lossless forwarding classes has changed in Junos OS Release 12.2.

By default, the **fcoe** and **no-loss** forwarding classes are mapped to output queue 3 and output queue 4, respectively. These are the only two forwarding classes (and the only two queues) that support lossless transport.

In Junos OS Release 12.1 and earlier, explicitly setting the lossless **fcoe** and **no-loss** forwarding classes resulted in the same CoS behavior as using the default configuration. However, in Junos OS Release 12.2, the behavior when you explicitly configure the lossless forwarding classes differs from the behavior when you use the default forwarding classes.



NOTE: The default behavior differs from the explicit configuration behavior even if the explicit configuration is exactly the same as the default configuration.

- If you use the default forwarding class configuration for the lossless queues (the configuration does not include explicit setting of the **fcoe** or the **no-loss** forwarding classes), then the **fcoe** and **no-loss** queues behave as lossless queues.

If your CoS configuration does not explicitly configure the **fcoe** and **no-loss** forwarding classes, you can upgrade from Junos OS Release 12.1 to Junos OS Release 12.2, and the behavior of the two lossless queues remains the lossless.

- If your configuration includes statements that explicitly configure the **fcoe** or the **no-loss** forwarding class (using the **[set class-of-service forwarding-classes class class-name queue-num queue-number]** statement), after you upgrade to Junos OS Release 12.2, those queues do *not* receive lossless treatment and behave as lossy (**best-effort**) queues.

If your CoS configuration explicitly configures the **fcoe** and **no-loss** forwarding classes, to retain the lossless behavior of those queues, you need to remove the explicit configuration for these two forwarding classes from the CoS configuration *before* you upgrade.

If you upgrade to Junos OS Release 12.2 and the **fcoe** and **no-loss** forwarding classes are explicitly configured, then those two queues continue to be used, but the traffic is treated as lossy traffic, not lossless traffic. To make the queues for these two forwarding classes lossless, you must delete the explicit forwarding class configuration.



CAUTION: If you explicitly configured the **fcoe** or the **no-loss** forwarding class and you upgrade to Junos OS Release 12.2, the system does not return an upgrade error or a commit error, or a generate a syslog message, to notify you that these forwarding classes are no longer lossless. Traffic mapped to these forwarding classes is not treated as lossless traffic until you remove the explicit forwarding class configuration.

Before you upgrade, delete the **fcoe** and **no-loss** forwarding classes from the explicit configuration to preserve the lossless behavior of traffic mapped to these forwarding classes.

- To delete the explicit **fcoe** forwarding class configuration:

```
[edit]
user@switch# delete class-of-service forwarding-class class fcoe queue-num 3
user@switch# commit
```

- To delete the explicit **no-loss** forwarding class configuration:

```
[edit]
user@switch# delete class-of-service forwarding-class class no-loss queue-num 4
user@switch# commit
```



NOTE: If you try to delete these forwarding classes and they have not been explicitly configured on the system, the system returns the message **warning: statement not found**. This simply means that there is no explicit configuration to delete and does not change the lossless behavior of the **fcoe** and **no-loss** forwarding classes.

After you delete the explicit configuration for the **fcoe** and **no-loss** forwarding classes, traffic mapped to those forwarding classes retains its lossless behavior after the upgrade to Junos OS Release 12.2.

Related Documentation

- [Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\) on page 7](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.3 \(QFX3500 and QFX3600 Switches\) or to Junos OS Release 13.1 \(QFabric Systems\) on page 10](#)
- [Overview of CoS Changes Introduced in Junos OS Release 11.3 on page 13](#)
- [Upgrading Software on QFX3500, QFX3600, and QFX5100 Switches](#)
- [Understanding CoS Forwarding Classes on page 68](#)
- [Example: Configuring Forwarding Classes on page 283](#)

[Overview of CoS Upgrade Requirements to Junos OS Release 12.3 \(QFX3500 and QFX3600 Switches\) or to Junos OS Release 13.1 \(QFabric Systems\)](#)

Before you upgrade to Junos OS Release 12.3 (QFX3500 and QFX3600 switches) or to Junos OS Release 13.1 (QFabric systems), you might need to edit the class-of-service (CoS) configuration, because the way the QFX Series handles lossless forwarding classes has changed from earlier Junos OS releases. (Throughout this document, changes introduced on standalone switches in Junos OS Release 12.3 are introduced on QFabric systems in Junos OS Release 13.1 unless otherwise noted.)

- [Support for Six Lossless Forwarding Classes on page 10](#)
- [Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 12](#)
- [Strict-High Priority Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 12](#)

Support for Six Lossless Forwarding Classes

By default, the **fcoe** and **no-loss** forwarding classes are mapped to output queue 3 and output queue 4, respectively, and to IEEE 802.1p priority 3 (code point 011) and priority 4 (code point 100), respectively. These are the only two forwarding classes (and the only two queues) that support lossless transport in the default configuration.

If you use the default CoS configuration, you do not need to edit the CoS configuration after upgrading to Junos OS Release 12.3 (QFX3500 and QFX3600 switches) or to Junos OS Release 13.1 (QFabric system) because the default CoS configuration is backward-compatible.

Junos OS Release 12.3 increases the support for lossless forwarding classes (priorities) from two forwarding classes to six forwarding classes. To support configuring lossless forwarding classes, Junos OS Release 12.3 introduces a new option to forwarding class configuration: the *no-loss* packet drop attribute.



NOTE: The new *no-loss* packet drop attribute and the previously existing *no-loss* default forwarding class have the same name, but they are not the same. You can use the *no-loss* packet drop attribute on any unicast forwarding class.

If you explicitly configure any lossless forwarding class (including explicitly configuring the default *fcoe* and *no-loss* forwarding classes), you *must* specify the *no-loss* packet drop attribute to obtain lossless behavior. If you do not explicitly configure the *fcoe* and *no-loss* forwarding classes, those forwarding classes remain lossless.

The addition of the *no-loss* packet drop attribute to forwarding class configuration means that when you upgrade from an earlier release to Junos OS Release 12.3, the new software might not preserve the lossless forwarding class configuration of the *fcoe* and *no-loss* forwarding classes.

If you used the default forwarding class configuration for the *fcoe* and *no-loss* forwarding classes, the CoS configuration is backward-compatible. You do not have to do anything to preserve the lossless behavior of traffic that uses those forwarding classes when you upgrade to Junos OS Release 12.3. (This is because the default configuration of these two forwarding classes includes the *no-loss* packet drop attribute.)

However, if you explicitly configured the *fcoe* or the *no-loss* forwarding class by including the **set forwarding-classes class forwarding-class-name queue-num queue-number** at the **[edit class-of-service]** hierarchy level, then those forwarding classes are no longer lossless, they are lossy. In Junos OS Release 12.3 and later, you must include the *no-loss* packet drop attribute in any explicit forwarding class configuration to configure a lossless forwarding class.

For example, before Junos OS Release 12.3, the following explicit configuration resulted in a lossless forwarding class:

```
user@switch# set class-of-service forwarding-classes class fcoe queue-num 3
```

However, in Junos OS Release 12.3, this configuration is lossy because it does not include the *no-loss* packet drop attribute. To preserve lossless behavior, after upgrading to Junos OS Release 12.3, you need to add the *no-loss* drop attribute:

```
user@switch# set class-of-service forwarding-classes class fcoe queue-num 3 no-loss
```

Alternatively, you can delete the explicit configuration before you upgrade to Junos OS Release 12.3 so that the system uses the default forwarding class, which is lossless:

```
user@switch# delete class-of-service forwarding-classes class fcoe queue-num 3
```



NOTE: The explicit configuration of other forwarding classes does not affect the lossless (or lossy) state of the fcoe and no-loss forwarding classes, because only the fcoe and no-loss forwarding classes are lossless forwarding classes before Junos OS Release 12.3. For example, if you explicitly configured the best-effort forwarding class but you used the default fcoe and no-loss forwarding classes in Junos OS Release 12.2, then when you upgrade to Junos OS Release 12.3, the fcoe and no-loss forwarding classes are still lossless (and the best-effort forwarding class retains its explicit configuration).



NOTE: To achieve lossless behavior for the traffic belonging to any forwarding class, you must also enable PFC on the IEEE 802.1p priority mapped to the forwarding class and ensure that DCBX exchanges the protocol TLVs for the application with the connected peer.

Scheduling on QFabric System Node Device Fabric (fte) Ports

Junos OS Release 13.1 introduces the ability to configure scheduling on the fabric (fte) ports of QFabric system Node devices. In earlier Junos OS releases, Node device fabric port scheduling was done by default, with no user configuration.

In Junos OS Release 13.1, the default fabric port scheduler configuration is similar to the default scheduler configuration on access interfaces. Similar to the access port default configuration, the default fabric port scheduler supports the five default forwarding classes (best-effort, fcoe, no-loss, network-control, and mcast). If you configure any new forwarding classes, you must configure scheduling on the fabric ports to allocate bandwidth to those forwarding classes, just as you must configure scheduling on the access ports for user-defined forwarding classes.

Strict-High Priority Scheduling on QFabric System Node Device Fabric (fte) Ports

If a fabric interface handles strict-high priority traffic, you must define a separate fc-set (priority group) for strict-high priority traffic. Strict-high priority traffic cannot be mixed with traffic of other priorities in an fc-set. For example, you might choose to create different fc-sets for best effort, lossless, strict-high priority, and multidestination traffic.

Related Documentation

- [Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\) on page 7](#)
- [Overview of CoS Upgrade Requirements to Junos OS Release 12.2 on page 8](#)
- [Upgrading Software on QFX3500, QFX3600, and QFX5100 Switches](#)

Overview of CoS Changes Introduced in Junos OS Release 11.3

Junos OS Release 11.3 introduces many changes to class-of-service (CoS) functionality and to the CoS default values. This overview summarizes the changes, which other documents describe in detail.



NOTE: Some of the CoS changes are not backward compatible with Junos OS Releases 11.1 and 11.2. [“Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\)” on page 7](#) describes how to upgrade to Junos OS Release 11.3 if you have configured CoS on your QFX3500 switch.

This topic describes the following changes in CoS default values and behavior:

- [CoS Default Value Changes on page 13](#)
- [Queue Priority Configuration Changes on page 18](#)
- [Minimum Guaranteed Bandwidth \(Transmit Rate and Guaranteed Rate\) Changes on page 19](#)
- [Excess Rate Statement Disabled on page 19](#)
- [Queue Scheduling \(Low and Strict-High Priority Queues\) on page 20](#)
- [Multidestination Traffic Changes on page 20](#)

CoS Default Value Changes

The default values of the following CoS components have changed in Junos OS Release 11.3:

- [Default Forwarding Classes on page 13](#)
- [Default IEEE 802.1p Unicast Classifiers on page 15](#)
- [Default IEEE 802.1p Multidestination Classifiers on page 16](#)
- [Default Scheduler on page 17](#)

Default Forwarding Classes

In Junos OS Releases 11.1 and 11.2, there were eight default forwarding classes, four unicast default forwarding classes and four default multidestination (multicast, broadcast, and destination lookup fail) forwarding classes. [Table 3 on page 13](#) shows the old default forwarding classes and default queue mapping:

Table 3: Junos OS Release 11.1 and 11.2 Default Forwarding Classes and Queue Mapping

Default Forwarding Class	Description	Default Queue Mapping
best-effort (be)	Unicast best-effort traffic	0
no-loss	Unicast guaranteed delivery for TCP no-loss traffic	2

Table 3: Junos OS Release 11.1 and 11.2 Default Forwarding Classes and Queue Mapping (*continued*)

Default Forwarding Class	Description	Default Queue Mapping
fcoe	Unicast guaranteed delivery for FCoE traffic	3
network-control	Unicast network control traffic	7
multicast-best-effort (mcast-be)	Multidestination best-effort traffic	8
multicast-expedited-forwarding (mcast-ef)	Multidestination low-loss, low-latency traffic	9
multicast-assured-forwarding (mcast-af)	Multidestination assured forwarding traffic	10
multicast-network-control (mcast-nc)	Multidestination network control traffic	11

Junos OS Release 11.3 changes the default forwarding classes and queue mapping in the following ways:

- Instead of eight default forwarding classes, there are five default forwarding classes.
- The same four unicast default forwarding classes remain valid, but the default queue mapping of the no-loss forwarding class has changed from queue 2 to queue 4.
- There is now only one default multidestination forwarding class instead of four default multidestination forwarding classes. All multidestination traffic is assigned by default to the default multidestination forwarding class.



NOTE: The rest of the forwarding class characteristics remain the same as before. For example, the QFX Series still supports 12 forwarding classes and 12 output queues. You can still configure a total of eight unicast forwarding classes and four multidestination forwarding classes. The unicast queues are still queues 0 through 7 and the multidestination queues are still queues 8 through 11. Unicast traffic must be mapped to unicast queues, and multidestination traffic must be mapped to multidestination queues. The queue to which a forwarding class is mapped determines whether the forwarding class is unicast or multidestination.

[Table 4 on page 14](#) shows the default forwarding classes and queue mapping in Junos OS 11.3 and later:

Table 4: Junos OS Release 11.3 Default Forwarding Classes and Queue Mapping

Default Forwarding Class	Description	Default Queue Mapping
best-effort (be)	Best-effort traffic class	0
fcoe	Guaranteed delivery for FCoE traffic	3

Table 4: Junos OS Release 11.3 Default Forwarding Classes and Queue Mapping (*continued*)

Default Forwarding Class	Description	Default Queue Mapping
no-loss	Guaranteed delivery for TCP no-loss traffic	4
network-control (nc)	Network control traffic	7
mcast	Multicast traffic	8

Default IEEE 802.1p Unicast Classifiers

In Junos OS Release 11.1 and 11.2, there were default unicast classifiers only for best-effort and network-control traffic, as shown in [Table 5 on page 15](#):

Table 5: Junos OS Release 11.1 and 11.2 Default IEEE 802.1 Unicast Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	best-effort	low
be1 (001)	best-effort	low
ef (010)	best-effort	low
ef1 (011)	best-effort	low
af11 (100)	best-effort	low
af12 (101)	best-effort	low
nc1 (110)	network-control	low
nc2 (111)	network-control	low

Junos OS Release 11.3 introduces new default classifiers for FCoE and no-loss traffic, replacing the best-effort classifiers mapped to IEEE 802.1p code points 011 and 100, respectively, as shown in [Table 6 on page 15](#):

Table 6: Junos OS Release 11.3 Default IEEE 802.1 Unicast Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	best-effort	low
be1 (001)	best-effort	low
ef (010)	best-effort	low
ef1 (011)	fcoe	low

Table 6: Junos OS Release 11.3 Default IEEE 802.1 Unicast Classifiers (*continued*)

Code Point	Forwarding Class	Loss Priority
af11 (100)	no-loss	low
af12 (101)	best-effort	low
nc1 (110)	network-control	low
nc2 (111)	network-control	low

Default IEEE 802.1p Multidestination Classifiers

In Junos OS Release 11.1 and 11.2, there were default multidestination classifiers for best-effort and network-control traffic, as shown in [Table 7 on page 16](#):

Table 7: Junos OS Release 11.1 and 11.2 Default IEEE 802.1 Multidestination Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	mcast-be	low
be1 (001)	mcast-be	low
ef (010)	mcast-be	low
ef1 (011)	mcast-be	low
af11 (100)	mcast-be	low
af12 (101)	mcast-be	low
nc1 (110)	mcast-nc	low
nc2 (111)	mcast-nc	low

Junos OS Release 11.3 replaces the best-effort and network-control multidestination classifiers and maps all IEEE 802.1p code points to the new default multidestination forwarding class, as shown in [Table 8 on page 16](#):

Table 8: Junos OS Release 11.3 Default IEEE 802.1 Multidestination Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	mcast	low
be1 (001)	mcast	low
ef (010)	mcast	low

Table 8: Junos OS Release 11.3 Default IEEE 802.1 Multidestination Classifiers (*continued*)

Code Point	Forwarding Class	Loss Priority
ef1 (011)	mcast	low
af11 (100)	mcast	low
af12 (101)	mcast	low
nc1 (110)	mcast	low
nc2 (111)	mcast	low

Default Scheduler

In Junos OS Release 11.1 and 11.2, there were four default schedulers:

- Unicast best effort
- Unicast network control
- Multidestination best effort
- Multidestination network control

[Table 9 on page 17](#) shows the default scheduler configuration in Junos OS Release 11.1 and 11.2:

Table 9: Junos OS Release 11.1 and 11.2 Default Schedulers

Default Scheduler and Queue Number	Guaranteed Rate (Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Rate (Extra Bandwidth Sharing)	Priority
Best-effort scheduler (queue 0)	75%	None	25%	Low
Network-control scheduler (queue 7)	5%	None	25%	Low
Best-effort multidestination scheduler (queue 8)	15%	None	25%	Low
Network-control multidestination scheduler (queue 11)	5%	None	25%	Low

Junos OS Release 11.3 replaces the four old classifiers with five new classifiers:

- Unicast best effort
- FCoE
- No loss

- Unicast network control
- Multidestination

There are now four different default unicast classifiers to provide default CoS for lossless queues (FCoE and no-loss traffic). Because there is only one default multidestination forwarding class in Junos OS Release 11.3, there is only one default multidestination classifier for all multidestination traffic. Also, the excess rate default value is removed from the scheduler because the **excess-rate** statement is no longer supported, as described elsewhere in this document. [Table 10 on page 18](#) shows the default scheduler configuration in Junos OS Releases 11.3:

Table 10: Default Schedulers

Default Scheduler and Queue Number	Guaranteed Rate (Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Bandwidth Sharing	Priority
Best-effort scheduler (queue 0)	5%	None	5%	Low
FCoE scheduler (queue 3)	35%	None	35%	Low
No-loss scheduler (queue 4)	35%	None	35%	Low
Network-control scheduler (queue 7)	5%	None	5%	Low
Multidestination scheduler (queue 8)	20%	None	20%	Low



NOTE: The minimum guaranteed bandwidth rate also determines the amount of excess (extra) bandwidth that the queue can share. Extra bandwidth is allocated to queues in proportion to the minimum guaranteed bandwidth rate of each queue.

Queue Priority Configuration Changes

In Junos OS Release 11.1 and 11.2, you could configure strict-high priority queues with a guaranteed minimum bandwidth and configure forwarding class sets (priority groups) with a mix of low priority and strict-high priority queues. In Junos OS Release 11.3 and later, these configurations are invalid, and several other changes have also been implemented:

- Priority configuration in Junos OS Release 11.1 and 11.2 provided three priority levels: **strict-high**, **high**, and **low**. In Junos OS Release 11.3, the **high** priority option has been removed. Only the **strict-high** and **low** priority options are valid in Release 11.3.
- Minimum guaranteed bandwidth (transmit rate) is not allowed on strict-high priority queues. Minimum guaranteed bandwidth (guaranteed rate) is not allowed on forwarding class sets that contain strict-high priority queues.

- You cannot configure a multidestination queue as a strict-high priority queue. You cannot configure a queue as a strict-high priority queue if it belongs to the multidestination forwarding class set.
- Only one forwarding class set can contain strict-high priority queues. If you want to configure a strict-high priority queue, you must also configure a separate forwarding class set for the strict-high priority queue. A forwarding class set cannot contain a mixture of low priority and strict-high priority queues.

The rest of the queue priority characteristics remain the same as before. For example, you can configure only one queue as a strict-high priority queue.



NOTE: If you have configured strict-high or high priority queues in Junos OS Release 11.1 or 11.2, the changes in Release 11.3 are not backward compatible. Please read [“Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\)”](#) on page 7 before you upgrade to Release 11.3.

Minimum Guaranteed Bandwidth (Transmit Rate and Guaranteed Rate) Changes

The following restrictions have been placed on minimum guaranteed bandwidth configuration in Junos OS Release 11.3:

- You cannot configure a guaranteed minimum bandwidth (transmit rate) for strict-high priority queues.
- Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.
- You cannot configure a guaranteed minimum bandwidth (guaranteed rate) for forwarding class sets that include strict-high priority queues.
- For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.

Excess Rate Statement Disabled

The **excess-rate** statement has been disabled in Junos OS Release 11.3. Excess rate was used to specify the way extra bandwidth was shared among queues.

The **excess-rate** statement was used at the **[edit class-of-service schedulers]** hierarchy level for queue scheduling configuration and at the **[edit class-of-service traffic-control-profiles]** hierarchy level for forwarding class set scheduling configuration.

In Junos OS Release 11.3, extra bandwidth sharing among queues is proportional to the minimum guaranteed bandwidth (transmit rate) of the queue. Extra bandwidth sharing among forwarding class sets (priority groups) is proportional to the minimum guaranteed bandwidth (guaranteed rate) of the forwarding class set.



NOTE: If you have configured the excess-rate option in Junos OS Release 11.1 or 11.2, the changes in Release 11.3 are not backward compatible. Please read [“Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\)”](#) on page 7 before you upgrade to Release 11.3.

Queue Scheduling (Low and Strict-High Priority Queues)

In Junos OS Release 11.1 and 11.2, if you configured a guaranteed minimum bandwidth (transmit rate) for low-priority queues, the low-priority queues received their guaranteed minimum bandwidth from the same bandwidth pool as the strict-high priority queue, using round-robin scheduling. Until the minimum bandwidth requirements of all queues were met, the strict-high priority queue and low-priority queues that had a guaranteed minimum bandwidth were treated equally. After the minimum bandwidth requirements of all queues were met, the strict-high priority queue received as much of the leftover bandwidth as it needed. This meant that the only way to ensure that a strict-high priority queue received all of the bandwidth it needed was not to configure a guaranteed minimum bandwidth for other queues.

In Junos OS Release 11.3 and later, queue scheduling has changed so that queues receive bandwidth in the following sequence:

1. The strict-high priority queue receives all of the bandwidth it needs before any other queue is served. The strict-high priority queue can take the full port bandwidth if necessary and can starve other queues on the port.
2. The guaranteed minimum bandwidth (transmit rate) of low-priority queues is served until the minimum is met or the queues are empty.
3. All other low-priority queues and needs that exceed the minimum bandwidth are served.

Multidestination Traffic Changes

The changes to the default forwarding classes and classifiers affects multidestination traffic handling in Junos OS Release 11.3:

- The number of default multidestination forwarding classes has been reduced from four default multidestination forwarding classes in Junos OS Release 11.1 and 11.2 to one default multidestination in Release 11.3 (see [Table 4 on page 14](#)).
- The default classifier configuration for multidestination traffic has changed so that there is now one default classifier for all multidestination traffic (see [Table 8 on page 16](#)).
- By default, all IEEE 802.1p code points map to the default multidestination forwarding class.

- The default scheduler for multidestination traffic has changed so that there is now one default scheduler for all multidestination traffic (see [Table 10 on page 18](#)).
- You cannot configure multidestination queues as strict-high priority queues and you cannot include strict-high priority queues in a forwarding class set that contains multidestination queues.

**Related
Documentation**

- [Overview of CoS Upgrade Requirements \(Junos OS Release 11.1 or 11.2 to a Later Release\) on page 7](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 21](#)
- [Understanding Default CoS Settings on page 37](#)
- [Understanding Default CoS Scheduling and Classification on page 74](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)

Overview of CoS Changes Introduced in Junos OS Release 12.2

Junos OS Release 12.2 introduces some changes to class-of-service (CoS) functionality and to the CoS default values. This overview summarizes the changes, which other documents describe in detail.

This topic describes the following changes in CoS default values and behavior:

- [Lossless Forwarding Classes \(fcoe and no-loss\) on page 21](#)
- [Default MTU for Headroom Buffer Calculation for Lossless Forwarding Classes on page 22](#)
- [CoS for Layer 3 Physical Interfaces on page 22](#)
- [DSCP IPv6 Classifiers and Rewrite Rules on page 22](#)

Lossless Forwarding Classes (fcoe and no-loss)

The way the QFX Series handles lossless forwarding classes (the **fcoe** and **no-loss** forwarding classes) changes in Junos OS Release 12.2. In Junos OS Release 12.2 and in earlier releases, by default, the **fcoe** and **no-loss** forwarding classes are mapped to output queue 3 and output queue 4, respectively. These are the only two forwarding classes (and the only two queues) that support lossless transport.

In earlier releases, explicitly setting the lossless **fcoe** and **no-loss** forwarding classes resulted in the same CoS behavior as using the default configuration. However, in Junos OS Release 12.2, the behavior when you explicitly configure the lossless forwarding classes differs from the behavior when you use the default forwarding classes.



NOTE: The default behavior differs from the explicit configuration behavior even if the explicit configuration is exactly the same as the default configuration.

If you use the default forwarding class settings for the lossless queues (the configuration does not include explicit setting of the **fcoe** or the **no-loss** forwarding classes), then the **fcoe** and **no-loss** queues behave as lossless queues. When you upgrade to Junos OS Release 12.2, traffic assigned to the **fcoe** and **no-loss** queues continues to be treated as lossless traffic.

If your configuration explicitly sets the **fcoe** or the **no-loss** forwarding class (**set class-of-service forwarding-classes class class-name queue-num queue-number**), after you upgrade to Junos OS Release 12.2, those queues do *not* receive lossless treatment and behave as lossy (**best-effort**) queues. To retain lossless treatment of the **fcoe** and **no-loss** queues, delete the explicit lossless forwarding class configuration before you upgrade to Junos OS Release 12.2.



CAUTION: If you explicitly configured the **fcoe** or the **no-loss** forwarding class, and you upgrade to Junos OS Release 12.2, the system does not return an upgrade error or a commit error, or a generate a syslog message, to notify you that these forwarding classes are no longer lossless. Traffic mapped to these forwarding classes is not treated as lossless traffic until you remove the explicit forwarding class configuration.

Default MTU for Headroom Buffer Calculation for Lossless Forwarding Classes

The default maximum transmission unit (MTU) the system uses for buffer headroom calculation is 2500 bytes for traffic classified into the **fcoe** forwarding class or the **no-loss** forwarding class.

In Junos OS Release 12.2, the default MTU used for buffer headroom calculation for the **fcoe** and **no-loss** forwarding classes remains 2500 bytes. However, if the buffer is filled, in Junos OS Release 12.2 you might experience commit failures.

CoS for Layer 3 Physical Interfaces

Before Junos OS Release 12.2, the QFX Series supported only Layer 2 CoS. Junos OS Release 12.2 introduces CoS support for Layer 3 traffic at the physical interface level.

If a physical Layer 3 interface has at least one logical interface configured on it, you can configure Layer 3 CoS for the physical interface. The CoS configured on the physical interface applies to all of the logical Layer 3 interfaces on that physical interface. The system does not support Layer 3 CoS configuration on individual Layer 3 logical interfaces.

DSCP IPv6 Classifiers and Rewrite Rules

Junos OS Release 12.2 introduces support for DSCP IPv6 classifiers and rewrite rules. The existing DSCP IP default classifier is now also the DSCP IPv6 default classifier.

You can configure and apply DSCP IPv6 classifiers and DSCP IPv6 rewrite rules to Layer 2 logical interfaces and to Layer 3 physical interfaces.



NOTE: DSCP IPv6 classifiers are not supported for multidestination (multicast, broadcast, and destination lookup fail) traffic.

Related Documentation

- [Overview of CoS Upgrade Requirements to Junos OS Release 12.2 on page 8](#)
- [Overview of CoS Changes Introduced in Junos OS Release 11.3 on page 13](#)
- [Understanding Default CoS Settings on page 37](#)
- [Understanding CoS Classifiers on page 51](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57](#)

Overview of Policers

A switch polices traffic by limiting the input or output transmission rate of a class of traffic according to user-defined criteria. Policing (or rate-limiting) traffic allows you to control the maximum rate of traffic sent or received on an interface and to provide multiple priority levels or classes of service.


- [Policer Overview on page 23](#)
- [Policer Types on page 24](#)
- [Policer Actions on page 25](#)
- [Policer Colors on page 25](#)
- [Filter-Specific Policers on page 26](#)
- [Suggested Naming Convention for Policers on page 26](#)
- [Policer Counters on page 27](#)
- [Policer Algorithms on page 27](#)
- [How Many Policers are Supported? on page 27](#)
- [Policers can Limit Egress Firewall Filters on page 27](#)

Policer Overview

You use policers to apply limits to traffic flow and set consequences for packets that exceed these limits—usually applying a higher loss priority—so that if packets encounter downstream congestion, they can be discarded first. Policers apply only to unicast packets.

Policers provide two functions: metering and marking. A policer meters (measures) each packet against traffic rates and burst sizes that you configure. It then passes the packet and the metering result to the marker, which assigns a packet loss priority that corresponds to the metering result. [Figure 2 on page 24](#) illustrates this process.

Figure 2: Flow of Tricolor Marking Policer Operation



9017049

After you name and configure a policer, you use it by specifying it as an action in one or more firewall filters.

Policer Types

A switch supports three types of policers:

- **Single-rate two-color marker**—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit with a specified PLP or simply discard them.

You can specify this type of policer in an ingress or egress firewall.



NOTE: A two-color policer is most useful for metering traffic at the port (physical interface) level.

- **Single-rate three-color marker**—This type of policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on one rate—the configured committed information rate (CIR) as well as the committed burst size (CBS) and the excess burst size (EBS). The CIR specifies the average rate at which bits are admitted to the switch. The CBS specifies the usual burst size in bytes and the EBS specifies the maximum burst size in bytes. The EBS must be greater than or equal to the CBS, and neither can be 0.

You can specify this type of policer in an ingress or egress firewall.



NOTE: A single-rate three-color marker (TCM) is most useful when a service is structured according to packet length and not peak arrival rate.

- **Two-rate three-color marker**—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding per-hop-behavior classification system for a Differentiated Services environment. This type of policer meters traffic based on two rates—the CIR and peak information rate (PIR) along with their associated burst sizes, the CBS and peak burst size (PBS). The PIR specifies the maximum rate at which bits are admitted to the network and must be greater than or equal to the CIR.

You can specify this type of policer in an ingress or egress firewall.



NOTE: A two-rate three-color policer is most useful when a service is structured according to arrival rates and not necessarily packet length.

See [Table 11 on page 25](#) for information about how metering results are applied for each of these policer types.

Policer Actions

Policer actions are implicit or explicit and vary by policer type. *Implicit* means that Junos OS assigns the loss priority automatically. [Table 11 on page 25](#) describes the policer actions.

Table 11: Policer Actions

Policer	Marking	Implicit Action	Configurable Action
Single-rate two-color	Green (conforming)	Assign low loss priority	None
	Red (nonconforming)	None	Discard
Single-rate three-color	Green (conforming)	Assign low loss priority	None
	Yellow (above the CIR and CBS)	Assign medium-high loss priority	None
	Red (above the EBS)	Assign high loss priority	Discard
Two-rate three-color	Green (conforming)	Assign low loss priority	None
	Yellow (above the CIR and CBS)	Assign medium-high loss priority	None
	Red (above the PIR and PBS)	Assign high loss priority	Discard



NOTE: If you specify a policer in an egress firewall filter, the only supported action is discard.

Policer Colors

Single-rate and two-rate three-color policers can operate in two modes:

- **Color-blind**—In color-blind mode, the three-color policer assumes that all packets examined have not been previously marked or metered. In other words, the three-color policer is “blind” to any previous coloring a packet might have had.
- **Color-aware**—In color-aware mode, the three-color policer assumes that all packets examined have been previously marked or metered. In other words, the three-color policer is “aware” of the previous coloring a packet might have had. In color-aware mode, the three-color policer can increase the PLP of a packet but cannot decrease it. For example, if a color-aware three-color policer meters a packet with a medium PLP marking, it can raise the PLP level to high but cannot reduce the PLP level to low.

Filter-Specific Policers

You can configure policers to be filter-specific, which means that Junos OS creates only one policer instance regardless of how many times the policer is referenced. When you do this, rate limiting is applied in aggregate, so if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms, the total bandwidth allowed by the filter is 1 Gbps. However, the behavior of a filter-specific policer is affected by how the firewall filter terms that reference the policer are stored in TCAM. If you create a filter-specific policer and reference it in multiple firewall filter terms, the policer allows more traffic than expected if the terms are stored in different TCAM slices. For example, if you configure a policer to discard traffic that exceeds 1 Gbps and reference that policer in three different terms that are stored in three separate memory slices, the total bandwidth allowed by the filter is 3 Gbps, not 1 Gbps.

To prevent this unexpected behavior from occurring, use the information about TCAM slices presented in *Planning the Number of Firewall Filters to Create* to organize your configuration file so that all the firewall filter terms that reference a given filter-specific policer are stored in the same TCAM slice.

Suggested Naming Convention for Policers

We recommend that you use the naming convention ***policertypeTCM#-color type*** when configuring three-color policers and ***policer#*** when configuring two-color policers. TCM stands for three-color marker. Because policers can be numerous and must be applied correctly to work, a simple naming convention makes it easier to apply the policers properly. For example, the first single-rate, color-aware three-color policer configured would be named ***srTCM1-ca***. The second two-rate, color-blind three-color configured would be named ***trTCM2-cb***. The elements of this naming convention are explained below:

- sr (single-rate)
- tr (two-rate)
- TCM (tricolor marking)
- 1 or 2 (number of marker)
- ca (color-aware)
- cb (color-blind)

Policer Counters

Each policer that you configure includes an implicit counter that counts the number of packets that exceed the rate limits that are specified for the policer. If you use the same policer in multiple terms—either within the same filter or in different filters—the implicit counter counts all the packets that are policed in all of these terms. If you want to obtain separate packet counts for each term, use these options:

- Configure a unique policer for each term.
- Configure only one policer, but use a unique, explicit counter in each term.

Policer Algorithms

Policing uses the *token-bucket algorithm*, which enforces a limit on average bandwidth while allowing bursts up to a specified maximum value. It offers more flexibility than the *leaky bucket algorithm* in allowing a certain amount of bursty traffic before it starts discarding packets.

How Many Policers are Supported?

You can configure and commit the following numbers of policers on QFX3500 and QFX3600 devices when they are operating as standalone switches:

- Two-color policers used in ingress firewall filters: 767
- Three-color policers used in ingress firewall filters: 767
- Two-color policers used in egress firewall filters: 1022
- Three-color policers used in egress firewall filters: 512

Policers can Limit Egress Firewall Filters

The number of egress policers that you configure can affect the total number of allowed egress firewall filters. Every policer has two implicit counters that consume two entries in a 1024-entry TCAM that is used for counters, including counters that are configured as action modifiers in firewall filter terms. (Policers consume two entries because one is used for green packets and one is used for nongreen packets regardless of policer type.) If the TCAM becomes full, you cannot commit any more egress firewall filters that have terms with counters. For example, if you configure and commit 512 egress policers (two-color, three-color, or a combination of both policer types), all of the memory entries for counters are used up. If later in your configuration file you insert additional egress firewall filters with terms that also include counters, *none* of the terms in those filters are committed because there is no available memory space for the counters.

Here are some additional examples:

- Assume that you configure egress filters that include a total of 512 policers and no counters. Later in your configuration file you include another egress filter with 10 terms, 1 of which has a counter action modifier. None of the terms in this filter are committed because there is not enough TCAM space for the counter.

- Assume that you configure egress filters that include a total of 500 policers, so 1000 TCAM entries are occupied. Later in your configuration file you include the following two egress filters:
 - Filter A with 20 terms and 20 counters. All the terms in this filter are committed because there is enough TCAM space for all the counters.
 - Filter B comes after Filter A and has five terms and five counters. *None* of the terms in this filter are committed because there is not enough memory space for *all* the counters. (Five TCAM entries are required but only four are available.)

You can prevent this problem by ensuring that egress firewall filter terms with counter actions are placed earlier in your configuration file than terms that include policers. In this circumstance, Junos OS commits policers even if there is not enough TCAM space for the implicit counters. For example, assume the following:

- You have 1024 egress firewall filter terms with counter actions.
- Later in your configuration file you have an egress filter with 10 terms. None of the terms have counters but one has a policer action modifier.

You can successfully commit the filter with 10 terms even though there is not enough TCAM space for the implicit counters of the policer. The policer is committed without the counters.

**Related
Documentation**

- *Understanding Color-Blind Mode for Single-Rate Tricolor Marking*
- *Understanding Color-Blind Mode for Two-Rate Tricolor Marking*
- *Understanding Color-Aware Mode for Single-Rate Tricolor Marking*
- *Understanding Color-Aware Mode for Two-Rate Tricolor Marking*
- *Configuring Two-Color and Three-Color Policers to Control Traffic Rates*

CHAPTER 2

CoS Overview

- [Understanding Junos CoS Components on page 30](#)
- [Understanding CoS Packet Flow on page 34](#)
- [CoS Inputs and Outputs Overview on page 36](#)
- [Understanding Default CoS Settings on page 37](#)
- [Understanding Host Inbound Traffic Classification on page 46](#)
- [Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 47](#)
- [Understanding CoS Code-Point Aliases on page 49](#)
- [Understanding CoS Classifiers on page 51](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57](#)
- [Understanding CoS Forwarding Classes on page 68](#)
- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 73](#)
- [Understanding Default CoS Scheduling and Classification on page 74](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Understanding CoS Priority Group Scheduling on page 95](#)
- [Understanding CoS Traffic Control Profiles on page 98](#)
- [Understanding CoS Priority Group and Queue Guaranteed Rates \(Minimum Bandwidth\) on page 99](#)
- [Understanding CoS Priority Group Shaping and Queue Shaping \(Maximum Bandwidth\) on page 102](#)
- [Understanding CoS Scheduling Behavior and Configuration Considerations on page 104](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 128](#)
- [Understanding CoS Buffer Configuration on page 131](#)
- [Understanding CoS WRED Drop Profiles on page 149](#)
- [Understanding CoS Rewrite Rules on page 154](#)
- [Understanding DCB Features and Requirements on page 156](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)

- [Understanding DCBX on page 170](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 179](#)

Understanding Junos CoS Components

This topic describes the Junos operating system (OS) class-of-service (CoS) components for the QFX Series:

- [Code-Point Aliases on page 30](#)
- [Policers on page 30](#)
- [Classifiers on page 30](#)
- [Forwarding Classes on page 31](#)
- [Forwarding Class Sets on page 31](#)
- [Flow Control \(Ethernet PAUSE, PFC, and ECN\) on page 31](#)
- [WRED Profiles on page 32](#)
- [Schedulers on page 33](#)
- [Rewrite Rules on page 33](#)

Code-Point Aliases

A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers and rewrite rules.

Policers

Policers limit traffic of a certain class to a specified bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a different forwarding class, a different loss priority, or both. You define policers with filters that you can associate with input interfaces.

Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. In Junos OS, *classifiers* associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class. Junos OS supports two general types of classifiers:

- Behavior aggregate (BA) or CoS value traffic classifiers—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value or IEEE 802.1p value.
- Multifield traffic classifiers—Examine multiple fields in the packet, such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

You can create unicast classifiers for unicast traffic and multideestination classifiers for multicast, broadcast, and destination lookup fail traffic. You cannot assign unicast traffic and multideestination traffic to the same classifier.

You can apply unicast classifiers to one or more interfaces. Multideestination classifiers apply to all of the switch interfaces and cannot be applied to individual interfaces.

Forwarding Classes

Forwarding classes group packets for transmission and CoS. You assign each packet to an output queue based on the packet's forwarding class. Forwarding classes affect the forwarding, scheduling, and rewrite marking policies applied to packets as they transit the switch.

The switch provides five default forwarding classes:

- fcoe—Fibre Channel over Ethernet traffic
- no-loss—Lossless traffic
- be—Best-effort traffic
- nc—Network control traffic
- mcast—Multicast traffic

The switch supports a total of 12 forwarding classes (8 unicast forwarding classes and 4 multicast forwarding classes), which provide flexibility in classifying traffic.

Forwarding Class Sets

You can group forwarding classes (output queues) into *forwarding class sets* in order to apply CoS to groups of traffic that require similar treatment. Forwarding class sets map traffic into priority groups to support enhanced transmission selection (ETS, described in IEEE 802.1Qaz).

You can configure up to three unicast forwarding class sets and one multicast forwarding class set. For example, you can configure different forwarding class sets to apply CoS to unicast groups of local area network (LAN) traffic, storage area network (SAN) traffic, and high-performance computing (HPC) traffic, and configure another group for multicast traffic.

Within each forwarding class set, you can configure special CoS treatment for the traffic mapped to each individual queue. This provides the ability to configure CoS in a two-tier hierarchical manner. At the forwarding class set tier, you configure CoS for groups of traffic using a *traffic control profile*. At the queue tier, you configure CoS for individual output queues within a forwarding class set using a *scheduler* that you map to a queue (forwarding class) using a *scheduler map*.

Flow Control (Ethernet PAUSE, PFC, and ECN)

Ethernet PAUSE (described in IEEE 802.3X) is a link-level flow control mechanism. During periods of network congestion, Ethernet PAUSE stops all traffic on a full-duplex Ethernet link for a period of time specified in the PAUSE message.

Priority-based flow control (PFC, described in IEEE 802.1Qbb) is part of the IEEE data center bridging (DCB) specifications for creating a lossless Ethernet environment to transport loss-sensitive flows such as Fibre Channel over Ethernet (FCoE) traffic.

PFC is a link-level flow control mechanism similar to Ethernet PAUSE. However, Ethernet PAUSE stops all traffic on a link for a period of time. PFC decouples the pause function from the physical link and divides the traffic on the link into eight priorities (3-bit IEEE 802.1p code points). You can think of the eight priorities as eight “lanes” of traffic. You can apply pause selectively to the traffic on any priority without pausing the traffic on other priorities on the same link.

The granularity that PFC provides allows you to configure different levels of CoS for different types of traffic on the link. You can create lossless lanes for traffic such as FCoE, LAN backup, or management, while using standard frame-drop methods of congestion management for IP traffic on the same link.



NOTE: If you transport FCoE traffic, you must enable PFC on the priority assigned to FCoE traffic (usually IEEE 802.1p code point 011 on interfaces that carry FCoE traffic).

Explicit congestion notification (ECN) enables end-to-end congestion notification between two endpoints on TCP/IP based networks. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality. ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets. RFC 3168, *The Addition of Explicit Congestion Notification (ECN) to IP*, defines ECN.

WRED Profiles

A WRED (weighted random early detection) profile (drop profile) defines parameters that enable the network to drop packets during periods of congestion. A drop profile defines the conditions under which packets of different loss priorities drop, by determining the probability of dropping a packet for each loss priority when output queues become congested. Drop profiles essentially set a value for a level of queue fullness—when the queue fills to the level of the queue fullness value, packets drop.

You can associate different drop profiles with different loss priorities to set the probability of dropping packets. You can apply a drop profile for each loss priority to a forwarding class (output queue) by applying a drop profile to a scheduler, and then mapping the scheduler to a forwarding class using a scheduler map. When the queue mapped to the forwarding class experiences congestion, the drop profile determines the level of packet drop for traffic of each loss priority in that queue.

Loss priority affects the scheduling of a packet without affecting the packet's relative ordering. Typically you mark packets exceeding a particular service level with a high loss priority.

Schedulers

Each switch interface has multiple queues assigned to store packets. The switch determines which queue to service based on a particular method of scheduling. This process often involves determining the sequence in which different types of packets should be transmitted.

You can define the priority (**priority**), minimum bandwidth (**transmit-rate**), maximum bandwidth (**shaping-rate**), and WRED profiles to be applied to a particular queue for packet transmission. Extra bandwidth is shared among queues in proportion to the minimum guaranteed bandwidth of each queue.

A scheduler map associates a specified forwarding class with a scheduler configuration. You can associate up to four user-defined scheduler maps with the interfaces.

Rewrite Rules

A *rewrite rule* sets the appropriate CoS bits in the outgoing packet. This allows the next downstream device to classify the packet into the appropriate service group. Rewriting (marking) outbound packets is useful when the switch is at the border of a network and must change the CoS values to meet the policies of the targeted peer.



NOTE: Ingress firewall filters can also rewrite forwarding class and loss priority values.

Related Documentation

- [Understanding CoS Packet Flow on page 34](#)
- [Understanding CoS Code-Point Aliases on page 49](#)
- [Overview of Policers on page 23](#)
- [Understanding CoS Classifiers on page 51](#)
- [Understanding CoS Forwarding Classes on page 68](#)
- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 73](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)
- [Understanding CoS Explicit Congestion Notification](#)
- [Understanding CoS WRED Drop Profiles on page 149](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Understanding CoS Rewrite Rules on page 154](#)
- [Understanding DCB Features and Requirements on page 156](#)

Understanding CoS Packet Flow

When a packet traverses a QFX Series Node device, the Node device provides the appropriate level of service to the packet using either default class of service (CoS) settings or CoS settings that you configure. On ingress ports, the Node device classifies packets into appropriate forwarding classes and assigns a loss priority to the packets. On egress ports, the Node device applies packet scheduling and (if you have configured them) rewrite rules to re-mark packets.

On the QFX Series, you can configure CoS on Layer 2 logical interfaces, and you can configure CoS on Layer 3 physical interfaces if you have defined at least one logical interface on the Layer 3 physical interface. You cannot configure CoS on Layer 2 physical interfaces and Layer 3 logical interfaces.

For Layer 2 traffic, either use the default CoS settings or configure CoS on each logical interface. You can apply different CoS settings to different Layer 2 logical interfaces.

For Layer 3 traffic, either use the default CoS settings or configure CoS on the physical interface (not on the logical unit). The QFX Series uses the CoS applied on the physical Layer 3 interface for all logical Layer 3 interfaces configured on the physical Layer 3 interface.

The QFX Series applies to CoS to packets as they flow through the system:

- An interface has one or more classifiers of different types applied to it (configure this at the **[edit class-of-service interfaces]** hierarchy level). The classifier types are based on the portion of the incoming packet that the classifier examines (IEEE 802.1p code point bits or DSCP code point bits).
- When a packet enters an ingress port, the classifier assigns the packet to a forwarding class and a loss priority based on the code point bits of the packet (configure this at the **[edit class-of-service classifiers]** hierarchy level).
- The QFX Series assigns each forwarding class to an output queue (configure this at the **[edit class-of-service forwarding-classes]** hierarchy level).
- Input (and output) policers meter traffic and can change the forwarding class and loss priority if a traffic flow exceeds its service level.
- A scheduler map is applied to each interface. When a packet exits an egress port, the scheduler map controls how it is treated (configure this at the **[edit class-of-service interfaces]** hierarchy level). A scheduler map assigns schedulers to forwarding classes (configure this at the **[edit class-of-service scheduler-maps]** hierarchy level).
- A scheduler defines how traffic is treated at the egress interface output queue (configure this at the **[edit class-of-service schedulers]** hierarchy level). You control the transmit rate, shaping rate, priority, and drop profile of each forwarding class by mapping schedulers to forwarding classes in scheduler maps, then applying scheduler maps to interfaces.

- A drop-profile defines how aggressively to drop packets that are mapped to a particular scheduler (configure this at the **[edit class-of-service drop-profiles]** hierarchy level).
- A rewrite rule takes effect as the packet leaves an interface that has a rewrite rule configured (configure this at the **[edit class-of-service rewrite-rules]** hierarchy level). The rewrite rule writes information to the packet (for example, a rewrite rule can re-mark the code point bits of outgoing traffic) according to the forwarding class and loss priority of the packet.

Figure 3 on page 35 is a high-level flow diagram of how packets from various sources enter QFX Series interfaces, are classified at the ingress, and then scheduled (provided bandwidth) at the egress queues.

Figure 3: CoS Classifier, Queues, and Scheduler





Figure 4 on page 35 shows the packet flow through the CoS components that you can configure.

Figure 4: Packet Flow Through Configurable CoS Components



The middle box ("Forwarding Class and Loss Priority") represents two values that you can use on ingress and egress interfaces. The system uses these values for classifying traffic on ingress interfaces and for rewrite rule re-marking on egress interfaces. Each outer box represents a process component. The components in the top row apply to incoming packets. The components in the bottom row apply to outgoing packets.

The solid-line arrows show the direction of packet flow from ingress to egress. The dotted-line arrows show inputs and outputs or show settings and actions based on those settings.

For example, the BA classifier sets the forwarding class and loss priority of incoming packets, so the forwarding class and loss priority are outputs of the classifier and the arrow points away from the classifier. The scheduler receives the forwarding class and loss priority settings, and queues the outgoing packets based on those settings, so the arrow points toward the scheduler.

Related Documentation

- [Understanding CoS Classifiers on page 51](#)
- [Overview of Policers on page 23](#)
- [Understanding CoS Forwarding Classes on page 68](#)
- [Understanding Default CoS Scheduling and Classification on page 74](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Understanding CoS Priority Group Scheduling on page 95](#)
- [Understanding CoS WRED Drop Profiles on page 149](#)
- [Understanding CoS Rewrite Rules on page 154](#)

CoS Inputs and Outputs Overview

Some CoS components map one set of values to another set of values. Each mapping contains one or more inputs and one or more outputs. When you configure a mapping, you set the outputs for a given set of inputs, as shown in [Table 12 on page 36](#).

Table 12: CoS Mappings—Inputs and Outputs

CoS Mappings	Inputs	Outputs	Comments
classifiers	code-points	forwarding-class , loss-priority	The map sets the forwarding class and packet loss priority (PLP) for a specific set of code points.
drop-profile-map	loss-priority , protocol	drop-profile	The map sets the drop profile for a specific PLP and protocol type.
rewrite-rules	loss-priority , forwarding-class	code-points	The map sets the code points for a specific forwarding class and PLP.
rewrite-value (Fibre Channel Interfaces)	forwarding-class	code-point	The map sets the code point for the forwarding class specified in the fixed classifier attached to the native Fibre Channel (NP_Port) interface.

Related Documentation

- [Understanding CoS Packet Flow on page 34](#)
- [Understanding CoS Classifiers on page 51](#)
- [Example: Configuring Unicast Classifiers on page 273](#)

- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 277](#)
- [Example: Configuring Forwarding Classes on page 283](#)
- [Configuring CoS Drop Profile Maps on page 408](#)
- [Defining CoS Rewrite Rules on page 427](#)

Understanding Default CoS Settings

If you do not configure CoS settings on the QFX Series, Junos OS performs some CoS functions to ensure that traffic and protocol packets are forwarded with minimum delay when the network experiences congestion. Some default mappings are automatically applied to each logical interface that you configure.

You can display default CoS settings by issuing the **show class-of-service** operational mode command.

This topic describes the default configurations for the following CoS components:

- [Default Forwarding Classes and Queue Mapping on page 37](#)
- [Default Forwarding Class Sets \(Priority Groups\) on page 38](#)
- [Default Code-Point Aliases on page 38](#)
- [Default Classifiers on page 40](#)
- [Default Rewrite Rules on page 43](#)
- [Default Drop Profile on page 43](#)
- [Default Schedulers on page 43](#)
- [Default Scheduler Maps on page 45](#)
- [Default Shared Buffer Configuration on page 45](#)

Default Forwarding Classes and Queue Mapping

[Table 13 on page 37](#) shows the default mapping of the default forwarding classes to queues and packet drop attribute.

Table 13: Default Forwarding Classes and Queue Mapping

Default Forwarding Class	Description	Default Queue Mapping	Packet Drop Attribute
best-effort (be)	Best-effort traffic class (priority 0, IEEE 802.1p code point 000)	0	drop
fcoe	Guaranteed delivery for FCoE traffic (priority 3, IEEE 802.1p code point 011)	3	no-loss

Table 13: Default Forwarding Classes and Queue Mapping (*continued*)

Default Forwarding Class	Description	Default Queue Mapping	Packet Drop Attribute
no-loss	Guaranteed delivery for TCP no-loss traffic (priority 4, IEEE 802.1p code point 100)	4	no-loss
network-control (nc)	Network control traffic (priority 7, IEEE 802.1p code point 111)	7	drop
mcast	Multidestination traffic	8	drop

NOTE: You cannot configure multidestination forwarding classes as no-loss (lossless) traffic classes.

Default Forwarding Class Sets (Priority Groups)

If you do not explicitly configure forwarding class sets, the system automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The system assigns 100 percent of the port output bandwidth to the default forwarding class set.

Ingress traffic is classified based on the default classifier settings. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default scheduler settings. Forwarding classes that are not part of the default scheduler receive no bandwidth.

The default forwarding class set is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange (DCBX) protocol advertisement.

Default Code-Point Aliases

Table 14 on page 38 shows the default mapping of code-point aliases to IEEE code points.

Table 14: Default IEEE 802.1 Code-Point Aliases

CoS Value Types	Mapping
be	000
be1	001
ef	010
ef1	011
af11	100

Table 14: Default IEEE 802.1 Code-Point Aliases (*continued*)

CoS Value Types	Mapping
af12	101
nc1	110
nc2	111

Table 15 on page 39 shows the default mapping of code-point aliases to DSCP and DCSP IPv6 code points.

Table 15: Default DSCP and DCSP IPv6 Code-Point Aliases

CoS Value Types	Mapping
ef	101110
af11	001010
af12	001100
af13	001110
af21	010010
af22	010100
af23	010110
af31	011010
af32	011100
af33	011110
af41	100010
af42	100100
af43	100110
be	000000
cs1	001000
cs2	010000
cs3	011000

Table 15: Default DSCP and DCSP IPv6 Code-Point Aliases (*continued*)

CoS Value Types	Mapping
cs4	100000
cs5	101000
nc1	110000
nc2	111000

Default Classifiers

The QFX Series applies default unicast IEEE 802.1, unicast DSCP, and multdestination classifiers to each interface that does not have explicitly configured classifiers. If you explicitly configure one type of classifier but not other types of classifiers, the system uses only the configured classifier and does not use default classifiers for other types of traffic. There are two different default unicast IEEE 802.1 classifiers, a trusted classifier for ports that are in trunk mode or tagged-access mode, and an untrusted classifier for ports that are in access mode.

[Table 16 on page 40](#) shows the default mapping of IEEE 802.1 code-point values to unicast forwarding classes and loss priorities for ports in trunk mode or tagged-access mode.

Table 16: Default IEEE 802.1 Unicast Classifiers for Ports in Trunk Mode or Tagged Access Mode (Trusted Classifier)

Code Point	Forwarding Class	Loss Priority
be (000)	best-effort	low
be1 (001)	best-effort	low
ef (010)	best-effort	low
ef1 (011)	fcoe	low
af11 (100)	no-loss	low
af12 (101)	best-effort	low
nc1 (110)	network-control	low
nc2 (111)	network-control	low

[Table 17 on page 41](#) shows the default mapping of IEEE 802.1p code-point values to unicast forwarding classes and loss priorities for ports in access mode (all incoming traffic is mapped to best-effort forwarding classes).

Table 17: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier)

Code Point	Forwarding Class	Loss Priority
000	best-effort	low
001	best-effort	low
010	best-effort	low
011	best-effort	low
100	best-effort	low
101	best-effort	low
110	best-effort	low
111	best-effort	low

[Table 18 on page 41](#) shows the default mapping of IEEE 802.1 code-point values to multidestination (multicast, broadcast, and destination lookup fail traffic) forwarding classes and loss priorities.

Table 18: Default IEEE 802.1 Multidestination Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	mcast	low
be1 (001)	mcast	low
ef (010)	mcast	low
ef1 (011)	mcast	low
af11 (100)	mcast	low
af12 (101)	mcast	low
nc1 (110)	mcast	low
nc2 (111)	mcast	low

[Table 19 on page 42](#) shows the default mapping of DSCP code-point values to unicast forwarding classes and loss priorities for DSCP IP and DCSP IPv6.

Table 19: Default DSCP IP and IPv6 Unicast Classifiers

Code Point	Forwarding Class	Loss Priority
ef (101110)	best-effort	low
af11 (001010)	best-effort	low
af12 (001100)	best-effort	low
af13 (001110)	best-effort	low
af21 (010010)	best-effort	low
af22 (010100)	best-effort	low
af23 (010110)	best-effort	low
af31 (011010)	best-effort	low
af32 (011100)	best-effort	low
af33 (011110)	best-effort	low
af41 (100010)	best-effort	low
af42 (100100)	best-effort	low
af43 (100110)	best-effort	low
be (000000)	best-effort	low
cs1 (001000)	best-effort	low
cs2 (010000)	best-effort	low
cs3 (011000)	best-effort	low
cs4 (100000)	best-effort	low
cs5 (101000)	best-effort	low
nc1 (110000)	network-control	low
nc2 (111000)	network-control	low



NOTE: There are no default DSCP IP or IPv6 classifiers for multidestination traffic. DSCP IPv6 classifiers are not supported for multidestination traffic.

Default Rewrite Rules

There are no default rewrite rules. If you do not explicitly configure rewrite rules, the switch does not reclassify egress traffic.

Default Drop Profile

Table 20 on page 43 shows the default drop profile configuration.

Table 20: Default Drop Profile

Fill Level	Drop Probability
100	100

Default Schedulers

Table 21 on page 43 shows the default scheduler configuration.

Table 21: Default Schedulers

Default Scheduler and Queue Number	Transmit Rate (Guaranteed Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Bandwidth Sharing	Priority	Buffer Size
Best-effort scheduler (queue 0)	5%	None	5%	low	5%
FCoE scheduler (queue 3)	35%	None	35%	low	35%
No-loss scheduler (queue 4)	35%	None	35%	low	35%
Network-control scheduler (queue 7)	5%	None	5%	low	5%
Multidestination scheduler (queue 8)	20%	None	20%	low	20%



NOTE: The minimum guaranteed bandwidth (transmit rate) also determines the amount of excess (extra) bandwidth that the queue can share. Extra bandwidth is allocated to queues in proportion to the transmit rate of each queue.

By default, only the five default schedulers shown in Table 21 on page 43 have traffic mapped to them. Only the queues associated with the default schedulers receive default bandwidth, based on the default scheduler transmit rate. (You can configure schedulers and forwarding classes to allocate bandwidth to other queues or to change the default bandwidth of a default queue.) In addition, multidestination queue 11 receives enough bandwidth from the default multidestination scheduler to handle CPU-generated multidestination traffic.

Default hierarchical scheduling divides the total port bandwidth between two groups of traffic: unicast traffic and multidestination traffic. By default, unicast traffic consists of queue 0 (**best-effort** forwarding class), queue 3 (**fcoe** forwarding class), queue 4 (**no-loss** forwarding class), and queue 7 (**network-control** forwarding class). Unicast traffic receives and shares a total of 80 percent of the port bandwidth. By default, multidestination traffic (**mcast** queue 8) receives a total of 20 percent of the port bandwidth. So on a 10-Gigabit port, unicast traffic receives 8-Gbps of bandwidth and multidestination traffic receives 2-Gbps of bandwidth.



NOTE: Multidestination queue 11 also receives a small amount of default bandwidth from the multidestination scheduler. CPU-generated multidestination traffic uses queue 11, so you might see a small number of packets egress from queue 11. In addition, in the unlikely case that firewall filter match conditions map multidestination traffic to a unicast forwarding class, that traffic uses queue 11.

Default scheduling uses weighted round-robin (WRR) scheduling. Each queue receives a portion (weight) of the total available interface bandwidth. The scheduling weight is based on the transmit rate of the default scheduler for that queue. For example, queue 7 receives a default scheduling weight of 5 percent of the available bandwidth, and queue 4 receives a default scheduling weight of 35 percent of the available bandwidth. Queues are mapped to forwarding classes, so forwarding classes receive the default bandwidth for the queues to which they are mapped.

You should explicitly map traffic to non-default (unconfigured) queues and create schedulers to allocate bandwidth to those queues if you want to use them to forward traffic. By default, unicast queues 1, 2, 5, and 6 are unconfigured, and multidestination queues 9, 10, and 11 are unconfigured. Unconfigured queues have a default scheduling weight of 1 so that they can receive a small amount of bandwidth in case they need to forward traffic. (However, queue 11 can use more of the default multidestination scheduler bandwidth if necessary to handle CPU-generated multidestination traffic.)



NOTE: All four multidestination queues have a scheduling weight of 1. Because by default multidestination traffic goes to queue 8, queue 8 receives almost all of the multidestination bandwidth. (There is no traffic on queue 9 and queue 10, and very little traffic on queue 11, so there is almost no competition for multidestination bandwidth.)

However, if you explicitly configure queue 9, 10, or 11 (by mapping code points to the unconfigured multidestination forwarding classes using the multidestination classifier), the explicitly configured queues share the multidestination scheduler bandwidth equally with default queue 8, because all of the queues have the same scheduling weight (1). To ensure that multidestination bandwidth is allocated to each queue properly and that the bandwidth allocation to the default queue (8) is not reduced too much, we strongly recommend that you configure a scheduler if you explicitly classify traffic into queue 9, 10, or 11.

If you map traffic to an unconfigured queue, the queue receives only the amount of group bandwidth proportional to its default weight (1). The actual amount of bandwidth an unconfigured queue receives depends on how much bandwidth the other queues in the group are using.

If the other unicast queues use less than their allocated amount of bandwidth, the unconfigured queues can share the unused bandwidth. Sharing unused bandwidth is one of the key advantages of hierarchical port scheduling. Configured queues have higher priority for bandwidth than unconfigured queues, so if a configured queue needs more bandwidth, then less bandwidth is available for unconfigured queues. Unconfigured queues always receive a minimum amount of bandwidth based on their scheduling weight (1). If you map traffic to an unconfigured queue, to allocate bandwidth to that queue, configure a scheduler for the forwarding class that is mapped to the queue.

Default Scheduler Maps

Table 22 on page 45 shows the default mapping of forwarding classes to schedulers.

Table 22: Default Scheduler Maps

Forwarding Class	Scheduler
best-effort	Default BE scheduler
fcoe	Default FCoE scheduler
no-loss	No-loss scheduler
network-control	Default network-control scheduler
mcast-be	Default multideestination scheduler

Default Shared Buffer Configuration

Table 23 on page 45 and Table 24 on page 45 show the default shared buffer allocations:

Table 23: Default Ingress Shared Buffer Configuration

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
100%	9%	45%	46%

Table 24: Default Egress Shared Buffer Configuration

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	50%	31%	19%

**Related
Documentation**

- [Overview of Junos OS CoS for the QFX Series on page 4](#)
- [Understanding Junos CoS Components on page 30](#)
- [Understanding Default CoS Scheduling and Classification on page 74](#)
- [Understanding CoS Classifiers on page 51](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57](#)
- [Understanding CoS Code-Point Aliases on page 49](#)
- [Understanding CoS Forwarding Classes on page 68](#)
- [Understanding CoS Rewrite Rules on page 154](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Understanding CoS WRED Drop Profiles on page 149](#)

Understanding Host Inbound Traffic Classification

The destination address of traffic that enters the switch can be an external device such as another switch, a router, or a server, or the destination can be the host (the switch Routing Engine or CPU). When the destination is an external device, the DSCP and IEEE 802.1p code-point bits of incoming traffic are preserved as the traffic travels through the switch to the egress port. At the egress port, the code-point bits are either preserved when the packets are sent to the next hop or they are rewritten according to the rewrite rule attached to the egress interface.

When the destination of incoming traffic is the host, DSCP bits are preserved. However, IEEE 802.1p bits are not preserved. The IEEE 802.1p bits of traffic destined for the host are set to zero (0). This does not affect system behavior because the switch prioritizes traffic destined for the host based on the protocol type. For example, the switch gives a higher priority to BPDU traffic than to ping traffic.

**Related
Documentation**

- [Understanding Default CoS Scheduling and Classification on page 74](#)
- [Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 47](#)

Understanding Host Routing Engine Outbound Traffic Queues and Defaults

The host Routing Engine and CPU generate outbound traffic that is transmitted using different protocols. You cannot configure a classifier to map different types of outbound traffic that the host generates to forwarding classes (queues). The traffic that the host generates is assigned to forwarding classes by default as shown in [Table 25 on page 47](#).

If you want to separate host outbound traffic from other traffic or if you want to assign that traffic to a particular queue, you can configure a single forwarding class for all traffic that the host generates. If you configure a forwarding class for outbound host traffic, that forwarding class is used globally for all traffic generated by the host. (That is, the host outbound traffic is mapped to the selected queue on all egress interfaces.) Configuring a forwarding class for host outbound traffic does not affect transit or incoming traffic.

Whether you use the default host outbound traffic forwarding class configuration or configure a forwarding class for all host outbound traffic, the configuration applies to all Layer 2 and Layer 3 protocols and to all application-level traffic such as FTP and ping operations.

If you configure a queue for host outbound traffic, the queue must be properly configured on all interfaces.



NOTE: Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) packets generated by the CPU are always transmitted on the `fcoe` queue (queue 3), even if you configure a queue for host outbound traffic. This helps to ensure lossless behavior for FCoE traffic. QFabric systems classify FIP control packets into the same traffic class (`fcoe`) across the Interconnect device (fabric) and the egress Node device.

By default, traffic generated by the host is sent to the best effort queue (queue 0) or to the network control queue (queue 7). [Table 25 on page 47](#) lists the default host traffic to output queue mapping.

Table 25: Routing Engine Protocol Default Queue Mapping

Routing Engine Protocol	Default Queue Mapping
Address Resolution Protocol (ARP) reply	Queue 0
ARP request	Queue 0
Border Gateway Protocol (BGP)	Queue 0
BGP TCP Retransmission	Queue 7
Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP)	Queue 3
File Transfer Protocol (FTP)	Queue 0

Table 25: Routing Engine Protocol Default Queue Mapping (*continued*)

Routing Engine Protocol	Default Queue Mapping
Internet Control Message Protocol (ICMP) reply	Queue 0
ICMP request	Queue 0
Internet Group Management Protocol (IGMP) query	Queue 7
IGMP report	Queue 0
Link Aggregation Control Protocol (LACP)	Queue 7
Open Shortest Path First (OSPF) hello	Queue 7
OSPF protocol data unit (PDU)	Queue 7
OSPF link state advertisements (LSAs)	Queue 7
Protocol Independent Multicast (PIM)	Queue 7
PIM hello	Queue 7
Simple Network Management Protocol (SNMP)	Queue 0
Secure Shell (SSH)	Queue 0
Telnet	Queue 0
Virtual Router Redundancy Protocol (VRRP)	Queue 7
VLAN Spanning Tree Protocol (VSTP)	Queue 7
xnm-clear-text	Queue 0
xnm-ssl	Queue 0

- Related Documentation**
- [Understanding CoS Forwarding Classes on page 68](#)
 - [Changing the Host Outbound Traffic Default Queue Mapping on page 415](#)
 - [Example: Configuring Forwarding Classes on page 283](#)

Understanding CoS Code-Point Aliases

A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.

Behavior aggregate classifiers use class-of-service (CoS) values such as Differentiated Services code points (DSCPs) and IEEE 802.1 bits to associate incoming packets with a particular CoS servicing level. You can assign a meaningful name or alias to the CoS values and use that alias instead of bits when configuring CoS components. These aliases are not part of the specifications but are well known through usage. For example, the alias for DSCP 101110 is widely accepted as ef (expedited forwarding).

When you configure classes and define classifiers, you can refer to the markers by alias names. You can configure alias names for user-defined classifiers. If the value of an alias changes, it alters the behavior of any classifier that references it.

You can configure code-point aliases for the following type of CoS markers:

- dscp or dscp-ipv6—Handles incoming IP and IPv6 packets.
- ieee-802.1—Handles Layer 2 CoS.

This topic covers:

- [Default Code-Point Aliases on page 49](#)

Default Code-Point Aliases

[Table 26 on page 49](#) shows the default mapping of code-point aliases to IEEE code points.

Table 26: Default IEEE 802.1 Code-Point Aliases

CoS Value Types	Mapping
be	000
be1	001
ef	010
ef1	011
af11	100
af12	101
nc1	110
nc2	111

[Table 27 on page 50](#) shows the default mapping of code-point aliases to DSCP and DSCP IPv6 code points.

Table 27: Default DSCP and DSCP IPv6 Code-Point Aliases

CoS Value Types	Mapping
ef	101110
af11	001010
af12	001100
af13	001110
af21	010010
af22	010100
af23	010110
af31	011010
af32	011100
af33	011110
af41	100010
af42	100100
af43	100110
be	000000
cs1	001000
cs2	010000
cs3	011000
cs4	100000
cs5	101000
nc1	110000
nc2	111000

Related Documentation

- [Understanding Junos CoS Components on page 30](#)

- [Defining CoS Code-Point Aliases on page 404](#)

Understanding CoS Classifiers

Packet classification associates incoming packets with a particular class-of-service (CoS) servicing level. Classifiers associate packets with a forwarding class and a loss priority, and assign packets to output queues based on the associated forwarding class. There are three general types of classifiers:

- Behavior aggregate (BA) classifiers—DSCP and DSCP IPv6 classify IP and IPv6 traffic, EXP classifies MPLS traffic, and IEEE 802.1p classifiers classify all other traffic. (Although this topic covers EXP classifiers, for more details about EXP classifiers, see *Understanding CoS MPLS EXP Classifiers and Rewrite Rules*. EXP classifiers are applied only on **family mpls** interfaces.)
 - Fixed classifiers—Fixed classifiers classify all ingress traffic on a physical interface into one forwarding class, regardless of the CoS bits in the packet header.
 - Multifield (MF) classifiers—MF classifiers classify traffic based on more than one field in the packet header and take precedence over BA and fixed classifiers.
- [Interfaces and Output Queues on page 51](#)
 - [Behavior Aggregate Classifiers on page 52](#)
 - [Fixed Classifiers on Ethernet Interfaces on page 55](#)
 - [Fixed Classifiers on Native Fibre Channel Interfaces \(NP_Ports\) on page 56](#)
 - [Multifield Classifiers on page 56](#)
 - [Packet Classification for Routed VLAN Interfaces \(RVIs\) on page 57](#)

Interfaces and Output Queues

On Gigabit Ethernet interfaces, 10-Gigabit Ethernet interfaces, and link aggregation (LAG) interfaces, you can apply classifiers to Layer 2 logical interfaces and to Layer 3 physical interfaces if the Layer 3 physical interface has at least one defined logical interface. Classifiers applied to Layer 3 physical interfaces are used on all logical interfaces on that physical interface. “[Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces](#)” on page 57 describes the interaction between classifiers and interfaces in greater detail.

You can configure both a BA classifier and an MF classifier on an interface. If you do this, the BA classification is performed first and then the MF classification is performed. If the two classification results conflict, the MF classification result overrides the BA classification result.

You cannot configure a fixed classifier and a BA classifier on the same interface.

You can configure both a DSCP or a DSCP IPv6 classifier and an IEEE 802.1p classifier on the same interface. IP traffic uses the DSCP or DSCP IPv6 classifier. All other traffic uses the IEEE classifier (except when you configure a global EXP classifier; in that case, MPLS traffic uses the EXP classifier providing that the interface is configured as **family**

mpls). You can configure only one DSCP classifier on a physical interface (either one DSCP classifier or one DSCP IPv6 classifier, but not both).

Although you can configure as many EXP classifiers as you want, the switch uses only one MPLS EXP classifier as a global classifier on all interfaces. After you configure an MPLS EXP classifier, you can configure it as the global EXP classifier by including the EXP classifier in the **[edit class-of-service system-defaults classifiers exp]** hierarchy. All switch interfaces that are configured as **family mpls** use the EXP classifier specified using this configuration statement to classify MPLS traffic.

You can create unicast BA classifiers for unicast traffic and multicast BA classifiers for multdestination traffic, which includes multicast, broadcast, and destination lookup fail (DLF) traffic. You cannot assign unicast traffic and multdestination traffic to the same BA classifier.

On each interface, the switch has separate output queues for unicast traffic and for multdestination traffic:

- The switch supports 12 output queues, with 8 queues dedicated to unicast traffic and 4 queues dedicated to multdestination traffic.
- Queues 0 through 7 are unicast traffic queues. You can apply only unicast BA classifiers to unicast queues. A unicast BA classifier should contain only forwarding classes that are mapped to unicast queues.
- Queues 8 through 11 are multdestination traffic queues. You can apply only multdestination BA classifiers to multdestination queues. A multdestination BA classifier should contain only forwarding classes that are mapped to multdestination queues.

You can apply unicast classifiers to one or more interfaces. Multdestination classifiers and EXP classifiers apply to all of the switch interfaces and cannot be applied to individual interfaces. Use the DSCP multdestination classifier for both IP and IPv6 multdestination traffic. The DSCP IPv6 classifier is not supported for multdestination traffic.

Behavior Aggregate Classifiers

The behavior aggregate classifier maps a class-of-service (CoS) value to a forwarding class and loss priority. The forwarding class determines the output queue. A scheduler uses the loss priority to control packet discard during periods of congestion by associating different drop profiles with different loss priorities.

The switch supports three types of BA classifiers:

- Differentiated Services Code Point (DSCP) for IP DiffServ (IP and IPv6)
- IEEE 802.1p CoS bits
- MPLS EXP (applies only to interfaces configured as **family mpls**)

BA classifiers are based on fixed-length fields, which makes them computationally more efficient than MF classifiers. Therefore, core devices, which handle high traffic volumes, are normally configured to perform BA classification.

Unicast and multicast traffic cannot share the same classifier. You can map unicast traffic and multicast traffic to the same classifier CoS value, but the unicast traffic must belong to a unicast classifier and the multicast traffic must belong to a multidestination classifier.

Default Behavior Aggregate Classification

Juniper Networks Junos OS automatically assigns implicit default classifiers to all logical interfaces based on the type of interface. Table 28 on page 53 lists different types of interfaces and the corresponding implicit default BA classifiers.

Table 28: Default BA Classification

Type of Interface	Default BA Classification
Layer 2 interface in trunk mode or in tagged-access mode	ieee8021p-default
Layer 3 interface	dscp-default
Layer 2 interface in access mode	ieee8021p-untrusted



NOTE: There are default BA classifiers for the **best-effort**, **fcoe**, **no-loss**, **network-control**, and **mcast** forwarding classes.



NOTE: There is no default MPLS EXP classifier. You must configure an EXP classifier and apply it globally to all interfaces that are configured as family mpls by including it in the [edit class-of-service system-defaults classifiers exp] hierarchy. On family mpls interfaces, if a fixed classifier is present on the interface, the EXP classifier overrides the fixed classifier.

If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

Because the EXP classifier is global, you cannot configure some ports to use a fixed IEEE 802.1p classifier for MPLS traffic on some interfaces and the global EXP classifier for MPLS traffic on other interfaces. When you configure a global EXP classifier, all MPLS traffic on all interfaces uses the EXP classifier, even interfaces that have a fixed classifier.

When you explicitly associate a unicast classifier with a logical interface, you override the default unicast classifier with the explicit unicast classifier.



NOTE: You can apply only one classifier of each type, DSCP and IEEE 802.1p, to a Layer 2 interface. If both types of classifiers are present, DSCP classifiers take precedence over IEEE 802.1p classifiers. (If you also configure a global EXP classifier, only MPLS traffic on interfaces configured as family `mpls` uses the EXP classifier, and other traffic uses the configured or default classifier for that traffic type.)

Importing a Classifier

You can use any existing classifier, including the default classifiers, as the basis for defining a new classifier. You accomplish this using the **import** statement.

The imported classifier is used as a template and is not modified. The modifications you make become part of a new classifier (and a new template) identified by the name of the new classifier. Whenever you commit a configuration that assigns a new class-name and loss-priority value to a code-point alias or set of bits, it replaces that entry in the new classifier template. As a result, you must explicitly specify every CoS value in every designation that requires modification.

Multidestination Classifiers

Multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces. You can configure both a DSCP multidestination classifier and an IEEE multidestination classifier. IP and IPv6 traffic use the DSCP classifier, and all other traffic uses the IEEE classifier.

DSCP IPv6 multidestination classifiers are not supported, so IPv6 traffic uses the DSCP multidestination classifier.

The default multidestination classifier is the IEEE 802.1p multidestination classifier.

PFC Priorities

The eight IEEE 802.1p code points correspond to the eight priorities that priority-based flow control (PFC) uses to differentiate traffic classes for lossless transport. When you map a forwarding class (which maps to an output queue) to an IEEE 802.1p CoS value, the IEEE 802.1p CoS value identifies the priority.

Although you can map a priority to any output queue (by mapping the priority to a forwarding class), we recommend that the priority and the unicast forwarding class match in a one-to-one correspondence in which priority 0 is assigned to queue 0, priority 1 is assigned to queue 1, and so on, as shown in [Table 29 on page 55](#). A one-to-one correspondence of queue and priority numbers makes it easier to configure and maintain the mapping of forwarding classes to priorities and queues.

Table 29: Default IEEE 802.1p Code Point to PFC Priority, Output Queue, and Forwarding Class Mapping

IEEE 802.1p Code Point	PFC Priority	Unicast Output Queue	Forwarding Class and Packet Drop Attribute
000	0	0	best-effort (drop)
001	1	1	best-effort (drop)
010	2	2	best-effort (drop)
011	3	3	fcoe (no-loss)
100	4	4	no-loss (no-loss)
101	5	5	best-effort (drop)
110	6	6	network-control (drop)
111	7	7	network-control (drop)



NOTE: By convention, deployments with converged server access typically use IEEE 802.1p priority 3 (011) for FCoE traffic. The default mapping of the fcoe forwarding class is to queue 3. Apply priority-based flow control (PFC) to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE requires. We recommend that you use priority 3 for FCoE traffic unless your network architecture requires that you use a different priority.

Fixed Classifiers on Ethernet Interfaces

Fixed classifiers map all traffic on an interface to a forwarding class and a loss priority. (As opposed to BA classifiers, which map traffic into multiple different forwarding classes based on the CoS field value in the packet header.) The forwarding class determines the output queue. Incoming traffic of all IEEE 802.1p priorities is classified into the forwarding class specified in the fixed classifier. A scheduler uses the loss priority to control packet discard during periods of congestion by associating different drop profiles with different loss priorities.

You cannot configure a fixed classifier and a DSCP or IEEE 802.1p BA classifier on the same interface. If you configure a fixed classifier on an interface, you cannot configure a DSCP or an IEEE classifier on that interface. If you configure a DSCP classifier, an IEEE classifier, or both classifiers on an interface, you cannot configure a fixed classifier on that interface.



NOTE: Because EXP classifiers are global, you can configure both a global EXP classifier and also apply fixed classifiers on interfaces. When both the global EXP classifier and a fixed classifier are applied to an interface, MPLS traffic on interfaces configured as family mpls uses the EXP classifier and all other traffic uses the fixed classifier.

To switch from a fixed classifier to a BA classifier or to switch from a BA classifier to a fixed classifier, deactivate the existing classifier attachment on the interface, and then attach the new classifier to the interface.



NOTE: If you configure a fixed classifier that classifies all incoming traffic into the fcoe forwarding class (or any forwarding class designed to handle FCoE traffic), you must ensure that all traffic that enters the interface is FCoE traffic and is tagged with the FCoE IEEE 802.1p code point (priority).

Fixed Classifiers on Native Fibre Channel Interfaces (NP_Ports)

Applying a fixed classifier to a native Fibre Channel (FC) interface (NP_Port) is a special case. By default, native FC interfaces classify incoming traffic from the FC SAN into the fcoe forwarding class and map the traffic to IEEE 802.1p priority 3 (code point 011). When you apply a fixed classifier to an FC interface, you also configure a priority rewrite value for the interface. The FC interface uses the priority rewrite value as the IEEE 802.1p tag value for all incoming packets instead of the default value of 3.

For example, if you specify a priority rewrite value of 5 (code point 101) for an FC interface, the interface tags all incoming traffic from the FC SAN with priority 5 and classifies the traffic into the forwarding class specified in the fixed classifier.



NOTE: The forwarding class specified in the fixed classifier on FC interfaces must be a lossless forwarding class.

Multifield Classifiers

Multifield classifiers examine multiple fields in a packet such as source and destination addresses and source and destination port numbers of the packet. With MF classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

MF classification is normally performed at the network edge because of the general lack of DiffServ code point (DSCP) support in end-user applications. On a switch at the edge of a network, an MF classifier provides the filtering functionality that scans through a variety of packet fields to determine the forwarding class for a packet. Typically, a classifier performs matching operations on the selected fields against a configured value.

Packet Classification for Routed VLAN Interfaces (RVIs)

You cannot apply classifiers directly to routed VLAN interfaces (RVIs) because the members of RVIs are VLANs, not ports. However, you can apply classifiers to the VLAN port members of an RVI. You can also apply MF classifiers to RVIs.

Related Documentation

- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules](#)
- [Understanding CoS Packet Flow on page 34](#)
- [Understanding Default CoS Settings on page 37](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57](#)
- [Example: Configuring Unicast Classifiers on page 273](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 277](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 404](#)
- [Configuring a Global MPLS EXP Classifier](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers](#)
- [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 406](#)

Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces

At ingress interfaces, classifiers group incoming traffic into classes based on the IEEE 802.1p, DSCP, or MPLS EXP class of service (CoS) code point bits in the packet header. At egress interfaces, you can use rewrite rules to change (re-mark) the code point bits before the interface forwards the packets. At ingress interfaces, classifiers group incoming traffic into classes based on the IEEE 802.1p, DSCP, or MPLS EXP CoS code point bits in the packet header. At egress interfaces, rewrite rules can change (re-mark) the code point bits before the interface forwards the packets.

You can apply classifiers and rewrite rules to interfaces to control the level of CoS applied to each packet as it traverses the system and the network. This topic describes:

- [Supported Classifier and Rewrite Rule Types on page 58](#)
- [Ethernet Interfaces Supported for Classifier and Rewrite Rule Configuration on page 59](#)
- [Default Classifiers on page 61](#)
- [Default Rewrite Rules on page 61](#)
- [Classifier Precedence on page 61](#)
- [Classifier Behavior and Limitations on page 63](#)
- [Rewrite Rule Precedence and Behavior on page 64](#)
- [Classifier and Rewrite Rule Configuration Interaction with Ethernet Interface Configuration on page 64](#)

Supported Classifier and Rewrite Rule Types

Table 30 on page 58 shows the types of classifiers and rewrite rules that the QFX Series supports:

Table 30: Supported Classifiers and Rewrite Rules

Classifier or Rewrite Rule Type	Description
Fixed classifier	Classifies all ingress traffic on a physical interface into one fixed forwarding class, regardless of the CoS bits in the packet header.
DSCP and DSCP IPv6 unicast classifiers	Classifies IP and IPv6 traffic into forwarding classes and assigns loss priorities to the traffic.
IEEE 802.1p unicast classifier	Classifies Ethernet traffic into forwarding classes and assigns loss priorities to the traffic.
MPLS EXP classifier	Classifies MPLS traffic into forwarding classes and assigns loss priorities to the traffic on interfaces configured as family mpls . The system uses one global EXP classifier on all family mpls switch interfaces.
DSCP multidestination classifier (also used for IPv6 multidestination traffic)	Classifies IP and IPv6 multicast, broadcast, and destination lookup fail (DLF) traffic into multidestination forwarding classes. Multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces.
IEEE 802.1p multidestination classifier	Classifies Ethernet multicast, broadcast, and destination lookup fail (DLF) traffic into multidestination forwarding classes. Multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces.
DSCP and DSCP IPv6 rewrite rules	Re-marks the DSCP code points of IP and IPv6 packets before forwarding the packets.
IEEE 802.1p rewrite rule	Re-marks the IEEE 802.1p code points of Ethernet packets before forwarding the packets.
MPLS EXP rewrite rule	Re-marks the EXP code points of MPLS packets before forwarding the packets on interfaces configured as family mpls .



NOTE: On native Fibre Channel (FC) interfaces (NP_Ports) only, you can specify a rewrite value to set the IEEE 802.1p code point of incoming FC traffic when the NP_Port encapsulates the FC packet in Ethernet before forwarding it to the FCoE network (see “[Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway](#)” on page 128).

DSCP, IEEE 802.1p, and MPLS EXP classifiers are behavior aggregate (BA) classifiers. Unlike DSCP and IEEE 802.1p classifiers, EXP classifiers are global and apply only to all interfaces that are configured as **family mpls**. Also unlike DSCP and IEEE 802.1p classifiers, for MPLS traffic only, EXP classifiers overwrite fixed classifiers. (An interface that has a fixed classifier uses the EXP classifier for MPLS traffic, not the fixed classifier.)

Multidestination classifiers are global and apply to all interfaces; you cannot apply a multidestination classifier to individual interfaces.

Classifying packets into forwarding classes assigns packets to the output queues associated with the forwarding classes. Classifying traffic into a forwarding class associates the CoS scheduling for the forwarding class with that traffic.



NOTE: In addition to BA classifiers and fixed classifiers, which classify traffic based on the CoS field in the packet header, you can use firewall filters to configure multifield (MF) classifiers. MF classifiers classify traffic based on more than one field in the packet header and take precedence over BA and fixed classifiers.

Ethernet Interfaces Supported for Classifier and Rewrite Rule Configuration

To apply a classifier to incoming traffic or a rewrite rule to outgoing traffic, you need to apply the classifier or rewrite rule to one or more interfaces. When you apply a classifier or rewrite rule to an interface, the interface uses the classifier to group incoming traffic into forwarding classes and uses the rewrite rule to re-mark the CoS code point value of each packet before it leaves the system.

Not all interfaces types support all types of CoS configuration. This section describes:

- [Interface Types That Support Classifier and Rewrite Rule Configuration on page 59](#)
- [Classifier and Rewrite Rule Physical and Logical Ethernet Interface Support on page 59](#)
- [Routed VLAN Interfaces \(RVIs\) and Integrated Routing and Bridging \(IRB\) Interfaces on page 61](#)

Interface Types That Support Classifier and Rewrite Rule Configuration

You can apply classifiers to all Ethernet interfaces. For Layer 3 LAGs, configure BA or fixed classifiers on the LAG (ae) interface. The classifier configured on the LAG is valid on all of the LAG member interfaces.

You can apply fixed classifiers to native FC interfaces (NP_Ports). You cannot apply other types of classifiers or rewrite rules to native FC interfaces. You can rewrite the value of the IEEE 802.1p code point of incoming FC traffic when the interface encapsulates it in Ethernet before forwarding it to the FCoE network as described in [“Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway” on page 128](#).

Classifier and Rewrite Rule Physical and Logical Ethernet Interface Support

The QFX Series Ethernet ports can function as:

- Layer 2 physical interfaces (family ethernet-switching)
- Layer 2 logical interfaces (family ethernet-switching)
- Layer 3 physical interfaces (family inet/inet6)
- Layer 3 logical interfaces (family inet/inet6)

You can apply CoS classifiers and rewrite rules only to the following interfaces:

- Layer 2 logical interfaces
- Layer 3 physical interfaces if at least one logical Layer 3 interface is configured on the physical interface



NOTE: The CoS you configure on a Layer 3 physical interface is applied to all of the Layer 3 logical interfaces on that physical interface. This means that each Layer 3 interface uses the same classifiers and rewrite rules for all of the Layer 3 traffic on that interface.

You cannot apply classifiers or rewrite rules to Layer 2 physical interfaces or to Layer 3 logical interfaces. [Table 31 on page 60](#) shows on which interfaces you can configure and apply classifiers and rewrite rules.

Table 31: Ethernet Interface Support for Classifier and Rewrite Rule Configuration

CoS Classifiers and Rewrite Rules	Layer 2 Physical Interfaces	Layer 2 Logical Interfaces	Layer 3 Physical Interfaces (If at Least One Logical Layer 3 Interface Is Defined)	Layer 3 Logical Interfaces
Fixed classifier	No	Yes	Yes	No
DSCP classifier	No	Yes	Yes	No
DSCP IPv6 classifier	No	Yes	Yes	No
IEEE 802.1p classifier	No	Yes	Yes	No
EXP classifier	Global classifier, applies only to all switch interfaces that are configured as family mpls . Cannot be configured on individual interfaces.			
DSCP rewrite rule	No	Yes	Yes	No
DSCP IPv6 rewrite rule	No	Yes	Yes	No
IEEE 802.1p rewrite rule	No	Yes	Yes	No
EXP rewrite rule	No	Yes	Yes	No



NOTE: IEEE 802.1p multidestination and DSCP multidestination classifiers are applied to all interfaces and cannot be applied to individual interfaces. No DSCP IPv6 multidestination classifier is supported. IPv6 multidestination traffic uses the DSCP multidestination classifier.

Routed VLAN Interfaces (RVIs) and Integrated Routing and Bridging (IRB) Interfaces

You cannot apply classifiers and rewrite rules directly to routed VLAN interfaces (RVIs) or integrated routing and bridging (IRB) interfaces because the members of RVIs and IRBs are VLANs, not ports. However, you can apply classifiers and rewrite rules to the VLAN port members of an RVI or an IRB. You can also apply MF classifiers to RVIs and IRBs.

Default Classifiers

If you do not explicitly configure classifiers on an Ethernet interface, the QFX Series applies default classifiers (see [“Understanding Default CoS Settings” on page 37](#)) so that the traffic receives basic CoS treatment. The factors that determine the default classifier applied to the interface include the interface type (Layer 2 or Layer 3), the port mode (trunk, tagged-access, or access), and whether logical interfaces have been configured. The system applies a default classifier using the following rules:

- If the physical interface has at least one Layer 3 logical interface configured, it uses the default DSCP classifier.
- If the physical interface has a Layer 2 logical interface in trunk mode or tagged-access mode, it uses the default trusted classifier.
- If the physical interface has a Layer 2 logical interface in access mode, it uses the default untrusted classifier.
- If the physical interface has no logical interface configured, no default classifier is applied.
- The default multdestination classifier is the IEEE 802.1p multdestination classifier.
- There is no default MPLS EXP classifier. If you want to classify traffic using EXP bits, you must configure an EXP classifier and configure it as the global system default EXP classifier.

Default Rewrite Rules

No default rewrite rules are applied to interfaces. If you want to re-mark packets at the egress interface, you must explicitly configure a rewrite rule.

Classifier Precedence

You can apply multiple unicast classifiers (MF, fixed, IEEE 802.1p, DSCP, or EXP) to a physical or logical Ethernet interface to handle different types of traffic. (EXP classifiers are global and apply only to all MPLS traffic on all **family mpls** interfaces.) When you apply more than one classifier to an interface, the system uses an order of precedence to determine which classifier to use on physical and logical interfaces:

- [Unicast Classifier Precedence on Physical Ethernet Interfaces on page 62](#)
- [Unicast Classifier Precedence on Logical Ethernet Interfaces on page 62](#)

Unicast Classifier Precedence on Physical Ethernet Interfaces

The precedence of unicast classifiers on physical interfaces, from the highest-priority classifier to the lowest-priority classifier, is:

- MF classifier on a logical interface (no classifier has a higher priority than MF classifiers)
- Fixed classifier on the physical interface
- DSCP or DSCP IPv6 classifier on the physical interface
- IEEE 802.1p classifier on the physical interface



NOTE: If an EXP classifier is configured, MPLS traffic uses the EXP classifier on all family mpls interfaces, even if an MF or fixed classifier is applied to the interface. If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic. DSCP classifiers are not applied to MPLS traffic.

You can apply a DSCP classifier, an IEEE 802.1p classifier, and an EXP classifier on a physical interface. When all three classifiers are on an interface, IP traffic uses the DSCP classifier, MPLS traffic uses the EXP classifier, and all other traffic uses the IEEE classifier.



NOTE: You cannot apply a fixed classifier and a DSCP or IEEE classifier to the same interface. If a DSCP classifier, an IEEE classifier, or both are on an interface, you cannot apply a fixed classifier to that interface unless you first delete the DSCP and IEEE classifiers. If a fixed classifier is on an interface, you cannot apply a DSCP classifier or an IEEE classifier unless you first delete the fixed classifier.

Unicast Classifier Precedence on Logical Ethernet Interfaces

The precedence of unicast classifiers on logical interfaces, from the highest priority classifier to the lowest priority classifier, is:

- MF classifier on a logical interface (no classifier has a higher priority than MF classifiers)
- Fixed classifier on the logical interface
- DSCP or DSCP IPv6 classifier on the physical interface
- IEEE 802.1p classifier on the physical interface



NOTE: If an EXP classifier is configured, MPLS traffic uses the EXP classifier on all **family mpls** interfaces, even if a fixed classifier is applied to the interface. If an EXP classifier is not configured, then if a fixed classifier is applied to the interface, the MPLS traffic uses the fixed classifier. If no EXP classifier and no fixed classifier is applied to the interface, MPLS traffic is treated as best-effort traffic.

You can apply both a DSCP classifier and an IEEE 802.1p classifier on a logical interface. When both a DSCP and an IEEE classifier are on an interface, IP traffic uses the DSCP classifier, and all other traffic uses the IEEE classifier. If an MPLS EXP classifier is also applied to the interface, only MPLS traffic uses the EXP classifier.

Classifier Behavior and Limitations

Consider the following behaviors and constraints when you apply classifiers to physical and logical Ethernet interfaces:

- You can configure only one DSCP classifier (IP or IPv6) on a physical interface. You cannot configure both types of DSCP classifier on one physical interface. Both IP and IPv6 traffic use whichever DSCP classifier is configured on the interface.
- When you configure a DSCP or a DSCP IPv6 classifier on a physical interface and the physical interface has at least one logical Layer 3 interface, all packets (IP, IPv6, and non-IP) use that classifier.
- An interface with both a DSCP classifier (IP or IPv6) and an IEEE 802.1p classifier uses the DSCP classifier for IP and IPv6 packets, and uses the IEEE classifier for all other packets.
- Fixed classifiers and BA classifiers (DSCP and IEEE classifiers) are not permitted simultaneously on an interface. If you configure a fixed classifier on an interface, you cannot configure a DSCP or an IEEE classifier on that interface. If you configure a DSCP classifier, an IEEE classifier, or both classifiers on an interface, you cannot configure a fixed classifier on that interface.
- When you configure an IEEE 802.1p classifier on a physical interface and a DSCP classifier is not explicitly configured on that interface, the interface uses the IEEE classifier for all types of packets. No default DSCP classifier is applied to the interface. (In this case, if you want a DSCP classifier on the interface, you must explicitly configure it.)
- The system does not apply a default classifier to a physical interface until you create a logical interface on that physical interface. If you configure a Layer 3 logical interface, the system uses the default DSCP classifier. If you configure a Layer 2 logical interface, the system uses the default IEEE 802.1p trusted classifier if the port is in trunk mode or tagged-access mode, or the default IEEE 802.1p untrusted classifier if the port is in access mode.
- MF classifiers configured on logical interfaces take precedence over BA and fixed classifiers, with the exception of the global EXP classifier, which is always used for MPLS traffic on **family mpls** interfaces. (Use firewall filters to configure MF classifiers.)

When BA or fixed classifiers are present on an interface, you can still configure an MF classifier on that interface.

- There is no default EXP classifier for MPLS traffic.
- You can configure as many EXP classifiers as you want, but the switch uses only one MPLS EXP classifier as a global classifier on all **family mpls** interfaces. After you configure an MPLS EXP classifier, you can configure it as the global EXP classifier by including the EXP classifier in the **[edit class-of-service system-defaults classifiers exp]** hierarchy. All **family mpls** switch interfaces use the EXP classifier specified using this configuration statement to classify MPLS traffic, even on interfaces that have a fixed classifier. No other traffic uses the EXP classifier.

Rewrite Rule Precedence and Behavior

The following rules apply on both physical and logical Ethernet interfaces for rewrite rules:

- If you configure both one DSCP (or DSCP IPv6) rewrite rule and one IEEE 802.1p rewrite rule on an interface, both rewrite rules take effect. Traffic with IP and IPv6 headers use the DSCP rewrite rule, and traffic with a VLAN tag uses the IEEE rewrite rule.
- If you do not explicitly configure a rewrite rule, there is no default rewrite rule, so the system does not apply any rewrite rule to the interface.
- You can apply a DSCP rewrite rule or a DSCP IPv6 rewrite rule to an interface, but you cannot apply both a DSCP and a DSCP IPv6 rewrite rule to the same interface. Both IP and IPv6 packets use the same DSCP rewrite rule, regardless if the configured rewrite rule is DSCP or DSCP IPv6.
- MPLS EXP rewrite rules apply only to logical interfaces on **family mpls** interfaces. You cannot apply to an EXP rewrite rule to a physical interface. You can configure as many EXP rewrite rules as you want, but you can only use 16 EXP rewrite rules at any time on the switch.
- A logical interface can use both DSCP (or DSCP IPv6) and EXP rewrite rules.
- DSCP and DSCP IPv6 rewrite rules are not applied to MPLS traffic.
- If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.

Classifier and Rewrite Rule Configuration Interaction with Ethernet Interface Configuration

You can apply classifiers and rewrite rules only on Layer 2 logical interfaces and Layer 3 physical interfaces (if the Layer 3 physical interface has at least one defined logical interface). This section focuses on BA classifiers, but the interaction between BA classifiers and interfaces described in this section also applies to fixed classifiers and rewrite rules.



NOTE: Multidestination classifiers, and EXP classifiers (only on family mpls interfaces), are global and apply to all switch interfaces. See [“Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\)” on page 404](#) for how to configure multidestination classifiers and see [“Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\)” on page 404](#) for how to configure EXP classifiers.

There are two components to applying classifiers or rewrite rules to interfaces:

1. Setting the interface family (inet, inet6, or ethernet-switching; ethernet-switching is the default interface family) in the **[edit interfaces]** configuration hierarchy.
2. Applying a classifier or rewrite rule to the interface in the **[edit class-of-service]** hierarchy.

These are separate operations that can be set and committed at different times. Because the type of classifier or rewrite rule you can apply to an interface depends on the interface family configuration, the system performs checks to ensure that the configuration is valid. The method the system uses to notify you of an invalid configuration depends on the **set** operation that causes the invalid configuration.

When applying the classifier or rewrite rule to the interface in the **[edit class-of-service]** hierarchy causes an invalid configuration, the system rejects the configuration and returns a commit check error.

When setting the interface family in the **[edit interfaces]** configuration hierarchy causes an invalid configuration, the system creates a syslog error message. When you receive the error message, you need to remove the classifier or rewrite rule configuration from the logical interface and apply it to the physical interface, or remove the classifier or rewrite rule configuration from the physical interface and apply it to the logical interface. For classifiers, if you do not take action to correct the error, the system programs the default classifier for the interface family on the interface. (There are no default rewrite rules. If the commit check fails, no rewrite rule is applied to the interface.)

Two scenarios illustrate these situations:

- [Scenario 1: Applying a Classifier to an Ethernet Interface Causes a Commit Check Error on page 66](#)
- [Scenario 2: Configuring the Ethernet Interface Family Causes a Syslog Error on page 66](#)



NOTE: Both of these scenarios also apply to fixed classifiers and rewrite rules.

Scenario 1: Applying a Classifier to an Ethernet Interface Causes a Commit Check Error

In Scenario 1, we set the interface family, and then specify an invalid classifier.

1. Set and commit the interface as a Layer 3 (family **inet**) interface:

```
[edit interfaces]
user@switch# set xe-0/0/20 unit 0 family inet
user@switch# commit
```

This commit operation succeeds.

2. Set and commit a DSCP classifier on the logical interface (this example uses a DSCP classifier named **dscp1**):

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
```

This configuration is not valid, because it attempts to apply a classifier to a Layer 3 logical interface. Because the failure is caused by the class-of-service configuration and not by the interface configuration, the system rejects the commit operation and issues a commit error, not a syslog message.

Note that the commit operation succeeds if you apply the classifier to the physical Layer 3 interface as follows:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 classifiers dscp dscp1
user@switch# commit
```

Because the logical unit is not specified, the classifier is applied to the physical Layer 3 interface in a valid configuration, and the commit check succeeds.

Scenario 2: Configuring the Ethernet Interface Family Causes a Syslog Error

In Scenario 2, we set the classifier first, then set an invalid interface type.

1. Set and commit a DSCP classifier on a Layer 3 logical interface, assuming that the interface has no existing configuration:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
```

This commit succeeds. Because no explicit configuration existed on the interface, it is by default a Layer 2 (**family ethernet-switching**) interface. Layer 2 logical interfaces support BA classifiers, so applying the classifier is a valid configuration.

2. Set and commit the interface as a Layer 3 interface (family **inet**) interface:

```
[edit interfaces]
user@switch# set xe-0/0/20 unit 0 family inet
user@switch# commit
```

This configuration is not valid because it attempts to change an interface from Layer 2 (**family ethernet-switching**) to Layer 3 (**family inet**) when a classifier has already been applied to a logical interface. Layer 3 logical interfaces do not support classifiers. Because the failure is caused by the interface configuration and not by the

class-of-service configuration, the system does not issue a commit error, but instead issues a syslog message.

When the system issues the syslog message, it programs the default classifier for the interface type on the interface. In this scenario, the interface has been configured as a Layer 3 interface, so the system applies the default DSCP profile to the physical Layer 3 interface.

In this scenario, to install a configured DSCP classifier, you remove the misconfigured classifier from the Layer 3 logical interface and apply it to the Layer 3 physical interface. For example:

```
[edit]
user@switch# delete class-of-service interfaces xe-0/0/20 unit 0 classifiers dscp dscp1
user@switch# commit
user@switch# set class-of-service interfaces xe-0/0/20 classifiers dscp dscp1
user@switch# commit
```

Related Documentation

- [Understanding CoS Packet Flow on page 34](#)
- [Understanding Default CoS Settings on page 37](#)
- [Understanding CoS Classifiers on page 51](#)
- [Understanding Default CoS Scheduling and Classification on page 74](#)
- [Understanding CoS Rewrite Rules on page 154](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 128](#)
- [Example: Configuring Unicast Classifiers on page 273](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 277](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 404](#)
- [Configuring a Global MPLS EXP Classifier](#)
- [Defining CoS Rewrite Rules on page 427](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers](#)
- [Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces \(NP_Ports\) on page 423](#)

Understanding CoS Forwarding Classes

Forwarding classes group traffic and assign the traffic to output queues. Each forwarding class is mapped to an output queue. Classification identifies the output queue for each incoming packet by mapping the packet code point bits to forwarding classes. The forwarding class to queue mapping defines the output queue used for the packet.

A classifier must associate each packet with one of the following five default forwarding classes or with a user-configured forwarding class in order to assign an output queue to the packet:

- **fcoe**—Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic.
- **no-loss**—Guaranteed delivery for TCP lossless traffic.
- **best-effort**—Provides best-effort delivery without a service profile. Loss priority is typically not carried in a class-of-service (CoS) value.
- **network-control**—Supports protocol control and is typically high priority.
- **mcast**—Provides no service profile for multdestination (multicast, broadcast, and destination lookup fail) packets.

The switch supports up to 12 forwarding classes, thus enabling flexible, differentiated, packet classification. For example, you can configure multiple classes of best-effort traffic such as **best-effort**, **best-effort1**, and **best-effort2**.

The switch supports up to 12 output queues: 8 output queues for unicast traffic (queues 0 through 7) and 4 output queues for multdestination traffic (queues 8 through 11). Forwarding classes mapped to unicast queues are associated with unicast traffic, and forwarding classes mapped to multdestination queues are associated with multdestination traffic. You cannot map unicast and multdestination traffic to the same queue. You cannot map a strict-high priority queue to a multdestination forwarding class (queues 8 through 11 do not support strict-high priority configuration).

- [Default Forwarding Classes on page 68](#)
- [Forwarding Class Configuration Rules on page 70](#)
- [Lossless Transport Support on page 71](#)

Default Forwarding Classes

[Table 32 on page 69](#) shows the four default forwarding classes defined for unicast traffic, and [Table 33 on page 69](#) shows the four default forwarding classes defined for multicast traffic.

If desired, you can rename the forwarding classes associated with the queues supported on your switch. Assigning a new class name to an output queue does not alter the default classification or scheduling that is applicable to that queue. CoS configurations can be quite complicated, so unless it is required by your scenario, we recommend that you not alter the default class names or queue number associations.

Table 32: Default Forwarding Classes for Unicast Packets

Forwarding Class Name	Default Queue Mapping	Comments
best-effort (be)	0	<p>The software does not apply any special CoS handling to packets with 000000 in the DiffServ field. This is a backward compatibility feature. These packets are usually dropped under congested network conditions.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of drop.</p>
fcoe	3	<p>By default, the fcoe forwarding class is a lossless forwarding class designed to handle Fibre Channel over Ethernet (FCoE) traffic. The no-loss packet drop attribute is applied by default.</p> <p>NOTE: By convention, deployments with converged server access typically use IEEE 802.1p priority 3 (011) for FCoE traffic. The default mapping of the fcoe forwarding class is to queue 3. Apply priority-based flow control (PFC) to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE requires.</p> <p>We recommend that you use priority 3 for FCoE traffic unless your network architecture requires that you use a different priority.</p>
no-loss	4	<p>By default, this is a lossless forwarding class with a packet drop attribute of no-loss.</p>
network-control (nc)	7	<p>The software delivers packets in this service class with a high priority. (These packets are not delay-sensitive.)</p> <p>Typically, these packets represent routing protocol hello or keepalive messages. Because loss of these packets jeopardizes proper network operation, packet delay is preferable to packet discard.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of drop.</p>

Table 33: Default Forwarding Classes for Multicast Packets

Forwarding Class Name	Default Queue Mapping	Comments
mcast	8	<p>The software does not apply any special CoS handling to the multideestination packets. These packets are usually dropped under congested network conditions.</p> <p>By default, this is a lossy forwarding class with a packet drop attribute of drop.</p>



NOTE: Mirrored traffic is always sent to the queue that corresponds to the multidestination forwarding class. The switched copy of the mirrored traffic is forwarded with the priority determined by the behavior aggregate classification process.

Forwarding Class Configuration Rules

Take the following rules into account when you configure forwarding classes:

- [Queue Assignment Rules on page 70](#)
- [Scheduling Rules on page 70](#)
- [Rewrite Rules on page 71](#)

Queue Assignment Rules

The following rules govern queue assignment:

- CoS configurations that specify more queues than the switch can support are not accepted. The commit operation fails with a detailed message that states the total number of queues available.
- All default CoS configurations are based on queue number. The name of the forwarding class that appears in the default configuration is the forwarding class currently associated with that queue.
- Only unicast forwarding classes can be mapped to unicast queues (0 through 7), and only multidestination forwarding classes can be mapped to multidestination queues (8 through 11).
- Strict-high priority queues cannot be mapped to multidestination forwarding classes. (Strict-high priority traffic cannot be mapped to queues 8 through 11).
- If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

In addition, if you configure a strict-high priority queue, we recommend that you always apply a shaping rate to prevent the strict-high priority queue from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

Scheduling Rules

When you define a forwarding class that is used on the switch (the behavior aggregate classifier has a forwarding class and you expect traffic for the forwarding class), you must also define a scheduling policy for the forwarding class. Defining a scheduling policy means:

- Mapping a scheduler to the forwarding class in a scheduler map
- Including the forwarding class in a forwarding class set
- Associating the scheduler map with a traffic control profile
- Attaching the traffic control profile to a forwarding class set and an interface

Rewrite Rules

On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

Lossless Transport Support

The QFX Series supports up to six lossless forwarding classes. For lossless transport, you must enable PFC on the IEEE 802.1p code point of lossless forwarding classes. The following limitations apply to support lossless transport:

- The external cable length from the QFX3500 or QFabric system Node device to other devices cannot exceed 300 meters.
- The internal cable length from the QFabric system Node device to the QFabric system Interconnect device cannot exceed 150 meters.
- For FCoE traffic, the interface maximum transmission unit (MTU) must be at least 2180 bytes to accommodate the packet payload, headers, and checks.
- Changing any portion of a PFC configuration on a port blocks the entire port until the change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Changing the PFC configuration means any change to a congestion notification profile that is configured on a port (enabling or disabling PFC on a code point, changing the MRU or cable-length value, or specifying an output flow control queue). Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.



.....

NOTE: Junos OS Release 12.2 introduces changes to the way the QFX Series handles lossless forwarding classes (the `fcoe` and `no-loss` forwarding classes).

In Junos OS Release 12.1, both explicitly configuring the `fcoe` and `no-loss` forwarding classes, and using the default configuration for these forwarding classes, resulted in the same lossless behavior for traffic mapped to those forwarding classes.

However, in Junos OS Release 12.2, if you explicitly configure the `fcoe` or the `no-loss` forwarding class, that forwarding class is no longer treated as a lossless forwarding class. Traffic mapped to these forwarding classes is treated as lossy (best-effort) traffic. This is true even if the explicit configuration is exactly the same as the default configuration.

If your CoS configuration from Junos OS Release 12.1 or earlier includes the explicit configuration of the `fcoe` or the `no-loss` forwarding class, then when you upgrade to Junos OS Release 12.2, those forwarding classes are not lossless. To preserve the lossless treatment of these forwarding classes, delete the explicit `fcoe` and `no-loss` forwarding class configuration before you upgrade to Junos OS Release 12.2.

See [“Overview of CoS Changes Introduced in Junos OS Release 12.2” on page 21](#) for detailed information about this change and how to delete an existing lossless configuration.

In Junos OS Release 12.3, the default behavior of the `fcoe` and `no-loss` forwarding classes is the same as in Junos OS Release 12.2. However, in Junos OS Release 12.3, you can configure up to six lossless forwarding classes. All explicitly configured lossless forwarding classes must include the new `no-loss` packet drop attribute or the forwarding class is lossy.

.....

**Related
Documentation**

- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 21](#)
- [Understanding Junos CoS Components on page 30](#)
- [Understanding CoS Packet Flow on page 34](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)
- [Example: Configuring Forwarding Classes on page 283](#)
- [Defining CoS Forwarding Classes on page 408](#)

Understanding CoS Forwarding Class Sets (Priority Groups)

A forwarding class set is the Junos OS configuration construct that equates to a priority group in enhanced transmission selection (ETS, described in IEEE 802.1Qaz). The switch implements ETS using a two-tier hierarchical scheduler.

A priority group is a group of queues (priorities). Mapping a forwarding class to a queue defines the traffic for that queue, so a priority equates to a queue (forwarding class). The queues in a priority group share the port bandwidth allocated to that priority group. The traffic for queues in one priority group usually share similar traffic-handling requirements.

You can configure up to three unicast forwarding class sets and one multicast forwarding class set. Only unicast forwarding classes can belong to unicast forwarding class sets. Only multicast forwarding classes can belong to the multicast forwarding class set.

If you configure a strict-high priority queue, you must observe the following rules when configuring forwarding class sets:

- You must create a separate forwarding class set for the strict-high priority queue.
- Only one forwarding class set can contain strict-high priority queues.
- Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.
- A strict-high priority queue cannot belong to a multidestination forwarding class set.
- You cannot configure a guaranteed minimum bandwidth (guaranteed rate) for a forwarding class set that includes a strict-high priority queue. (You also cannot configure a guaranteed minimum bandwidth for a strict-high queue.)
- We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

You must use hierarchical scheduling to define CoS for output queues. The two-tier hierarchical scheduler defines bandwidth resources for the priority group, and then allocates those resources among the priorities that belong to the priority group.

If you do not explicitly configure forwarding class sets, the system automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The system assigns 100 percent of the port output bandwidth to the default forwarding class set. Ingress traffic is classified based on the default classifier settings. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default scheduler settings. Forwarding classes that are not part of the default scheduler receive no bandwidth. The default priority group is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange Protocol (DCBX) advertisement.

When you explicitly configure forwarding class sets and map them to an interface, any forwarding class that you do not map to a forwarding class set receives no guaranteed bandwidth on that interface. Forwarding classes that belong to the default forwarding class set might receive bandwidth if the other forwarding class sets are not using all of the port bandwidth. However, the amount of bandwidth forwarding classes that are not in explicitly configured forwarding class sets receive is not guaranteed. The bandwidth for the default forwarding class depends on whether extra port bandwidth is available and therefore is not deterministic.

To guarantee bandwidth for forwarding classes in a predictable manner, be sure to map all forwarding classes that you expect to carry traffic on an interface to a forwarding class set and map the forwarding class set to the interface.

**Related
Documentation**

- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Forwarding Class Sets on page 286](#)
- [Defining CoS Forwarding Class Sets on page 410](#)

Understanding Default CoS Scheduling and Classification

If you do not configure hierarchical scheduling on an interface, the switch uses the default classifiers for ingress traffic and the default schedulers for egress traffic. Default scheduling and classification handle all traffic types (best-effort, FCoE, no-loss, network-control, and multidestination traffic).

Hierarchical scheduling groups egress queues (priorities, configured as forwarding classes) into priority groups (forwarding class sets). If you use only the default traffic scheduling and classification, the switch automatically creates a default priority group that contains all of the priorities (forwarding classes, which represent output queues) and assigns 100 percent of the port output bandwidth to that priority group. The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default classifier settings. The default priority group is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange (DCBX) protocol advertisement.



NOTE: If you explicitly configure one or more priority groups on an interface, any forwarding class that is not assigned to a priority group on that interface receives *no bandwidth*. This means that if you configure hierarchical scheduling on an interface, every forwarding class (priority) that you want to forward traffic on that interface must belong to a forwarding class set (priority group).

The following sections describe:

- [Default Classification on page 75](#)
- [Default Scheduling on page 78](#)

- [Default DCBX Advertisement on page 80](#)
- [Default Scheduling and Classification Summary on page 80](#)

Default Classification

The default classifiers assign unicast and multicast best-effort and network-control ingress traffic to forwarding classes and loss priorities. The QFX Series applies default unicast IEEE 802.1, unicast DSCP, and multideestination classifiers to each interface that does not have explicitly configured classifiers. If you explicitly configure one type of classifier but not other types of classifiers, the system uses only the configured classifier and does not use default classifiers for other types of traffic. There are two different default unicast IEEE 802.1 classifiers, a trusted classifier for ports that are in trunk mode or tagged-access mode, and an untrusted classifier for ports that are in access mode.

[Table 34 on page 75](#) shows the default mapping of IEEE 802.1 code-point values to unicast forwarding classes and loss priorities for ports in trunk mode or tagged-access mode.

Table 34: Default IEEE 802.1 Unicast Classifiers for Ports in Trunk Mode or Tagged-Access Mode (Trusted Classifier)

Code Point	Forwarding Class	Loss Priority
be (000)	best-effort	low
be1 (001)	best-effort	low
ef (010)	best-effort	low
ef1 (011)	fcoe	low
af11 (100)	no-loss	low
af12 (101)	best-effort	low
nc1 (110)	network-control	low
nc2 (111)	network-control	low

[Table 35 on page 75](#) shows the default mapping of IEEE 802.1p code-point values to unicast forwarding classes and loss priorities for ports in access mode (all incoming traffic is mapped to best-effort forwarding classes).

Table 35: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier)

Code Point	Forwarding Class	Loss Priority
000	best-effort	low
001	best-effort	low

Table 35: Default IEEE 802.1 Unicast Classifiers for Ports in Access Mode (Untrusted Classifier) (*continued*)

Code Point	Forwarding Class	Loss Priority
010	best-effort	low
011	best-effort	low
100	best-effort	low
101	best-effort	low
110	best-effort	low
111	best-effort	low

Table 36 on page 76 shows the default mapping of IEEE 802.1 code-point values to multdestination (multicast, broadcast, and destination lookup fail traffic) forwarding classes and loss priorities.

Table 36: Default IEEE 802.1 Multidestination Classifiers

Code Point	Forwarding Class	Loss Priority
be (000)	mcast	low
be1 (001)	mcast	low
ef (010)	mcast	low
ef1 (011)	mcast	low
af11 (100)	mcast	low
af12 (101)	mcast	low
nc1 (110)	mcast	low
nc2 (111)	mcast	low

Table 37 on page 76 shows the default mapping of DSCP code-point values to unicast forwarding classes and loss priorities for DSCP IP and DCSP IPv6.

Table 37: Default DSCP IP and IPv6 Unicast Classifiers

Code Point	Forwarding Class	Loss Priority
ef (101110)	best-effort	low
af11 (001010)	best-effort	low

Table 37: Default DSCP IP and IPv6 Unicast Classifiers (*continued*)

Code Point	Forwarding Class	Loss Priority
af12 (001100)	best-effort	low
af13 (001110)	best-effort	low
af21 (010010)	best-effort	low
af22 (010100)	best-effort	low
af23 (010110)	best-effort	low
af31 (011010)	best-effort	low
af32 (011100)	best-effort	low
af33 (011110)	best-effort	low
af41 (100010)	best-effort	low
af42 (100100)	best-effort	low
af43 (100110)	best-effort	low
be (000000)	best-effort	low
cs1 (001000)	best-effort	low
cs2 (010000)	best-effort	low
cs3 (011000)	best-effort	low
cs4 (100000)	best-effort	low
cs5 (101000)	best-effort	low
nc1 (110000)	network-control	low
nc2 (111000)	network-control	low



NOTE: There are no default DSCP IP or IPv6 multdestination classifiers for multdestination traffic. DSCP IPv6 multdestination classifiers are not supported for multdestination traffic.

Default Scheduling

The default schedulers allocate egress bandwidth resources to unicast and multicast egress traffic as shown in [Table 38 on page 78](#):

Table 38: Default Scheduler Configuration

Default Scheduler and Queue Number	Transmit Rate (Minimum Guaranteed Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Bandwidth Sharing	Priority	Buffer Size
Best-effort scheduler (queue 0)	5%	None	5%	low	5%
FCoE scheduler (queue 3)	35%	None	35%	low	35%
No-loss scheduler (queue 4)	35%	None	35%	low	35%
Network-control scheduler (queue 7)	5%	None	5%	low	5%
Multidestination scheduler (queue 8)	20%	None	20%	low	20%



NOTE: The minimum guaranteed bandwidth rate also determines the amount of excess (extra) bandwidth that the queue can share. Extra bandwidth is allocated to queues in proportion to the minimum guaranteed bandwidth (transmit rate) of each queue.

By default, only the five default schedulers shown in [Table 38 on page 78](#) have traffic mapped to them. Only the queues associated with the default schedulers receive default bandwidth, based on the default scheduler transmit rate. (You can configure schedulers and forwarding classes to allocate bandwidth to other queues or to change the default bandwidth of a default queue.) In addition, multidestination queue 11 receives enough bandwidth from the default multidestination scheduler to handle CPU-generated multidestination traffic.

Default hierarchical scheduling divides the total port bandwidth between two groups of traffic: unicast traffic and multidestination traffic. By default, unicast traffic consists of queue 0 (**best-effort** forwarding class), queue 3 (**fcoe** forwarding class), queue 4 (**no-loss** forwarding class), and queue 7 (**network-control** forwarding class). Unicast traffic receives and shares a total of 80 percent of the port bandwidth. By default, multidestination traffic (**mcast** queue 8) receives a total of 20 percent of the port bandwidth. So on a 10-Gigabit port, unicast traffic receives 8-Gbps of bandwidth and multidestination traffic receives 2-Gbps of bandwidth.



NOTE: Multidestination queue 11 also receives a small amount of default bandwidth from the multidestination scheduler. CPU-generated multidestination traffic uses queue 11, so you might see a small number of packets egress from queue 11. In addition, in the unlikely case that firewall filter match conditions map multidestination traffic to a unicast forwarding class, that traffic uses queue 11.

Default scheduling uses weighted round-robin (WRR) scheduling. Each queue receives a portion (weight) of the total available interface bandwidth. The scheduling weight is based on the transmit rate of the default scheduler for that queue. For example, queue 7 receives a default scheduling weight of 5 percent of the available bandwidth, and queue 4 receives a default scheduling weight of 35 percent of the available bandwidth. Queues are mapped to forwarding classes, so forwarding classes receive the default bandwidth for the queues to which they are mapped.

You should explicitly map traffic to non-default (unconfigured) queues and create schedulers to allocate bandwidth to those queues if you want to use them to forward traffic. By default, unicast queues 1, 2, 5, and 6 are unconfigured, and multidestination queues 9, 10, and 11 are unconfigured. Unconfigured queues have a default scheduling weight of 1 so that they can receive a small amount of bandwidth in case they need to forward traffic. (However, queue 11 can use more of the default multidestination scheduler bandwidth if necessary to handle CPU-generated multidestination traffic.)



NOTE: All four multidestination queues have a scheduling weight of 1. Because by default multidestination traffic goes to queue 8, queue 8 receives almost all of the multidestination bandwidth. (There is no traffic on queue 9 and queue 10, and very little traffic on queue 11, so there is almost no competition for multidestination bandwidth.)

However, if you explicitly configure queue 9, 10, or 11 (by mapping code points to the unconfigured multidestination forwarding classes using the multidestination classifier), the explicitly configured queues share the multidestination scheduler bandwidth equally with default queue 8, because all of the queues have the same scheduling weight (1). To ensure that multidestination bandwidth is allocated to each queue properly and that the bandwidth allocation to the default queue (8) is not reduced too much, we strongly recommend that you configure a scheduler if you explicitly classify traffic into queue 9, 10, or 11.

If you map traffic to an unconfigured queue, the queue receives only the amount of group bandwidth proportional to its default weight (1). The actual amount of bandwidth an unconfigured queue receives depends on how much bandwidth the other queues in the group are using.

If the other unicast queues use less than their allocated amount of bandwidth, the unconfigured queues can share the unused bandwidth. Sharing unused bandwidth is one of the key advantages of hierarchical port scheduling. Configured queues have higher

priority for bandwidth than unconfigured queues, so if a configured queue needs more bandwidth, then less bandwidth is available for unconfigured queues. Unconfigured queues always receive a minimum amount of bandwidth based on their scheduling weight (1). If you map traffic to an unconfigured queue, to allocate bandwidth to that queue, configure a scheduler for the forwarding class that is mapped to the queue.

Default DCBX Advertisement

When you configure hierarchical scheduling on an interface, DCBX advertises each priority group, the priorities in each priority group, and the bandwidth properties of each priority and priority group.

If you do not configure hierarchical scheduling on an interface, DCBX advertises the automatically created default priority group and its priorities. DCBX also advertises the default bandwidth allocation of the priority group, which is 100 percent of the port bandwidth.

Default Scheduling and Classification Summary

If you do not configure hierarchical scheduling on an interface:

- Default classifiers classify ingress traffic.
- Default schedulers schedule egress traffic.
- DCBX advertises a single default priority group with 100 percent of the port bandwidth allocated to that priority group. All priorities (forwarding classes) are assigned to the default priority group and receive bandwidth based on their default schedulers. The default priority group is generated automatically and is not user-configurable.

Related Documentation

- [Understanding CoS Packet Flow on page 34](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Understanding Default CoS Settings on page 37](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57](#)
- [Understanding DCB Features and Requirements on page 156](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 203](#)
- [Example: Configuring Unicast Classifiers on page 273](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 277](#)
- [Example: Configuring Queue Schedulers on page 289](#)

Understanding CoS Hierarchical Port Scheduling (ETS)

Scheduling defines the class-of-service (CoS) properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the queue priority, and the drop profiles associated with the queue.

Hierarchical port scheduling is a two-tier process that provides better port bandwidth utilization and greater flexibility to allocate resources to queues and to groups of queues. Hierarchical scheduling includes the Junos OS implementation of enhanced transmission selection (ETS, described in IEEE 802.1Qaz).



Video: [What is Enhanced Transmission Selection?](#)

This topic describes:

- [Hierarchical Scheduling Tiers on page 81](#)
- [Hierarchical Scheduling and ETS on page 82](#)
- [ETS Advertisement in DCBX on page 83](#)
- [Hierarchical Scheduling Process on page 83](#)
- [Strict-High Priority Queues and Hierarchical Scheduling on page 84](#)
- [Default Hierarchical Scheduling on page 85](#)

Hierarchical Scheduling Tiers

The two tiers used in hierarchical scheduling are priorities and priority groups, as shown in [Table 39 on page 81](#).

Table 39: Hierarchical Scheduling Tiers

Junos OS Configuration Construct	Equivalent ETS Construct	Description
Forwarding class	Priority	<p>Think about priorities (forwarding classes) as output queues. You map forwarding classes to queues, so each forwarding class is in essence an output queue.</p> <p>When you use a classifier to map a forwarding class to an IEEE 802.1p code point, the code point identifies that traffic's priority for priority-based flow control (PFC). Thus the forwarding class, the queue mapped to the forwarding class, and the priority mapped to the forwarding class all identify the same traffic.</p>
Forwarding class set	Priority group	<p>Priority groups (forwarding class sets) are groups of priorities. Forwarding class membership in a forwarding class set defines the priority group to which each priority belongs.</p> <p>You can configure up to three unicast priority groups and one multicast forwarding class set.</p>



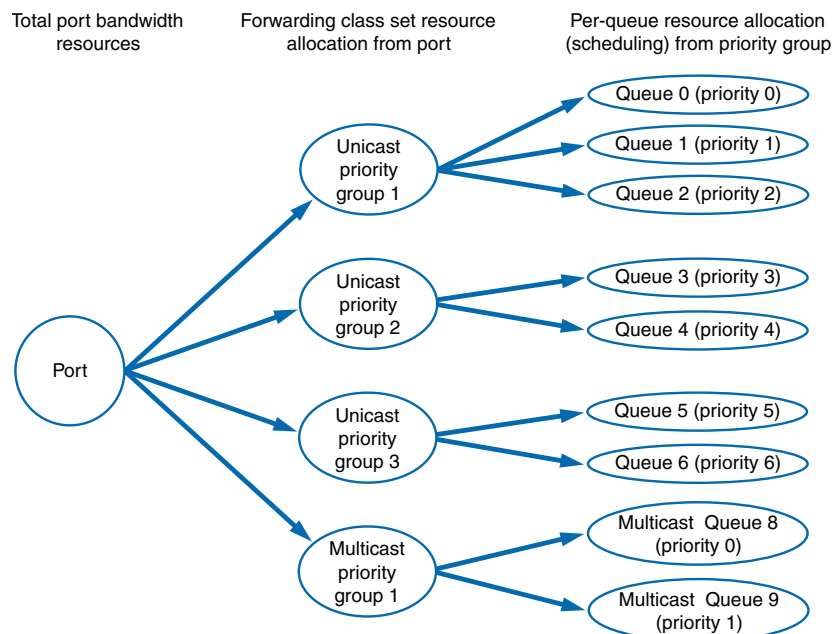
NOTE: If you explicitly configure one or more priority groups on an interface, any priority that is not assigned to a priority group on that interface is assigned to an automatically generated default priority group and receives *no bandwidth*. This means that if you configure hierarchical scheduling on an interface, every forwarding class that you want to forward traffic on that interface must belong to a forwarding class set.

Hierarchical Scheduling and ETS

Two-tier hierarchical scheduling enables you to manage bandwidth efficiently by enabling you to define the CoS properties for each priority group and for each priority. One tier of the hierarchical scheduler allocates port bandwidth to a priority group. The other tier of the hierarchical scheduler determines the portion of the priority group bandwidth that a queue can use.

The CoS properties you configure for a priority group define the port bandwidth resources available to the queues in that priority group. The CoS properties you configure for each queue specify the portion or percentage of the total bandwidth configured for the priority group that is available to the queue. [Figure 5 on page 82](#) shows the relationship of port resource allocation to priority groups and priority group resource allocation to queues (priorities).

Figure 5: Hierarchical Scheduling Tiers



g040722

If a queue is not using its allocated bandwidth, ETS shares the unused bandwidth among the other queues in the priority group. If link bandwidth is available or if a priority group on a link is not using its allocated bandwidth, ETS shares the unused bandwidth with other priority groups on the link. In this way ETS improves link bandwidth utilization and provides each queue with the maximum bandwidth. Priorities that consist of bursty traffic

can share bandwidth during periods of low traffic transmission instead of consuming their entire bandwidth allocation when traffic loads are light.



NOTE: The available link bandwidth is the bandwidth remaining after servicing strict-high priority flows.

ETS Advertisement in DCBX

When you configure hierarchical scheduling on a port, Data Center Bridging Capability Exchange Protocol (DCBX) advertises:

- Each priority group
- The priorities in each priority group
- The bandwidth properties of each priority group and priority

When you configure hierarchical scheduling on a port, any priority that is not part of an explicitly configured priority group is assigned to the automatically generated default priority group and receives no bandwidth. The default priority group is transparent. It does not appear in the configuration.

Hierarchical Scheduling Process

Hierarchical scheduling consists of multiple configuration steps that create the priorities and the priority groups, schedule their resources, and assign them to interfaces. The steps below correspond to the six blocks in the packet flow diagram shown in [Figure 6 on page 84](#):

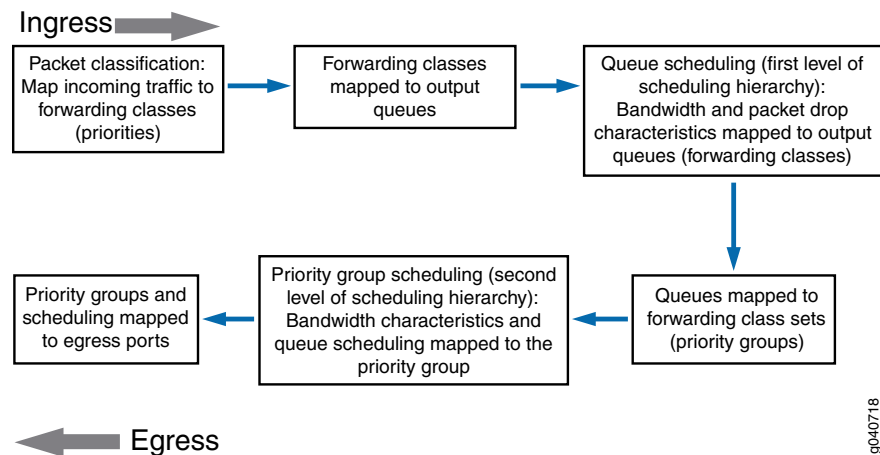
1. Packet classification:
 - Classify incoming traffic into priorities. This consists of either using the default classifiers or configuring classifiers to map IEEE 802.1p code points and loss priorities to the forwarding classes.
 - Apply the classifiers to ingress interfaces. This groups incoming traffic into forwarding classes (priorities) by mapping code points to forwarding classes and loss priorities on the specified interface.
2. Configure the output queues for the forwarding classes (priorities). This consists of either using the default forwarding classes and forwarding-class-to-queue mapping or creating your own forwarding classes and mapping them to queues.
3. Allocate resources to the forwarding classes:
 - Define resources for the priorities. This consists of configuring schedulers to set minimum guaranteed bandwidth, maximum bandwidth, drop profiles for Weighted Random Early Detection (WRED), and bandwidth priority to apply to a forwarding class. Extra bandwidth is shared among queues in proportion to the minimum guaranteed bandwidth of each queue.
 - Map resources to priorities. This consists of mapping forwarding classes to schedulers by using a scheduler map.

4. Configure priority groups. This consists of mapping forwarding classes (priorities) to forwarding class sets (priority groups) to define the priorities that belong to each priority group.
5. Define resources for the priority groups. This consists of configuring traffic control profiles to set minimum guaranteed bandwidth and maximum bandwidth for a priority group. Traffic control profiles also specify a scheduler map, which defines the resources (schedulers) for the priorities in the priority group. Extra port bandwidth is shared among priority groups in proportion to the minimum guaranteed bandwidth of each priority group.

The traffic control profile bandwidth settings determine the port resources available to the priority group, and the schedulers specified in the scheduler map determine the amount of the priority group resources that each priority receives.

6. Apply the hierarchical scheduling to a port. This consists of attaching one or more priority groups to a port interface. For each priority group, you also attach a traffic control profile. Different priority groups on the same port can use different traffic control profiles.

Figure 6: Hierarchical Scheduling Packet Flow



Strict-High Priority Queues and Hierarchical Scheduling

If you configure a strict-high priority queue, you must observe the following rules:

- You must create a separate forwarding class set (priority group) for the strict-high priority queue.
- Only one forwarding class set can contain strict-high priority queues.
- Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.
- A strict-high priority queue cannot belong to a multidestination forwarding class set.
- We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high

priority queue can use all of the available port bandwidth and starve other queues on the port.



NOTE: On a QFabric system, if a fabric (fte) interface handles strict-high priority traffic, you must define a separate fc-set (priority group) for strict-high priority traffic. Strict-high priority traffic cannot be mixed with traffic of other priorities in an fc-set. For example, you might choose to create different fc-sets for best effort, lossless, strict-high priority, and multideestination traffic.

Default Hierarchical Scheduling

If you do not explicitly configure hierarchical scheduling, the switch uses the default settings:

- The switch automatically creates a default forwarding class set that contains all of the forwarding classes on the switch. The switch assigns 100 percent of the port output bandwidth to the default forwarding class set. The default forwarding class set is transparent. It does not appear in the configuration and is used for Data Center Bridging Capability Exchange Protocol (DCBX) advertisement.
- Ingress traffic is classified based on the default classifier settings.
- The forwarding classes (queues) in the default forwarding class set receive bandwidth based on the default scheduler settings.

Related Documentation

- [Understanding CoS Packet Flow on page 34](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Understanding CoS Priority Group Scheduling on page 95](#)
- [*Benefits of Configuring CoS Hierarchical Port Scheduling*](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)
- [Understanding CoS Classifiers on page 51](#)
- [Understanding Default CoS Scheduling and Classification on page 74](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 199](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 203](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 300](#)
- [Example: Configuring Maximum Output Bandwidth on page 305](#)

Understanding CoS Output Queue Schedulers

Output queue scheduling defines the class-of-service (CoS) properties of output queues (output queues are mapped to forwarding classes and IEEE 802.1p priorities). Queue scheduling works with priority group scheduling to create a two-tier hierarchical scheduler. The hierarchical scheduler allocates port bandwidth to a group of queues called a priority group (forwarding class set), and queue scheduling determines the portion of the priority group's bandwidth that a particular queue can use.

Scheduler maps associate queue schedulers with forwarding classes, which are mapped to output queues. You can associate each scheduler map with a traffic control profile, and then associate each traffic control profile with a forwarding class set (priority group) and a port interface. In conjunction with the priority group scheduling configured in the traffic control profile, queue scheduling configures the output queues, packet schedulers, and weighted random early detection (WRED) packet drop processes that operate according to this mapping.



NOTE: When you configure bandwidth for a queue (forwarding class) or a priority group (forwarding class set), the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.

- [Output Queue Scheduling Components on page 86](#)
- [Default Schedulers on page 87](#)
- [Transmit Rate \(Minimum Guaranteed Bandwidth\) on page 90](#)
- [Sharing Extra Bandwidth on page 90](#)
- [Shaping Rate \(Maximum Bandwidth\) on page 91](#)
- [Scheduling Priority on page 91](#)
- [Scheduler Drop-Profile Maps on page 92](#)
- [Buffer Size on page 92](#)
- [Explicit Congestion Notification on page 93](#)
- [Scheduler Maps on page 94](#)

Output Queue Scheduling Components

[Table 40 on page 87](#) provides a quick reference to the scheduler components you can configure to determine the bandwidth properties of output queues, and [Table 41 on page 87](#) provides a quick reference to some related scheduling configuration components.

Table 40: Output Queue Scheduler Components

Output Queue Scheduler Component	Description
Buffer size	Sets the size of the queue buffer.
Drop profile map	Maps a drop profile to a loss priority. Drop profile map components include: <ul style="list-style-type: none"> Drop profile—Sets the probability of dropping packets as the queue fills up. Loss priority—Sets the traffic loss priority to which a drop profile applies.
Explicit congestion notification	Enables explicit congestion notification (ECN) on the queue.
Priority	Sets the scheduling priority applied to the queue.
Shaping rate	Sets the maximum bandwidth the queue can consume.
Transmit rate	Sets the minimum guaranteed bandwidth for the queue. Extra bandwidth is shared among queues in proportion to the minimum guaranteed bandwidth of each queue.

Table 41: Other Scheduling Components

Other Scheduling Components	Description
Scheduler map	Maps schedulers to forwarding classes (forwarding classes are mapped to queues, so a forwarding class represents a queue)
Forwarding class	Maps traffic to a queue. Classifiers map forwarding classes to IEEE 802.1p priorities. A forwarding class, an output queue, and an IEEE 802.1p priority are mapped to each other and identify the same traffic. (The IEEE 802.1p priority identifies incoming traffic, which is classified into the forwarding class, and the forwarding class is in turn mapped to an output queue for the traffic.)
Traffic control profile	Configures scheduling for the forwarding class set (priority group), and associates a scheduler map with the forwarding class set to apply queue scheduling to the forwarding classes in the forwarding class set. Extra port bandwidth is shared among forwarding class sets in proportion to the minimum guaranteed bandwidth of each forwarding class set.
Forwarding class set	Name of a priority group. You map forwarding classes to forwarding class sets. A forwarding class set consists of one or more forwarding classes.

Default Schedulers

Each forwarding class requires an associated scheduler. The default configuration uses only five forwarding classes, unicast best-effort (queue 0), fcoe (queue 3), no-loss (queue 4), network-control (queue 7), and multidestination (queue 8). You can use the default schedulers or you can define new schedulers for these five forwarding classes. For any other forwarding class, you must explicitly configure a scheduler.

[Table 42 on page 88](#) shows the default schedulers.

Table 42: Default Schedulers

Default Scheduler and Queue Number	Guaranteed Rate (Minimum Bandwidth)	Shaping Rate (Maximum Bandwidth)	Excess Bandwidth Sharing	Priority	Buffer Size
Best-effort scheduler (queue 0)	5%	None	5%	Low	5%
FCoE scheduler (queue 3)	35%	None	35%	Low	35%
No-loss scheduler (queue 4)	35%	None	35%	Low	35%
Network-control scheduler (queue 7)	5%	None	5%	Low	5%
Multidestination scheduler (queue 8)	20%	None	20%	Low	20%



NOTE: The minimum guaranteed bandwidth rate also determines the amount of excess (extra) bandwidth that the queue can share. Extra bandwidth is allocated to queues in proportion to the minimum guaranteed bandwidth rate of each queue.

By default, only the five default schedulers shown in [Table 38 on page 78](#) have traffic mapped to them. Only the queues associated with the default schedulers receive default bandwidth, based on the default scheduler transmit rate. (You can configure schedulers and forwarding classes to allocate bandwidth to other queues or to change the default bandwidth of a default queue.) In addition, multidestination queue 11 receives enough bandwidth from the default multidestination scheduler to handle CPU-generated multidestination traffic.

Default hierarchical scheduling divides the total port bandwidth between two groups of traffic: unicast traffic and multidestination traffic. By default, unicast traffic consists of queue 0 (**best-effort** forwarding class), queue 3 (**fcoe** forwarding class), queue 4 (**no-loss** forwarding class), and queue 7 (**network-control** forwarding class). Unicast traffic receives and shares a total of 80 percent of the port bandwidth. By default, multidestination traffic (**mcast** queue 8) receives a total of 20 percent of the port bandwidth. So on a 10-Gigabit port, unicast traffic receives 8-Gbps of bandwidth and multidestination traffic receives 2-Gbps of bandwidth.



NOTE: Multidestination queue 11 also receives a small amount of default bandwidth from the multidestination scheduler. CPU-generated multidestination traffic uses queue 11, so you might see a small number of packets egress from queue 11. In addition, in the unlikely case that firewall filter match conditions map multidestination traffic to a unicast forwarding class, that traffic uses queue 11.

Default scheduling uses weighted round-robin (WRR) scheduling. Each queue receives a portion (weight) of the total available interface bandwidth. The scheduling weight is based on the transmit rate of the default scheduler for that queue. For example, queue 7 receives a default scheduling weight of 5 percent of the available bandwidth, and queue 4 receives a default scheduling weight of 35 percent of the available bandwidth. Queues are mapped to forwarding classes, so forwarding classes receive the default bandwidth for the queues to which they are mapped.

You should explicitly map traffic to non-default (unconfigured) queues if you want to use them to forward traffic. By default, unicast queues 1, 2, 5, and 6 are unconfigured, and multidestination queues 9, 10, and 11 are unconfigured. Unconfigured queues have a default scheduling weight of 1 so that they can receive a small amount of bandwidth in case they need to forward traffic. (However, queue 11 can use more of the default multidestination scheduler bandwidth if necessary to handle CPU-generated multidestination traffic.)



NOTE: All four multidestination queues have a scheduling weight of 1. Because by default multidestination traffic goes to queue 8, queue 8 receives almost all of the multidestination bandwidth. (There is no traffic on queue 9 and queue 10, and very little traffic on queue 11, so there is almost no competition for multidestination bandwidth.)

However, if you explicitly configure queue 9, 10, or 11 (by mapping code points to the unconfigured multidestination forwarding classes using the multidestination classifier), the explicitly configured queues share the multidestination scheduler bandwidth equally with default queue 8, because all of the queues have the same scheduling weight (1). To ensure that multidestination bandwidth is allocated to each queue properly and that the bandwidth allocation to the default queue (8) is not reduced too much, we strongly recommend that you configure a scheduler if you explicitly classify traffic into queue 9, 10, or 11.

If you map traffic to an unconfigured queue, the queue receives only the amount of group bandwidth proportional to its default weight (1). The actual amount of bandwidth an unconfigured queue receives depends on how much bandwidth the other queues in the group are using.

If the other unicast queues use less than their allocated amount of bandwidth, the unconfigured queues can share the unused bandwidth. Sharing unused bandwidth is one of the key advantages of hierarchical port scheduling. Configured queues have higher priority for bandwidth than unconfigured queues, so if a configured queue needs more bandwidth, then less bandwidth is available for unconfigured queues. Unconfigured queues always receive a minimum amount of bandwidth based on their scheduling weight (1). If you map traffic to an unconfigured queue, to allocate bandwidth to that queue, configure a scheduler for the forwarding class that is mapped to the queue.

Transmit Rate (Minimum Guaranteed Bandwidth)

The transmit rate determines the minimum guaranteed bandwidth for each forwarding class. It also determines how much excess (extra) bandwidth each low-priority queue can share; each queue shares extra bandwidth in proportion to its transmit rate. You specify the rate in bits per second as a fixed value such as 1 Mbps or as a percentage of the total forwarding class set minimum guaranteed bandwidth (the guaranteed rate set in the traffic control profile). Either the default scheduler or a scheduler you configure allocates a portion of the outgoing interface bandwidth to each forwarding class.



NOTE: For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.

You cannot configure a transmit rate for strict-high priority queues. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.

The allocated bandwidth can exceed the configured minimum rate if additional bandwidth is available from other queues in the forwarding class set. In case of congestion, the configured transmit rate is guaranteed for the queue. This property enables you to ensure that each queue receives the amount of bandwidth appropriate to its level of service.



NOTE: Configuring the minimum guaranteed bandwidth (transmit rate) for a forwarding class does not work unless you also configure the minimum guaranteed bandwidth (guaranteed rate) for the forwarding class set in the traffic control profile.

Additionally, the sum of the transmit rates of the queues in a forwarding class set should not exceed the guaranteed rate for the forwarding class set. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

Sharing Extra Bandwidth

Extra bandwidth is available to low-priority queues when the minimum guaranteed bandwidth of the queues does not use the full amount of forwarding class set bandwidth. This extra bandwidth is shared among the forwarding classes in the set based on the minimum guaranteed bandwidth of each queue.

For example, in a forwarding class set, Queue A has a transmit rate of 1 Gbps, Queue B has a transmit rate of 1 Gbps, and Queue C has a transmit rate of 2 Gbps. After servicing the minimum guaranteed bandwidth of these queues, the forwarding class set has an extra 2 Gbps of bandwidth available, and all three queues still have packets to forward. The queues receive the extra bandwidth in proportion to their transmit rates, so Queue A

receives an extra 500 Mbps, Queue B receives an extra 500 Mbps, and Queue C receives an extra 1 Gbps.

Shaping Rate (Maximum Bandwidth)

The shaping rate determines the maximum bandwidth each forwarding class can consume. You specify the rate in bits per second as a fixed value such as 3 Mbps or as a percentage of the total forwarding class set maximum bandwidth (the shaping rate set in the traffic control profile).

The maximum bandwidth for a queue depends on the total bandwidth available to the forwarding class set to which the queue belongs and how much bandwidth the other queues in the forwarding class set consume.



NOTE: On QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets. A strict-high priority queue (or several queues with higher priorities than the starved queue) can consume all of the port bandwidth and prevent another queue from transmitting packets. To prevent a queue from being starved for bandwidth, you can configure a shaping rate on the queue or queues to prevent them from consuming all of the port bandwidth.



NOTE: On QFabric systems and QFX switches, we recommend that you always configure a shaping rate in the scheduler for strict-high priority queues to prevent them from starving other queues.

Scheduling Priority

Scheduling priority determines the order in which an output interface transmits traffic from the queues, thus ensuring that queues containing important traffic receive better access to the outgoing interface. The priority setting in the scheduler determines the priority for the queue.

Two levels of scheduling priority are supported:

- **Low**—Low-priority queues transmit traffic based on the weighted round robin (WRR) algorithm. The scheduler first determines if an individual queue is within its defined bandwidth profile. The scheduler then regularly reevaluates whether each individual queue is within its defined bandwidth profile and compares the amount of data the queue transmits to the amount of bandwidth the scheduler allocates to the queue. When the transmitted amount is less than the allocated amount, the queue is considered to be in profile. A queue is out of profile when its transmitted amount is larger than its allocated amount. Out of profile queue data is transmitted only if bandwidth is available. Otherwise, it is buffered if buffer space is available. If no buffer space is available, the traffic may be dropped.
- **Strict-high**—You can configure only one queue as **strict-high** priority. The other 11 queues are **low** priority.

The **strict-high** priority queue receives preferential treatment over the low-priority queues. The **strict-high** priority queue receives all of its configured bandwidth before low-priority queues are serviced. Low-priority queues do not transmit traffic until the strict-high priority queue is empty. Carefully consider how much bandwidth you want to allocate to the **strict-high** priority queue to avoid starving the low-priority queues.

If you configure a strict-high priority queue, you must observe the following rules:

- You must create a separate forwarding class set (priority group) for the strict-high priority queue.
- Only one forwarding class set can contain strict-high priority queues.
- Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.
- A strict-high priority queue cannot belong to a multidestination forwarding class set.
- You cannot configure a minimum guaranteed bandwidth for a strict-high priority queue. (You cannot configure a transmit rate for a strict-high priority queue scheduler, and you cannot configure a guaranteed rate for a forwarding class set that has a strict-high priority queue.)
- We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

Junos OS performs priority queueing using the following steps:

1. Services the strict-high priority queue before any other queues are served
2. Services the minimum bandwidth (transmit rate) of low-priority queues until the minimum is met or the queues are empty
3. Services all other low-priority queues and needs that exceed the minimum bandwidth

Scheduler Drop-Profile Maps

Drop-profile maps associate drop profiles with a scheduler. A drop-profile map sets the drop profile for a specific packet loss priority (PLP) and protocol type:

- PLP—Low, medium-high, high. You configure the PLP during classifier configuration. When you use a scheduler map to associate a forwarding class with a scheduler, you can use a drop-profile map to map different drop profiles to the forwarding class for different PLPs.
- Protocol type—Drop profiles match all protocol types.

Buffer Size

Most of the total system buffer space is divided into two buffer pools, shared buffers and dedicated buffers. Shared buffers are a global pool that the ports share dynamically as needed. Dedicated buffers are a reserved portion of the buffer pool that is distributed

evenly to all of the ports. Each port receives an equal allocation of dedicated buffer space. The dedicated buffer allocation to ports is not configurable because it is reserved for the ports.

The queue buffers are allocated from the dedicated buffer pool assigned to the port. By default, ports divide their allocation of dedicated buffers among the egress queues in the same proportion as the default scheduler sets the minimum guaranteed transmission rates (**transmit-rate**) for traffic. Only the queues included in the default scheduler receive dedicated buffers.

If you do not use the default configuration, you can explicitly configure the queue buffer size in either of two ways:

- As a percentage—The queue receives the specified percentage of dedicated port buffers when the queue is mapped to the scheduler and the scheduler is mapped to a port.
- As a remainder—After the port services the queues that have an explicit percentage buffer size configuration, the remaining port dedicated buffer space is divided equally among the other queues to which a scheduler is attached. (No default or explicit scheduler means no dedicated buffer allocation for the queue.) If you configure a scheduler and you do not specify a buffer size as a percentage, *remainder* is the default setting.



NOTE: The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.

For a complete discussion about queue buffer configuration in the context of ingress and egress port buffer configuration, see [“Understanding CoS Buffer Configuration” on page 131](#).

Explicit Congestion Notification

ECN enables end-to-end congestion notification between two endpoints on TCP/IP based networks. The two endpoints are an ECN-enabled sender and an ECN-enabled receiver. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. Any device in the transmission path that does not support ECN breaks the end-to-end ECN functionality. ECN notifies networks about congestion with the goal of reducing packet loss and delay by making the sending device decrease the transmission rate until the congestion clears, without dropping packets.

ECN is disabled by default. Normally, you enable ECN only on queues that handle best-effort traffic because other traffic types use different methods of congestion notification—lossless traffic uses priority-based flow control (PFC) and strict-high priority traffic receives all of the port bandwidth it requires up to the point of a configured maximum rate.

Scheduler Maps

A scheduler map associates a specified forwarding class with a scheduler configuration. After configuring a scheduler, you must include it in a scheduler map, associate the scheduler map with a traffic control profile, and then associate the traffic control profile with an interface and a forwarding class set.

You can associate up to four user-defined scheduler maps with traffic control profiles.

Related Documentation

- [Understanding Junos CoS Components on page 30](#)
- [Understanding CoS Priority Group Scheduling on page 95](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Understanding CoS Buffer Configuration on page 131](#)
- [*Understanding CoS Explicit Congestion Notification*](#)
- [Understanding CoS Scheduling Behavior and Configuration Considerations on page 104](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 199](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 203](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 300](#)
- [Example: Configuring Maximum Output Bandwidth on page 305](#)
- [Example: Configuring Queue Scheduling Priority on page 295](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
- [Example: Configuring WRED Drop Profiles on page 279](#)
- [Example: Configuring Drop Profile Maps on page 281](#)
- [*Example: Configuring ECN*](#)

Understanding CoS Priority Group Scheduling

Priority group scheduling defines the class-of-service (CoS) properties of a group of output queues (priorities). Priority group scheduling works with output queue scheduling to create a two-tier hierarchical scheduler. The hierarchical scheduler allocates bandwidth to a group of queues (a priority group, called a forwarding class set in Junos OS configuration). Queue scheduling determines the portion of the priority group bandwidth that the particular queue can use.

You configure priority group scheduling in a traffic control profile and then associate the traffic control profile with a forwarding class set and an interface. You attach a scheduler map to the traffic control profile to specify the queue scheduling characteristics.



NOTE: When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.

- [Priority Group Scheduling Components on page 95](#)
- [Default Traffic Control Profile on page 96](#)
- [Guaranteed Rate \(Minimum Guaranteed Bandwidth\) on page 96](#)
- [Sharing Extra Bandwidth on page 96](#)
- [Shaping Rate \(Maximum Bandwidth\) on page 97](#)
- [Scheduler Maps on page 97](#)

Priority Group Scheduling Components

[Table 43 on page 95](#) provides a quick reference to the traffic control profile components you can configure to determine the bandwidth properties of priority groups, and [Table 44 on page 96](#) provides a quick reference to some related scheduling configuration components.

Table 43: Priority Group Scheduler Components

Traffic Control Profile Component	Description
Guaranteed rate	Sets the minimum guaranteed port bandwidth for the priority group. Extra port bandwidth is shared among priority groups in proportion to the guaranteed rate of each priority group on the port.
Shaping rate	Sets the maximum port bandwidth the priority group can consume.
Scheduler map	Maps schedulers to queues (forwarding classes, also called priorities). This determines the portion of the priority group bandwidth that a queue receives.

Table 44: Other Scheduling Components

Other Scheduling Components	Description
Forwarding class	Maps traffic to a queue (priority).
Forwarding class set	Name of a priority group. You map forwarding classes to priority groups. A forwarding class set consists of one or more forwarding classes.
Scheduler	Sets the bandwidth and scheduling priority of individual queues (forwarding classes).

Default Traffic Control Profile

There is no default traffic control profile.

Guaranteed Rate (Minimum Guaranteed Bandwidth)

The guaranteed rate determines the minimum guaranteed bandwidth for each priority group. It also determines how much excess (extra) port bandwidth the priority group can share; each priority group shares extra port bandwidth in proportion to its guaranteed rate. You specify the rate in bits per second as a fixed value such as 3 Mbps or as a percentage of the total port bandwidth.

The minimum transmission bandwidth can exceed the configured rate if additional bandwidth is available from other priority groups on the port. In case of congestion, the configured guaranteed rate is guaranteed for the priority group. This property enables you to ensure that each priority group receives the amount of bandwidth appropriate to its level of service.



NOTE: Configuring the minimum guaranteed bandwidth (transmit rate) for a forwarding class does not work unless you also configure the minimum guaranteed bandwidth (guaranteed rate) for the forwarding class set in the traffic control profile.

Additionally, the sum of the transmit rates of the queues in a forwarding class set should not exceed the guaranteed rate for the forwarding class set. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

You cannot configure a guaranteed rate for forwarding class sets that include strict-high priority queues.

Sharing Extra Bandwidth

Extra bandwidth is available to priority groups when the priority groups do not use the full amount of available port bandwidth. This extra port bandwidth is shared among the priority groups based on the minimum guaranteed bandwidth of each priority group.

For example, Port A has three priority groups: fc-set-1, fc-set-2, and fc-set-3. Fc-set-1 has a guaranteed rate of 2 Gbps, fc-set-2 has a guaranteed rate of 2 Gbps, and fc-set-3 has a guaranteed rate of 4 Gbps. After servicing the minimum guaranteed bandwidth of these priority groups, the port has an extra 2 Gbps of available bandwidth, and all three priority groups have still have packets to forward. The priority groups receive the extra bandwidth in proportion to their guaranteed rates, so fc-set-1 receives an extra 500 Mbps, fc-set-2 receives an extra 500 Mbps, and fc-set-3 receives an extra 1 Gbps.

Shaping Rate (Maximum Bandwidth)

The shaping rate determines the maximum bandwidth the priority group can consume. You specify the rate in bits per second as a fixed value such as 5 Mbps or as a percentage of the total port bandwidth.

The maximum bandwidth for a priority group depends on the total bandwidth available on the port and how much bandwidth the other priority groups on the port consume.

Scheduler Maps

A scheduler map maps schedulers to queues. When you associate a scheduler map with a traffic control profile, then associate the traffic control profile with an interface and a forwarding class set, the scheduling defined by the scheduler map determines the portion of the priority group resources that each individual queue can use.

You can associate up to four user-defined scheduler maps with traffic control profiles.

Related Documentation

- [Understanding Junos CoS Components on page 30](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Understanding CoS Scheduling Behavior and Configuration Considerations on page 104](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 199](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 203](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 300](#)
- [Example: Configuring Maximum Output Bandwidth on page 305](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
- [Example: Configuring WRED Drop Profiles on page 279](#)
- [Example: Configuring Drop Profile Maps on page 281](#)

Understanding CoS Traffic Control Profiles

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (queues) mapped to a forwarding class set share the bandwidth that you assign to the forwarding class set in the traffic control profile.

This two-tier hierarchical scheduling architecture provides flexibility in allocating resources among queues and:

- Assigns a portion of port bandwidth to a priority group. You define the port resources for the priority group in a traffic control profile.
- Allocates priority group bandwidth among the queues that belong to the priority group. A scheduler map attached to the traffic control profile defines the amount of the priority group's resources that each queue can use.

Attaching a priority group and traffic control profile to a port defines the hierarchical scheduling properties of the group and the queues that belong to the group.

The ability to create priority groups supports enhanced transmission selection (ETS, described in IEEE 802.1Qaz). When a priority group does not use its allocated port bandwidth, ETS shares the excess port bandwidth among other priority groups on the port in proportion to their guaranteed minimum bandwidth (guaranteed rate). This utilizes the port bandwidth better than scheduling schemes that require setting strict priorities that reserve bandwidth for all groups whether it is needed or not. ETS allows traffic groups that need extra bandwidth to use it if the bandwidth is available, while preserving the ability to specify the minimum guaranteed bandwidth for traffic groups.

Traffic control profiles define the following CoS properties for priority groups:

- Minimum guaranteed bandwidth—Also known as the committed information rate (CIR). This is the minimum amount of port bandwidth the priority group receives. Priorities in the priority group receive their minimum guaranteed bandwidth as a portion of the priority group's minimum guaranteed bandwidth. The **guaranteed-rate** statement defines the minimum guaranteed bandwidth.



NOTE: You cannot apply a traffic control profile with a minimum guaranteed bandwidth to a priority group that includes strict-high priority queues.

- Shared excess (extra) bandwidth—When the priority groups on a port do not consume the full amount of bandwidth allocated to them or there is unallocated link bandwidth available, priority groups can contend for that extra bandwidth if they need it. Priorities in the priority group contend for extra bandwidth as a portion of the priority group's extra bandwidth. The amount of extra bandwidth for which a priority group can contend is proportional to the priority group's guaranteed minimum bandwidth (guaranteed rate).

- **Maximum bandwidth**—Also known as peak information rate (PIR). This is the maximum amount of port bandwidth the priority group receives. Priorities in the priority group receive their maximum bandwidth as a portion of the priority group's maximum bandwidth. The **shaping-rate** statement defines the maximum bandwidth.
- **Queue scheduling**—Each traffic control profile includes a scheduler map. The scheduler map maps priorities (forwarding classes) to schedulers to define the scheduling characteristics of the individual priorities in the priority group. The resources scheduled for each priority represent portions of the resources that the traffic control profile schedules for the entire priority group, not portions of the total link bandwidth. The **scheduler-maps** statement defines the mapping of forwarding classes to schedulers.

Related Documentation

- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 416](#)

Understanding CoS Priority Group and Queue Guaranteed Rates (Minimum Bandwidth)

You can set a guaranteed minimum bandwidth for individual forwarding classes (queues) and for groups of forwarding classes called forwarding class sets (priority groups). Setting a minimum guaranteed bandwidth ensures that priority groups and queues receive the bandwidth required to support the expected traffic.

This topic covers:

- [Guaranteeing Bandwidth Using Hierarchical Scheduling on page 99](#)
- [Priority Group Guaranteed Rate \(Minimum Bandwidth\) on page 101](#)
- [Queue Transmit Rate \(Minimum Bandwidth\) on page 101](#)

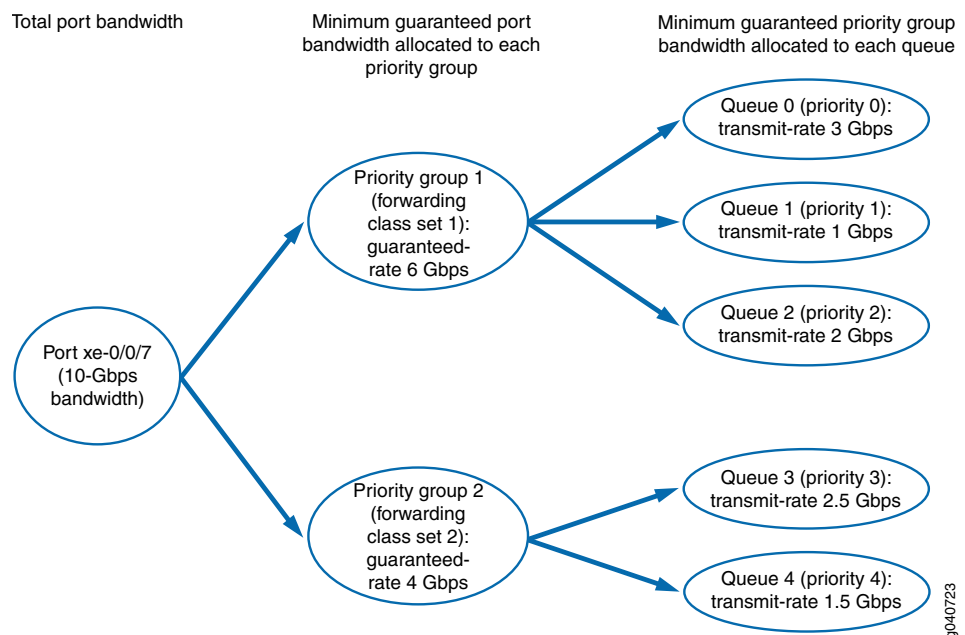
Guaranteeing Bandwidth Using Hierarchical Scheduling

The **guaranteed-rate** value for the priority group defines the minimum amount of bandwidth allocated to a forwarding class set on a port, whereas the **transmit-rate** value of the queue defines the minimum amount of bandwidth allocated to a particular queue in a priority group. The queue bandwidth is a portion of the priority group bandwidth.



NOTE: You cannot configure a minimum guaranteed bandwidth (transmit rate) for a forwarding class that is mapped to a strict-high priority queue, and you cannot configure a minimum guaranteed bandwidth (guaranteed rate) for a priority group that includes strict-high priority queues.

[Figure 7 on page 100](#) shows how the total port bandwidth is allocated to priority groups (forwarding class sets) based on the guaranteed rate of each priority group. It also shows how the guaranteed bandwidth of each priority group is allocated to the queues in the priority group based on the transmit rate of each queue.

Figure 7: Allocating Guaranteed Bandwidth Using Hierarchical Scheduling

The sum of the priority group guaranteed rates cannot exceed the total port bandwidth. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.

The sum of the queue transmit rates cannot exceed the total guaranteed rate of the priority group to which the queues belong. If you configure transmit rates whose sum exceeds the priority group guaranteed rate, the commit check fails and the system rejects the configuration.



NOTE: You must set both the priority group **guaranteed-rate** value and the queue **transmit-rate** value in order to configure the minimum bandwidth for individual queues. If you set the **transmit-rate** value but do not set the **guaranteed-rate** value, the configuration fails.

You can set the **guaranteed-rate** value for a priority group without setting the **transmit-rate** value for individual queues in the priority group. However, queues that do not have a configured **transmit-rate** value can become starved for bandwidth if other higher-priority queues need the priority group's bandwidth. To avoid starving a queue, it is a good practice to configure a **transmit-rate** value for most queues.

If you configure the guaranteed rate of a priority group as a percentage, configure all of the transmit rates associated with that priority group as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.

Priority Group Guaranteed Rate (Minimum Bandwidth)

Setting a priority group **guaranteed-rate** enables you to reserve a portion of the port bandwidth for the forwarding classes (queues) in that forwarding class set. The minimum bandwidth (**guaranteed-rate**) that you configure for a priority group sets the minimum bandwidth available to all of the forwarding classes in the forwarding class set.

The combined **guaranteed-rate** value of all of the forwarding class sets associated with an interface cannot exceed the amount of bandwidth available on that interface.

You configure the priority group **guaranteed-rate** in the traffic control profile. You cannot apply a traffic control profile that has a guaranteed rate to a priority group that includes strict-high priority queues.

Queue Transmit Rate (Minimum Bandwidth)

Setting a queue **transmit-rate** enables you to reserve a portion of the priority group bandwidth for the individual queue. For example, a queue that handles Fibre Channel over Ethernet (FCoE) traffic might require a minimum rate of 4 Gbps to ensure the class of service that storage area network (SAN) traffic requires.

The priority group **guaranteed-rate** sets the aggregate minimum amount of bandwidth available to the queues that belong to the priority group. The cumulative total minimum bandwidth the queues consume cannot exceed the minimum bandwidth allocated to the priority group to which they belong. (The combined transmit rates of the queues in a priority group cannot exceed the priority group's guaranteed rate.)

You must configure the **guaranteed-rate** value of the priority group in order to set a **transmit-rate** value for individual queues that belong to the priority group. The reason is that if there is no guaranteed bandwidth for a priority group, there is no way to guarantee bandwidth for queues in that priority group.

You configure the queue **transmit-rate** in the scheduler configuration. You cannot configure a transmit rate for strict-high priority queues.

**Related
Documentation**

- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Understanding CoS Traffic Control Profiles on page 98](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
- [Defining CoS Queue Schedulers on page 410](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 416](#)

Understanding CoS Priority Group Shaping and Queue Shaping (Maximum Bandwidth)

If the amount of traffic on an interface exceeds the maximum bandwidth of the interface, it leads to congestion. You can use priority group (forwarding class set) shaping and queue shaping to manage the excess traffic and avoid congestion.

The maximum bandwidth sets the most bandwidth a priority group or a queue can use after all of the priority group and queue minimum bandwidth requirements are met, even if more bandwidth is available.

This topic covers:

- [Priority Group Shaping on page 102](#)
- [Queue Shaping on page 102](#)
- [Shaping Maximum Bandwidth Using Hierarchical Scheduling on page 103](#)

Priority Group Shaping

Priority group shaping enables you to shape the aggregate traffic of a forwarding class set on a port to a maximum rate that is less than the line or port rate. The maximum bandwidth (**shaping-rate**) that you configure for a priority group sets the maximum bandwidth available to all of the forwarding classes (queues) in the forwarding class set.

If a port has more than one priority group and the combined **shaping-rate** value of the priority groups is greater than the amount of port bandwidth available, the bandwidth is shared proportionally among the priority groups.

You configure the priority group **shaping-rate** in the traffic control profile.

Queue Shaping

Queue shaping throttles the rate at which queues transmit packets. For example, using queue shaping, you can rate-limit a strict-high priority queue so that the strict-priority queue does not lock out (or starve) low-priority queues.



NOTE: We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

Similarly, for any queue, you can configure queue shaping (**shaping-rate**) to set the maximum bandwidth for a particular queue.

The **shaping-rate** value of the priority group sets the aggregate maximum amount of bandwidth available to the queues that belong to the priority group. The cumulative total bandwidth the queues consume cannot exceed the maximum bandwidth of the priority group to which they belong on a port.


If a priority group has more than queue and the combined **shaping-rate** value of the queues is greater than the amount of bandwidth available to the priority group, the bandwidth is shared proportionally among the queues.

You configure the queue **shaping-rate** in the scheduler configuration.

Shaping Maximum Bandwidth Using Hierarchical Scheduling

Priority group shaping defines the maximum bandwidth allocated to a forwarding class set on a port, whereas queue shaping defines a limit on maximum bandwidth usage per queue. The queue bandwidth is a portion of the priority group bandwidth.

Figure 8 on page 104 shows how the port bandwidth is allocated to priority groups (forwarding class sets) based on the shaping rate of each priority group, and how the bandwidth of each priority group is allocated to the queues in the priority group based on the shaping rate of each queue.

Figure 8: Setting Maximum Bandwidth Using Hierarchical Scheduling**Related Documentation**

- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Understanding CoS Traffic Control Profiles on page 98](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
- [Defining CoS Queue Schedulers on page 410](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 416](#)

Understanding CoS Scheduling Behavior and Configuration Considerations

Many factors affect scheduling configuration and bandwidth requirements, including:

- When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.
- When you define a forwarding class that will be used on the switch (the behavior aggregate classifier has a forwarding class and you expect traffic for the forwarding class), you must also define a scheduling policy for the forwarding class. Defining a scheduling policy means:
 - Mapping a scheduler to the forwarding class in a scheduler map
 - Including the forwarding class in a forwarding class set

- Associating the scheduler map with a traffic control profile
- Attaching the traffic control profile to a forwarding class set and an interface
- On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.
- For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rule rewrites only the outer VLAN tag.
- Configuring the minimum guaranteed bandwidth (**transmit-rate**) for a queue (forwarding class) does not work unless you also configure the minimum guaranteed bandwidth (**guaranteed-rate**) for the priority group (forwarding class set) in the traffic control profile.

Additionally, the sum of the transmit rates of the queues in a forwarding class set should not exceed the guaranteed rate for the forwarding class set. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.) If you configure transmit rates whose sum exceeds the guaranteed rate of the forwarding class set, the commit check fails and the system rejects the configuration.

- The sum of the priority group guaranteed rates cannot exceed the total port bandwidth. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.
- If you configure the **guaranteed-rate** of a priority group as a percentage, configure all of the transmit rates associated with that priority group as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.
- There are several factors to consider if you want to configure strict-high priority queues:
 - You cannot configure a minimum guaranteed bandwidth (**transmit-rate**) for a strict-high priority queue. You cannot configure a minimum guaranteed bandwidth (**guaranteed-rate**) for a forwarding class set that includes a strict-high priority queue.
 - You must create a separate forwarding class set for the strict-high priority queue.
 - Only one forwarding class set can contain strict-high priority queues.
 - Strict-high priority queues cannot belong to the same forwarding class set as queues that are not strict-high priority.
 - A strict-high priority queue cannot belong to a multidestination forwarding class set.
 - We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- In QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets. Failure of a queue to transmit packets for 12 consecutive seconds may be due to:
 - A strict-high priority queue consuming all of the port bandwidth
 - Several queues consuming all of the port bandwidth
 - Any queue or port receiving continuous priority-based flow control (PFC) or 802.3x Ethernet PAUSE messages (received PFC and PAUSE messages prevent a queue or a port, respectively, from transmitting packets because of network congestion)
 - Other conditions that prevent a queue from obtaining port bandwidth for 12 consecutive seconds

If the cause is a strict-high priority queue consuming all of the port bandwidth, use rate shaping to configure a maximum rate for the strict-high priority queue and prevent it from using all of the port bandwidth. To configure rate shaping, include the **shaping-rate (rate | percent percentage)** statement at the **[edit class-of-service schedulers scheduler-name]** hierarchy level and apply the shaping rate to the strict-high priority scheduler. We recommend that you always apply a shaping rate to strict-high priority traffic to prevent the strict-high priority queue from starving other queues.

If several queues consume all of the port bandwidth, you can use a scheduler to rate shape those queues and prevent them from using all of the port bandwidth.

- For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.
- When you set the maximum bandwidth for a queue or for a priority group (**shaping-rate**) at 100 Kbps or lower, the traffic shaping behavior is accurate only within +/– 20 percent of the configured **shaping-rate**.
- Ingress port congestion can occur during periods of egress port congestion if an ingress port forwards traffic to more than one egress port, and at least one of those egress ports experiences congestion. If this occurs, the congested egress port can cause the ingress port to exceed its fair allocation of ingress buffer resources. When the ingress port exceeds its buffer resource allocation, frames are dropped at the ingress. Ingress port frame drop affects not only the congested egress ports, but also all of the egress ports to which the congested ingress port forwards traffic.

If a congested ingress port drops traffic that is destined for one or more uncongested egress ports, configure a weighted random early detection (WRED) drop profile and apply it to the egress queue that is causing the congestion. The drop profile prevents the congested egress queue from affecting egress queues on other ports by dropping frames at the egress instead of causing congestion at the ingress port.



NOTE: Do not configure drop profiles for the **fcoe** and **no-loss** forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

- On an ingress port, do not configure classifiers that map the same IEEE 802.1p code point to both a multidestination traffic flow and a lossless unicast traffic flow (such as the default lossless **fcoe** or **no-loss** forwarding classes). Any code point used for multidestination traffic on a port should not be used to classify unicast traffic into a lossless forwarding class on the same port.

If a multidestination traffic flow and a lossless unicast traffic flow use the same code point on a port, the multidestination traffic is treated the same way as the lossless traffic. For example, if priority-based flow control (PFC) is applied to the lossless traffic, the multidestination traffic of the same code point is also paused. During periods of congestion, treating multidestination traffic the same as lossless unicast traffic can create ingress port congestion for the multidestination traffic and affect the multidestination traffic on all of the egress ports the multidestination traffic uses.

For example, the following configuration can cause ingress port congestion for the multidestination flow:

1. For unicast traffic, IEEE 802.1p code point 011 is classified into the **fcoe** forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1p ucast-cl forwarding-class fcoe
loss-priority low code-points 011
```

2. For multidestination traffic, IEEE 802.1p code point 011 is classified into the **mcast** forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1p mcast-cl forwarding-class mcast
loss-priority low code-points 011
```

3. The unicast classifier that maps traffic with code point 011 to the **fcoe** forwarding class is mapped to interface **xe-0/0/1**:

```
user@switch# set class-of-service interfaces xe-0/0/1 unit 0 classifiers ieee-802.1p ucast-cl
```

4. The multidestination classifier that maps traffic with code point 011 to the **mcast** forwarding class is mapped to all interfaces (multidestination traffic maps to all interfaces and cannot be mapped to individual interfaces):

```
user@switch# set class-of-service multi-destination classifiers ieee-802.1p mcast-cl
```

Because the same code point (011) maps unicast traffic to a lossless traffic flow and also maps multidestination traffic to a multidestination traffic flow, the multidestination traffic flow might experience ingress port congestion during periods of congestion.

To avoid ingress port congestion, do not map the code point used by the multidestination traffic to lossless unicast traffic. For example:

1. Instead of classifying code point **011** into the **fcoe** forwarding class, classify code point **011** into the **best-effort** forwarding class:

```
user@switch# set class-of-service classifiers ieee-802.1 ucast-cl forwarding-class  
best-effort loss-priority low code-points 011
```

2. user@switch# set class-of-service classifiers ieee-802.1 mcast-cl forwarding-class mcast
loss-priority low code-points 011

3. user@switch# set class-of-service interfaces xe-0/0/1 unit 0 classifiers ieee-802.1 ucast-cl

4. user@switch# set class-of-service multi-destination classifiers ieee-802.1 mcast-cl

Because the code point **011** does not map unicast traffic to a lossless traffic flow, the multidestination traffic flow does not experience ingress port congestion during periods of congestion.

The best practice is to classify unicast traffic with IEEE 802.1p code points that are also used for multidestination traffic into best-effort forwarding classes.

Related Documentation

- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Understanding CoS Priority Group Scheduling on page 95](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [*Benefits of Configuring CoS Hierarchical Port Scheduling*](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)

Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows

Junos OS Release 12.3 increased support for lossless priorities from two lossless forwarding classes to up to six lossless forwarding classes on QFX 3500 and QFX3600 switches. Each forwarding class is mapped to an IEEE 802.1p code point (priority).



NOTE: Junos OS Release 13.1 introduced support for up to six lossless forwarding classes on QFabric systems. Throughout this document, features introduced on standalone switches in Junos OS Release 12.3 are introduced on QFabric systems in Junos OS Release 13.1 unless otherwise noted.

Junos OS Release 13.2 is the first QFX5100 switch release. The QFX5100 also supports up to six lossless forwarding classes. However, because the QFX5100 switch has no native Fibre Channel (FC) interfaces, the QFX5100 switch does not support native FC traffic and does not support configuration as an FCoE-FC gateway. Throughout this document, features that pertain to native FC traffic and to FCoE-FC gateway configuration do not apply to QFX5100 switches.

Earlier Junos OS software releases supported two lossless forwarding classes, the default *fc* and *no-loss* forwarding classes, which are mapped by default to IEEE 802.1p priorities 3 (code point 011) and 4 (code point 100), respectively. Junos OS Release 12.3 also introduced a new output stanza in the congestion notification profile (CNP) to configure priority-based flow control (PFC) on output queues.



Video: [Why Use PFC in a Data Center Network?](#)

The default configuration is the same as the default configuration in Junos OS Release 12.2 and is backward-compatible. If you need only two (or fewer) lossless forwarding classes, use the default configuration. If you need more than two lossless forwarding classes, you can use the two default forwarding classes and configure additional lossless forwarding classes. If you do not want to use the default lossless forwarding classes, you can change them or use only the lossless forwarding classes that you explicitly configure.

- [Lossless Transport Features Introduced in Junos OS Release 12.3 on page 110](#)
- [Default Lossless Priority Configuration on page 110](#)
- [Configuring Lossless Priorities on page 113](#)
- [Backward Compatibility with Junos OS Releases Earlier Than Release 12.3 on page 125](#)
- [Configuration Rules and Recommendations on page 126](#)

Lossless Transport Features Introduced in Junos OS Release 12.3

Support for lossless transport introduced in Junos OS Release 12.3 includes:

- Configuring up to six lossless forwarding classes.
- Configuring PFC pause on output queues to program the output queues that can respond to PFC pause messages received from the connected peer. The priorities you pause on output queues must match the priorities on which you enable PFC on the corresponding ingress interfaces. For example, if you program output queues to pause priorities 3 (011) and 5 (101), then you must also enable pause on priorities 3 and 5 on the corresponding ingress interfaces. Configuring flow control on the output queues and enabling PFC on the corresponding input queues allows you to pause up to six priorities (forwarding classes).
- Controlling the headroom buffer on Ethernet interfaces by configuring the maximum receive unit (MRU) size for the traffic mapped to an IEEE 802.1p priority (configured per priority) and the length of the attached cable (configured per interface). The MRU size can range up to full jumbo packet size (9216 bytes).
- Remapping (rewriting) IEEE 802.1p priorities on native Fibre Channel (FC) interfaces when the system is acting as an FCoE-FC gateway. If the Ethernet (FCoE) network uses a different IEEE 802.1p priority than priority 3 (011) for FCoE traffic, then you can use priority remapping to classify FCoE traffic into a lossless forwarding class mapped to that different priority (see [“Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway” on page 128](#)).

Lossless transport still requires configuring previously existing features, including enabling PFC on the lossless priorities on ingress interfaces, and configuring classifiers to classify incoming traffic into lossless forwarding classes based on the IEEE 802.1p priority tag of the packet.



NOTE: If you expect a large amount of lossless traffic on your network and configure multiple lossless traffic classes, ensure that you reserve enough scheduling resources (bandwidth) and lossless headroom buffer space to support the lossless flows. ([“Understanding CoS Buffer Configuration” on page 131](#) describes how to configure buffers and provides a recommended buffer configuration for networks with larger amounts of lossless traffic.)

Default Lossless Priority Configuration

If you do not explicitly configure forwarding classes, the system uses the default forwarding class configuration, which provides two default lossless forwarding classes (*fcoe* and *no-loss*). (If you change the forwarding class configuration on a QFX Series switch or on a QFabric Node device, the changes apply to all traffic on that device because forwarding classes are global to a particular device.)

If you do not explicitly configure classifiers, and you do not explicitly configure flow control to pause output queues (configured in the output stanza of the CNP), the default classifier and the default output queue pause configuration are applied to all Ethernet interfaces

on QFX Series switches and Node devices. You can override the default classifier and the default output queue pause configuration on a per-interface basis by applying an explicit configuration to an Ethernet interface. The default configuration is used on all Ethernet interfaces that do not have an explicit configuration.



NOTE: If you do not configure flow control on output queues, the default configuration uses a one-to-one mapping of IEEE 802.1p code points (priorities) to output queues by number. For example, priority 0 (code point 000) is mapped to queue 0, priority 1 (code point 001) is mapped to queue 1, and so on. If you do not use the default configuration, you must explicitly configure flow control on each output queue that you want to enable for PFC pause in the output stanza of the CNP.

In the default configuration, only queue 3 and queue 4 are enabled to respond to pause messages from the connected peer. For queue 3 to respond to pause messages, priority 3 (code point 011) must be enabled for PFC in the input stanza of the CNP. For queue 4 to respond to pause messages, priority 4 (code point 100) must be enabled for PFC in the input stanza of the CNP.

The default configuration is the same as the default configuration in software releases earlier than Junos OS Release 12.3, and provides the same lossless behavior:

- There are two default lossless forwarding classes (the no-loss packet drop attribute is applied automatically):
fcoe—Mapped to output queue 3
no-loss—Mapped to output queue 4
- The default classifier maps the fcoe forwarding class to IEEE 802.1p priority 3 (011) and the no-loss forwarding class to IEEE 802.1p priority 4 (100)
- Priority-based flow control (PFC) is enabled on Ethernet interface output queues 3 and 4 when those queues carry lossless traffic (traffic that is mapped to the fcoe and no-loss forwarding classes, respectively). In Junos OS software releases earlier than Release 12.3, output queue flow control was not user-configurable.

On native FC interfaces (NP_Ports), default flow control is enabled on output queue 3 (IEEE 802.1p priority 3) for FCoE/FC traffic.

- PFC must be enabled explicitly on the lossless IEEE 802.1p priorities (code points) on ingress Ethernet interfaces; no default PFC configuration is applied at ingress interfaces. If you do not enable PFC on lossless priorities, those priorities might experience packet loss during periods of congestion. For example, if you want lossless FCoE traffic and you are using the default fcoe forwarding class, you use a CNP to enable PFC on priority 3 (code point 011), and apply that CNP to all ingress interfaces that carry FCoE traffic.
- On Ethernet ports, PFC buffer calculations use the following default values to determine the headroom buffer size:
Cable length—100 meters (approximately 328 feet)
MRU for priority 3 traffic—2500 bytes
MRU for priority 4 traffic—9216 bytes

Maximum transmission unit (MTU)—1522 (or the configured MTU value for the interface)



NOTE: If you configure flow control on a priority that is not one of the default flow control priorities, the default MRU value is 2500 bytes. For example, if you configure flow control on priority 5 and you do not configure an MRU value, the default MRU value is 2500 bytes.

- DCBX is enabled on all interfaces in autonegotiation mode, and automatically exchanges FCoE application protocol type, length, and values (TLVs) on interfaces that carry FCoE traffic. However, if you explicitly configure DCBX protocol TLV exchange for any application, then you must explicitly configure protocol TLV exchange for every application for which you want DCBX to exchange TLVs, including FCoE.

The default CoS configuration is backward-compatible with the *default* CoS configuration of software releases before Junos OS Release 12.3. If you explicitly configure lossless transport, ensure that the input and output queues corresponding to the lossless forwarding classes are explicitly configured for PFC pause.



NOTE: If you *explicitly* configured the lossless fcoe or no-loss forwarding classes before upgrading from a release earlier than Junos OS Release 12.3, those forwarding classes are *not* lossless after the upgrade to Junos OS Release 12.3 or later. To regain lossless behavior, you can delete the explicit configuration and use the default lossless forwarding classes, or you can use the no-loss packet drop attribute introduced in Junos OS Release 12.3 to configure the forwarding classes for lossless behavior.

Table 45 on page 112 summarizes the default unicast forwarding classes and their mapping to output queues, IEEE 802.1p priorities, and drop attributes.

Table 45: Mapping of Default Unicast Forwarding Class to Queue, IEEE 802.1p Priority, and Drop Attribute

Forwarding Class Name	Output Queue	Priority	Drop Attribute
best-effort	0	0	drop
fcoe	3	3	no-loss
no-loss	4	4	no-loss
network-control	7	7	drop

There is one default multdestination forwarding class named *mcast* for multicast, broadcast, and destination lookup fail (DLF) traffic that is mapped to output queue 8 with a drop attribute of drop. (Incoming multdestination traffic on all IEEE 802.1p priorities is mapped to the mcast forwarding class by default.)

Configuring Lossless Priorities

Configuring more than two lossless priorities (forwarding classes), or changing the default mapping of lossless forwarding classes to priorities and paused output queues, requires explicit configuration. Configuring lossless priorities includes:

- Configuring forwarding classes with the no-loss packet drop attribute
- Using a CNP to configure PFC on ingress interfaces and flow control (PFC) on egress interfaces
- Configuring a classifier to map IEEE 802.1p priorities (code points) to the correct forwarding classes (the forwarding classes for which you want lossless transport)

In addition, on Ethernet interfaces, DCBX must exchange the appropriate application protocol TLVs for the lossless traffic, and when the switch acts as an FCoE-FC gateway, you need to remap the FCoE priority on native FC interfaces if your network uses a priority other than 3 (IEEE code point 011) for FCoE traffic. This section describes:

- [Configuring Lossless Forwarding Classes \(Packet Drop Attribute\) on page 113](#)
- [Congestion Notification Profiles \(PFC Configuration\) on page 115](#)
- [Configuring DCBX \(Application Protocol TLV Exchange\) on page 121](#)
- [Fate Sharing Among Traffic Classes on page 121](#)
- [Transit Switch Configuration Versus FCoE-FC Gateway Configuration on page 123](#)
- [Configuration Results and Commit Checks on page 123](#)

Configuring Lossless Forwarding Classes (Packet Drop Attribute)

Junos OS Release 12.3 introduced the *no-loss* parameter for forwarding class configuration. (Although it uses the same name, this is not the no-loss default forwarding class. It is a packet drop attribute you can specify to configure any unicast forwarding class as a lossless forwarding class.)

You can configure up to six forwarding classes (depending on system architecture and the availability of system resources) as lossless forwarding classes by including the **no-loss** drop attribute at the **[edit class-of-service forwarding-classes class forwarding-class-name queue-num queue-number]** hierarchy level.

If you use the default fcoe or no-loss forwarding classes, they include the no-loss drop attribute by default. If you explicitly configure the fcoe or no-loss forwarding classes and you want to retain their lossless behavior, you *must* include the no-loss drop attribute in the configuration.



NOTE: All forwarding classes mapped to the same output queue must have the same packet drop attribute. (All forwarding classes mapped to the same output queue must be either lossy or lossless. You cannot map both a lossy and a lossless forwarding class to the same queue.)

To avoid fate sharing (different flows receiving the same CoS treatment), use a one-to-one mapping of lossless forwarding classes to IEEE 802.1p code points (priorities) and queues. (Each forwarding class should be mapped to a different queue and classified into a different priority.) The classifier attached to the interface determines the forwarding class to priority mapping.

The fcoe and no-loss forwarding classes are special cases, because in the default configuration, they are configured for lossless behavior (providing that you also enable PFC on the priorities mapped to the fcoe and no-loss forwarding classes in the CNP input stanza).

Table 46 on page 114 summarizes the possible configurations of the fcoe and no-loss forwarding classes in Junos OS Release 12.3 and later, and the result of those configurations in terms of lossless traffic behavior. It is assumed that PFC, DCBX, and classifiers are properly configured.

Table 46: FCoE and No-Loss Forwarding Class Configuration in Junos OS Release 12.3

Explicit (User-Configured) or Default Forwarding Class Configuration	Packet Drop Attribute	Result and Notes
Default	Default	The fcoe and no-loss forwarding classes are lossless. NOTE: Even if you explicitly configure other forwarding classes (lossy or lossless forwarding classes), the fcoe and no-loss forwarding classes remain lossless because they are not explicitly configured.
Explicit	Not specified in the explicit forwarding class configuration	The fcoe and no-loss forwarding classes are lossy because they do not include the no-loss drop attribute.
Explicit	No-loss	The fcoe and no-loss forwarding classes are lossless.
Explicit, configured in Junos OS Release 12.2 or earlier	Not specified (packet drop attribute was not available before Junos OS Release 12.3)	The fcoe and no-loss forwarding classes are lossy in Junos OS Release 12.3 and later because they do not include the no-loss drop attribute. NOTE: To retain lossless behavior, before you upgrade to Junos OS Release 12.3, delete the explicit configuration so that the system uses the default configuration. Alternatively, you can reconfigure the forwarding classes with the no-loss packet drop attribute after upgrading to Junos OS Release 12.3 or later.

For all other forwarding classes, you must explicitly configure lossless transport by specifying the no-loss packet drop attribute, because the default configuration for all other forwarding classes is lossy.

Congestion Notification Profiles (PFC Configuration)

Use CNPs to configure lossless PFC characteristics on input and output interfaces.

The input stanza of a CNP enables PFC on specified IEEE 802.1p priorities (code points) and fine-tunes headroom buffer settings by configuring the maximum receive unit (MRU) value and cable length on ingress interfaces.

The output stanza of a CNP enables PFC (flow control) on output queues for specified IEEE 802.1p priorities so that the queues can respond to PFC pause messages from the connected peer on the priority of your choice. (By default, output queues 3 and 4 respond to received PFC messages when those queues carry lossless traffic in the fcoe and no-loss forwarding classes, respectively.)

To achieve lossless transport, the priority paused at the ingress interfaces must match the priority paused at the egress interfaces for a given traffic flow. For example, if you configure ingress interfaces to pause traffic tagged with IEEE 802.1p priority 5 (code point 101) and priority 5 traffic is mapped to output queue 5, then you must also configure the corresponding output interfaces to pause priority 5 on queue 5. In addition, the forwarding class mapped to queue 5 must be configured as a lossless forwarding class (using the no-loss drop attribute).



CAUTION: Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

A change to the PFC configuration means any change to a CNP, including changing the input portion of the CNP (enabling or disabling PFC on a priority, or changing the MRU or cable-length values) or changing the output portion the CNP that enables or disables output flow control on a queue. A PFC configuration change only affects ports that use the changed CNP.

The following actions change the PFC configuration:

- Deleting or disabling a PFC configuration (input or output) in a CNP that is in use on one or more interfaces. For example:
 1. An existing CNP with an input stanza that enables PFC on priorities 3, 5, and 6 is configured on interfaces xe-0/0/20 and xe-0/0/21.
 2. We disable the PFC configuration for priority 6 in the input CNP, and then commit the configuration.

3. The PFC configuration change causes all traffic on interfaces xe-0/0/20 and xe-0/0/21 to stop until the PFC change has been implemented. When the PFC change has been implemented, traffic resumes.

- Configuring a CNP on an interface. (This changes the PFC state by enabling PFC on one or more priorities.)
- Deleting a CNP from an interface. (This changes the PFC state by disabling PFC on one or more priorities.)

Configuring Input Interface Flow Control (PFC and Headroom Buffer Calculation)

On Ethernet interfaces, the input stanza of the CNP enables PFC on specified priorities so that the ingress interface can send a pause message to the connected peer during periods of congestion. Input CNPs also fine-tune the headroom buffers used for PFC support by allowing you to configure the MRU value and cable length (if you do not want to use the default configuration).

Headroom buffers support lossless transport by storing the traffic that arrives at an interface after the interface sends a PFC flow control message to pause incoming traffic. Until the connected peer receives the flow control message and pauses traffic, the interface continues to receive traffic and must buffer it (and the traffic that is still on the wire after the peer pauses) to prevent packet loss.

The system uses the MRU and the length of the attached physical cable to calculate buffer headroom allocation. The default configuration values are:

- MRU for priority 3 traffic—2500 bytes
- MRU for priority 4 traffic—9216 bytes
- Cable length—100 meters (approximately 328 feet)



NOTE: If you configure flow control on a priority that is not one of the default flow control priorities, the default MRU value is 2500 bytes. For example, if you configure flow control on priority 5 and you do not explicitly configure an MRU value, the default MRU value is 2500 bytes.

You can fine-tune the MRU and the cable length to adjust the size of the headroom buffer on an interface. The QFX Series has a shared global buffer pool and dynamically allocates headroom buffer space to lossless queues as needed.

A lower MRU or a shorter cable length reduces the amount of headroom buffer required on an interface and leaves more headroom buffer space for other interfaces. A higher MRU or a longer cable length increases the amount of headroom buffer space required on an interface and leaves less headroom buffer space for other interfaces.

In many cases, you can better utilize the headroom buffers by reducing the MRU value (for example, an MRU of 2180 is sufficient for most FCoE networks) and by reducing the cable length value if the physical cable is less than 100 meters long.



NOTE: When you configure the headroom buffers by changing the MRU or the cable length, and commit the configuration, the system performs a commit check and rejects the configuration if sufficient headroom buffer space is not available.

However, the system does not perform a commit check but instead returns a syslog error if:

- The buffers are configured on a LAG interface.
- The default classifier is used on the interface (instead of a user-configured classifier).
- The interface has not been created yet.

Configuring Output Interface Flow Control (PFC)

On Ethernet interfaces, you can use the output stanza of the CNP to configure flow control on unicast output queues and enable PFC pause response on specified IEEE 802.1p priorities. By default, output queues 3 and 4 are enabled for PFC pause on priorities 3 (IEEE 802.1p code point 011) and 4 (IEEE 802.1p code point 100). The default PFC pause response supports the default lossless forwarding class configuration, which maps the fcoe forwarding class to queue 3 and priority 3, and maps the no-loss forwarding class to queue 4 and priority 4.

Configuring PFC on output queues enables you to pause any priority on any unicast output queue on any Ethernet interface. Output flow control enables you to use more than two output queues to support lossless traffic flows (you can configure up to six lossless forwarding classes and map them to different output queues that are enabled for PFC pause). Output queue flow control also enables you to support multiple lossless forwarding classes (each mapped to a different priority and output queue) for one class of traffic.



NOTE: Output flow control only works when PFC is enabled in the CNP input stanza on the corresponding priorities on the interface.

For example, if the converged Ethernet network uses two different priorities for FCoE traffic (for example, priority 3 and priority 5), then you can classify those priorities into different lossless forwarding classes that are mapped to different output queues by:

1. Configuring two lossless forwarding classes for FCoE traffic, with each forwarding class mapped to a different output queue. For example, you could use the default fcoe forwarding class, which is mapped to queue 3, and you could configure a second lossless forwarding class called fcoe1 and map it to queue 5. The fcoe forwarding class is for priority 3 FCoE traffic (code point 011), and the fcoe1 forwarding class is for priority 5 (code point 101) FCoE traffic.
2. Configuring a classifier that maps each forwarding class to the desired IEEE 802.1p code point (priority). If FCoE traffic on both priorities uses one interface, the classifier

must classify both forwarding classes to the correct priorities. If FCoE traffic of different priorities uses different interfaces, the classifier configuration on each interface must map the correct priority to the corresponding lossless forwarding class.

3. Applying the classifier to the interfaces that carry FCoE traffic. The classifier determines the mapping of forwarding classes to priorities on each interface.

To configure lossless transport for these forwarding classes, you also need to:

- Enable PFC on the two priorities (3 and 5 in this example) at the ingress interfaces in the CNP input stanza.
- Configure PFC on the output queues and priorities for the forwarding classes in the CNP output stanza so that the interface can respond to pause messages received from the connected peer.



NOTE: When you configure the CNP on an interface, all ingress and egress traffic is blocked until the configuration is implemented, then the interface is unblocked and traffic resumes. During the time the interface is blocked, all queues on the interface experience packet loss.

- Configure DCBX to exchange application protocol TLVs on both FCoE priorities.



NOTE: If you do not configure flow control to pause output queues, the default configuration uses a one-to-one mapping of IEEE 802.1p code points (priorities) to output queues by number. For example, priority 0 (code point 000) is mapped to queue 0, priority 1 (code point 001) is mapped to queue 1, and so on. By default, only queues 3 and 4 are enabled to respond to pause messages from the connected peer, and you must explicitly enable PFC on the corresponding priorities in the CNP input stanza to achieve lossless behavior.

If you do not use the default configuration, you must explicitly configure flow control on each output queue that you want to enable for PFC pause. For example, if you explicitly configure flow control on output queue 5, the default configuration is no longer valid, and only output queue 5 is enabled for PFC pause. Output queues 3 and 4 are no longer enabled for PFC pause, so traffic using those queues no longer responds to PFC pause messages even if the corresponding forwarding class is configured with the no-loss drop attribute. To retain the pause configuration on output queues 3 and 4 and configure flow control on queue 5, you need to explicitly configure flow control on queues 3, 4, and 5.

You cannot configure flow control to pause a multidestination output queue. You can configure flow control to pause only unicast output queues.

Output Interface Flow Control Profiles

Configuring the CNP output stanza creates an output flow control profile that tells egress ports the queues on which the Ethernet interface should respond to PFC pause messages. Although you can create an unlimited number of CNPs that contain input stanzas only, the number of CNPs that you can configure with output stanzas is limited:

- For QFX3500 and QFX3600 standalone switches that are not part of a QFabric system, you can configure up to two output interface flow control profiles. (You can configure up to two CNPs with output stanzas.)
- For QFabric systems, you can configure one output interface flow control profile per Node device. (You can configure one CNP with an output stanza per Node device.)

There are a total of four output flow control profiles.

The system has a default output flow control profile that is applied to all Ethernet interfaces when the CNP attached to the interface has only an input stanza and does not include an output stanza. The default profile responds to PFC pause messages received on queue 3 (for priority 3, for the default fcoe forwarding class) and on queue 4 (for priority 4, for the default no-loss forwarding class), and is effective only if PFC is configured on those priorities in the CNP input stanza.

Additionally, the system has two internal output flow control profiles that it applies automatically to fabric (FTE) ports and to native FC interfaces (NP_Ports). When the QFX3500 switch or the QFX3600 switch is not part of a QFabric system, the profile normally used for FTE ports is available for user configuration and provides a second user-configurable profile. (That is why standalone QFX3500 and QFX3600 switches have two user-configurable output flow control profiles, but QFX3500 and QFX3600 switches that are part of a QFabric have only one user-configurable output flow control profile.)

Because one output CNP can configure PFC pause response on multiple output queues (priorities), one user-configurable output CNP is usually flexible enough to specify the desired PFC response on all programmed interfaces.



NOTE: Each port can use one output flow control profile. You cannot apply more than one profile to one port.

Output flow control profiles can be expressed in table format. For example, [Table 47 on page 119](#) shows the default output flow control profile that pauses priorities 3 and 4 on queues 3 and 4 (remember that PFC must also be enabled on code points 3 and 4 in the CNP input stanza in order for PFC to work):

Table 47: Default Output Flow Control Profile

IEEE 802.1p Priority Specified in Received PFC Frame	Paused Output Queue
0 (000)	—

Table 47: Default Output Flow Control Profile (*continued*)

IEEE 802.1p Priority Specified in Received PFC Frame	Paused Output Queue
1 (001)	—
2 (010)	—
3 (011)	3
4 (100)	4
5 (101)	—
6 (110)	—
7 (111)	—

[Table 48 on page 120](#) is an example of a user-configured output flow control profile. Using the example from the preceding section, the CNP output stanza configures flow control on output queue 5, and also explicitly configures output flow control on queues 3 and 4 for the fcoe and no-loss forwarding classes. (If you explicitly configure an output CNP, you must explicitly configure every output queue that you want to respond to PFC messages, because the user-configured profile overrides the default profile. If this example did not include queues 3 and 4, those queues would no longer respond to received PFC messages.)

Table 48: User-Configured Output Flow Control Profile

IEEE 802.1p Priority Specified in Received PFC Frame	Paused Output Queue
0 (000)	—
1 (001)	—
2 (010)	—
3 (011)	3
4 (100)	4
5 (101)	5
6 (110)	—
7 (111)	—

Remember that you must also enable PFC on code points 3, 4, and 5 in the CNP input stanza for this configuration to work. When you configure the CNP on an interface, all ingress and egress traffic is blocked until the configuration is implemented, then the

interface is unblocked and traffic resumes. During the time the interface is blocked, all queues on the interface experience packet loss.

Configuring PFC Across Layer 3 Interfaces on QFX5100 Switches

Enabling PFC on traffic flows is based on the IEEE 802.1p code point (priority) in the priority code point (PCP) field of the Ethernet frame header (sometimes known as the CoS bits). To enable PFC on traffic that crosses Layer 3 interfaces, the traffic must be classified by its IEEE 802.1p code point, not by its DSCP (or DSCP IPv6) code point.

See *Understanding PFC Functionality Across Layer 3 Interfaces* for a conceptual overview of how to enable PFC on traffic across Layer 3 interfaces. See *Example: Configuring PFC Across Layer 3 Interfaces* for an example of how to configure PFC on traffic that traverses Layer 3 interfaces.

Configuring DCBX (Application Protocol TLV Exchange)

For applications that require lossless transport, DCBX exchanges application protocol TLVs with the connected peer interface. By default, DCBX advertises FCoE application protocol TLVs on all interfaces that are enabled for DCBX, and by default, DCBX is enabled on all interfaces. DCBX advertises no other applications by default.

For each application (for example, iSCSI) that you want to configure for lossless transport, you must enable the interfaces which carry that application traffic to exchange DCBX protocol TLVs with the connected peer. The TLV exchange allows the peer interfaces to negotiate a compatible configuration to support the application.

If you configure DCBX to advertise any application, the default DCBX advertisement is overridden, and DCBX advertises only the configured applications. If you want an interface to advertise only the FCoE application, you do not have to configure DCBX application protocol TLV exchange; instead, you can use the default configuration.

If you want DCBX to advertise other applications, you must explicitly configure an application map and apply it to the interfaces on which you want to exchange protocol TLVs for those applications. If you want to exchange FCoE application protocol TLVs in addition to other application protocol TLVs, you must also explicitly configure the FCoE application in the application map. “[Understanding DCBX Application Protocol TLV Exchange](#)” on page 179 describes how application mapping works.



NOTE: Lossless transport also requires that you enable PFC on the correct priority (IEEE 802.1p code point) on the ingress interfaces using an input CNP. If the priority you pause at the ingress interfaces is not mapped to queue 3 or queue 4 (the two output queues that are enabled for PFC pause flow control by default), then you must also enable the output queues that correspond to paused input priorities to pause using the output stanza of the CNP.

Fate Sharing Among Traffic Classes

You can configure different lossless (or lossy) traffic flows to share fate—that is, to receive the same CoS treatment.

Fate sharing is not desirable for I/O convergence. Instead of independent control of the fate of each type of flow, different types of flows receive the same treatment. Fate sharing is particularly undesirable for lossless flows. If one lossless flow experiences congestion and must be paused, that affects flows that share fate with the congested flow even if the other flows are not experiencing congestion, and also can cause ingress port congestion. If your network requires that all 802.1p priorities be lossless, you can achieve that by allowing some fate sharing among the eight priorities by spreading them across up to six lossless forwarding classes.

If the number of lossless priorities is less than or equal to the number of configured lossless forwarding classes, then you can avoid fate sharing by configuring a one-to-one mapping of forwarding classes to IEEE 802.1p code points (priorities) and output queues. (Each forwarding class should be mapped to a different output queue and classified to a different priority.)

If you want to configure different traffic flows to share fate, the QFX Series supports two fate-sharing configurations: mapping one forwarding class to more than one IEEE 802.1p code point (priority), and mapping two forwarding classes to the same output queue:

1. If you map one lossless forwarding class to more than one priority, the traffic tagged with each of the priorities uses the same CoS properties associated (the CoS properties associated with the forwarding class). For example, configuring a forwarding class called `fc1`, mapping it to queue 1, and mapping it to code points 101 and 110 using a classifier named `classify1` results in the traffic tagged with priorities 101 and 110 sharing fate:

```
user@switch# set class-of-service forwarding-classes class fc1 queue-num 1 no-loss
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc1
loss-priority low code-points 101
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc1
loss-priority low code-points 110
```

In this case, if the traffic mapped to either priority experiences congestion, both priorities are paused because they are mapped to the same forwarding class and are therefore treated similarly.

2. If you map multiple lossless forwarding classes to the same output queue, the traffic mapped to the forwarding classes uses the same output queue. This increases the amount of traffic the queue needs to buffer and forward, and can create congestion that affects all of the traffic flows that are mapped to the queue. For example, configuring two forwarding classes called `fc1` and `fc2`, mapping both forwarding classes to queue 1, and mapping the forwarding classes to code points 101 and 110 (respectively) using a classifier named `classify1` results in the traffic tagged with priorities 101 and 110 sharing fate on the same output queue:

```
user@switch# set class-of-service forwarding-classes class fc1 queue-num 1 no-loss
user@switch# set class-of-service forwarding-classes class fc2 queue-num 1 no-loss
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc1
loss-priority low code-points 101
user@switch# set class-of-service classifiers ieee-802.1 classify1 forwarding class fc2
loss-priority low code-points 110
```

In this case, even though the two forwarding classes use different IEEE 802.1p priorities, if one forwarding class experiences congestion, it affects the other forwarding class.

The reason is that if the output queue is paused because of congestion on either forwarding class, all traffic that uses that queue is paused. Since both forwarding classes are mapped to the queue, the traffic mapped to both forwarding classes is paused.



NOTE: If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).

Transit Switch Configuration Versus FCoE-FC Gateway Configuration

On a transit switch (all Ethernet ports, no native FC ports) that forwards FCoE traffic (or other traffic that requires lossless transport across the Ethernet network), the configuration of classifiers, lossless forwarding classes, DCBX, and PFC on ingress and egress interfaces to support lossless transport is as described in this document.

When the QFX Series acts as an FCoE-FC gateway, the system uses native FC interfaces (NP_Ports) to connect to the FC switch (or FCoE forwarder) at the FC network edge. You cannot apply CNPs or DCBX to native FC interfaces, only to Ethernet interfaces.

On an FCoE-FC gateway, the Ethernet interface configuration of classifiers, DCBX, and PFC is the same as the Ethernet interface configuration on a transit switch. The configuration of lossless forwarding classes is also the same.

However, supporting lossless transport on native FC interfaces requires that you rewrite the IEEE 802.1p priority value *if* your network uses any priority other than 3 (IEEE code point 011) for FCoE traffic. If your network uses priority 3 for FCoE traffic, you can and should use the default configuration on native FC interfaces.

By default, native FC interfaces tag packets with priority 3 when they encapsulate the incoming FC packets in Ethernet. If your FCoE network uses a different priority than 3 for FCoE traffic, you need to rewrite the priority value to the value that your network uses on the FC interface, classify the FCoE traffic to the correct priority on the Ethernet interfaces, and enable PFC on the correct priority on the Ethernet interfaces, as described in [“Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway” on page 128](#).

Configuration Results and Commit Checks

Different configurations of forwarding classes and their drop attributes, classifiers, CNPs (PFC flow control), and Ethernet PAUSE (IEEE 802.3X flow control) result in different system behaviors.

[Table 49 on page 124](#) describes the results of the possible lossless transport configurations in each case. The assumption in the *Result* column is that the system’s buffer headroom calculation resulted in a successful configuration.

However, if the system calculates that there is insufficient buffer space to support the configuration, a commit check prevents you from committing the configuration on an

individual Ethernet interface. For LAG interfaces, the system does not issue a commit check error but instead issues a syslog message.



NOTE: After you configure lossless transport for a LAG interface, be sure to check the syslog messages to confirm that the commit was successful.

Table 49: Results of Lossless Priority Configuration

Classifier Configuration	Congestion Notification Profile Configuration	Ethernet PAUSE (IEEE 802.3X) Configuration	Result
None (default classifier)	None	None	System default configuration. No flows are lossless. To achieve lossless behavior for the default fcoe and no-loss forwarding classes, you must configure a CNP to enable PFC on their IEEE 802.1p code points (011 and 100 respectively).
Classifier with no lossless forwarding classes	None	None	No lossless traffic flows are configured; all traffic is best effort.
Classifier with at least one lossless forwarding class	None	None	Because no CNP is attached to interfaces, PFC is not enabled on the code point of the lossless traffic and no headroom buffer is allocated to the lossless queue, so packets can drop during periods of congestion. This configuration does not achieve lossless behavior.
None (default classifier)	PFC enabled on the fcoe and no-loss forwarding class code points (priorities)	None	The default classifier classifies traffic into two lossless forwarding classes, fcoe and no-loss. The CNP enables PFC on the priorities mapped to both lossless forwarding classes, resulting in lossless behavior for traffic mapped to the fcoe and no-loss forwarding classes.
None (default classifier)	None	Flow control enabled	The system calculates buffer headroom for the physical link based on the interface MTU and the default cable length. The system does not calculate buffer headroom for individual output queues. Because Ethernet PAUSE is enabled on the link instead of PFC being enabled on the lossless priorities, the entire link is paused during periods of congestion. This configuration results in lossless behavior for all of the forwarding classes on the link, but because all traffic is paused, this can cause greater overall network congestion.

Table 49: Results of Lossless Priority Configuration (*continued*)

Classifier Configuration	Congestion Notification Profile Configuration	Ethernet PAUSE (IEEE 802.3X) Configuration	Result
Classifier with at least one lossless forwarding class	PFC enabled on the lossless forwarding class code points (priorities)	None	Headroom buffer allocated only to priorities that are mapped to the lossless forwarding classes and on which PFC is enabled. This configuration achieves lossless behavior for the lossless forwarding classes.
Classifier with no lossless forwarding classes	None	Flow control enabled	The system calculates buffer headroom for the physical link based on the interface MTU and the default cable length, and it pauses all traffic on the link during periods of congestion.
Classifier with at least one lossless forwarding class	None	Flow control enabled	The system calculates buffer headroom for the physical link based on the interface MTU and the default cable length, and it pauses all traffic on the link during periods of congestion.
Classifier with at least one lossless forwarding class	PFC enabled on the lossless forwarding class code points (priorities)	Flow control enabled on a <i>different</i> interface than the interface with the CNP	The system checks the available buffer space for both the PFC-enabled priorities and for the other link. If sufficient buffer space is available, the lossless forwarding classes configured with PFC on one interface and also all of the traffic on the link with Ethernet PAUSE enabled achieve lossless behavior.



NOTE: If you attempt to configure both PFC and Ethernet PAUSE on a link, the system returns a commit error. PFC and Ethernet PAUSE are mutually exclusive configurations on an interface.

Backward Compatibility with Junos OS Releases Earlier Than Release 12.3

The addition of the no-loss packet drop attribute to forwarding class configuration means that when you upgrade from an earlier release to Junos OS Release 12.3, the new software might not preserve the lossless forwarding class configuration of the fcoe and no-loss forwarding classes.

If you used the default forwarding class configuration for the fcoe and no-loss forwarding classes, the CoS configuration is backward-compatible. You do not have to do anything to preserve the lossless behavior of traffic that uses those forwarding classes when you upgrade to Junos OS Release 12.3. (This is because the default configuration of these two forwarding classes includes the no-loss packet drop attribute.)

However, if you explicitly configured the fcoe or the no-loss forwarding class by including the **set forwarding-classes class *forwarding-class-name* queue-num *queue-number*** statement at the **[edit class-of-service]** hierarchy level, then those forwarding classes are no longer lossless, they are lossy. (They are lossy because explicit configuration in releases earlier than Junos OS Release 12.3 did not use the no-loss packet drop attribute.) In Junos OS Release 12.3 and later, you must include the no-loss packet drop attribute in explicit forwarding class configurations to configure a lossless forwarding class.

For example, before Junos OS Release 12.3, the following explicit configuration resulted in a lossless forwarding class:

```
user@switch# set class-of-service forwarding-classes class fcoe queue-num 3
```

However, in Junos OS Release 12.3, this configuration is lossy because it does not include the no-loss packet drop attribute. To preserve lossless behavior, after upgrading to Junos OS Release 12.3, you need to add the no-loss drop attribute:

```
user@switch# set class-of-service forwarding-classes class fcoe queue-num 3 no-loss
```

Alternatively, you can delete the explicit configuration before you upgrade to Junos OS Release 12.3 so that the system uses the default forwarding class, which is lossless:

```
user@switch# delete class-of-service forwarding-classes class fcoe queue-num 3
```



NOTE: The explicit configuration of other forwarding classes does not affect the lossless (or lossy) state of the fcoe and no-loss forwarding classes, because only the fcoe and no-loss forwarding classes were lossless forwarding classes before Junos OS Release 12.3. For example, if you explicitly configured the best-effort forwarding class but you used the default fcoe and no-loss forwarding classes in Junos OS Release 12.2, then when you upgrade to Junos OS Release 12.3, the fcoe and no-loss forwarding classes are still lossless (and the best-effort forwarding classes retains its explicit configuration).



NOTE: To achieve lossless behavior for the traffic belonging to any forwarding class, you must also use a CNP to enable PFC on the IEEE 802.1p priority mapped to the forwarding class and apply the CNP to the relevant interfaces, and ensure that DCBX exchanges the protocol TLVs for the application with the connected peer.

Configuration Rules and Recommendations

Keep in mind the following configuration rules and recommendations when you configure lossless traffic flows:

- You can configure a maximum of six lossless forwarding classes (forwarding classes with the no-loss packet drop attribute).
- All forwarding classes that you map to the same queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes must be lossless).
- You cannot configure flow control to pause a multideestination output queue. You can configure PFC flow control only to pause unicast output queues.
- Forwarding classes mapped to multideestination queues (queues 8 through 11) cannot have the no-loss packet drop attribute. (Multideestination forwarding classes cannot be configured as lossless forwarding classes.)
- Do not configure weighted random early detection (WRED) on lossless forwarding classes. (Do not associate a drop profile with a forwarding class that has the no-loss packet drop attribute.)

Related Documentation

- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 128](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 179](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)
- [Understanding CoS Buffer Configuration on page 131](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 199](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 203](#)
- [*Understanding PFC Functionality Across Layer 3 Interfaces*](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 309](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 317](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 340](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 356](#)
- [*Example: Configuring PFC Across Layer 3 Interfaces*](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)

Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

When the QFX Series acts as an FCoE-FC gateway, it connects an Ethernet network that carries Fibre Channel over Ethernet (FCoE) traffic to a Fibre Channel (FC) network. Ethernet interfaces connect to the FCoE network. Native FC interfaces (NP_Ports) connect to the FC network.

FCoE traffic typically uses IEEE 802.1p priority 3 (code point 011). The QFX Series default configuration maps priority 3 traffic to the FCoE forwarding class. If your FCoE network uses priority 3 for FCoE traffic, you do not need to remap priorities, because the default configuration maps priority 3 to the FCoE forwarding class. (But you do need to enable PFC on IEEE 802.1p code point 3 on the Ethernet interfaces to achieve lossless behavior.)

However, if the FCoE network uses a different IEEE 802.1p priority than priority 3 for FCoE traffic, then you can use priority remapping to classify FCoE traffic into a lossless forwarding class mapped to that priority (and classified to that priority on the FCoE Ethernet interfaces in the ingress classifier). You specify the lossless forwarding class used for the FCoE traffic by configuring a fixed classifier and applying it to the native FC (NP_Port) interface. All traffic received from the FC SAN on that NP_Port interface is classified into the forwarding class specified in the fixed classifier.

When native FC interfaces on the FCoE-FC gateway encapsulate incoming FC traffic in Ethernet to create FCoE frames, by default they assign IEEE 802.1p code point 011 to the FCoE traffic, forward the traffic internally to the gateway Ethernet interfaces, and then forward the traffic to the FCoE network. Setting a rewrite value for the IEEE 802.1p code point configures the gateway native FC interface to assign the rewrite value priority to the FCoE frames when the native FC interface forwards the FCoE frames to the gateway Ethernet interface. Instead of a priority of 3, the FCoE frames use the priority specified in the rewrite value.

You can configure one rewrite value for each local FCoE-FC gateway fabric. All of the native FC interfaces in a particular fabric must use the same rewrite value. Native FC interfaces that belong to different FCoE-FC gateway fabrics can use different rewrite values.

- [Priority Remapping Configuration on page 128](#)
- [Configuration Rules on page 129](#)
- [Fate Sharing on page 130](#)

Priority Remapping Configuration

Native FC interfaces on an FCoE-FC gateway receive native FC traffic from the FC SAN and encapsulate it in Ethernet to create FCoE frames. Priority remapping enables you to map the encapsulated FC traffic (the FCoE traffic) to any IEEE 802.1p priority. (This is similar to the rewrite rules you can configure to remap forwarding classes to code points on Ethernet egress interfaces, but the rewrite takes place at the ingress FC interface so that the QFX Series uses the correct priority for FCoE traffic on the converged Ethernet network.)

To support lossless traffic flows, you must configure the remapped priority correctly on the native FC interfaces and also on the Ethernet interfaces that connect to the FCoE network. Achieving lossless behavior for FCoE traffic when you remap the FCoE priority requires configuring:

- A lossless forwarding class for FCoE traffic (or using the default *fcoe* forwarding class)
- A behavior aggregate (BA) classifier on the FCoE Ethernet interfaces to map the FCoE forwarding class to the IEEE 802.1p code points (priority) used for FCoE traffic on the FCoE network (the ingress classifier priority for the forwarding class must be the same as the rewrite value priority)
- A fixed classifier on the FCoE-FC gateway FC interface that maps all traffic from the FC network into the lossless FCoE forwarding class (the forwarding class must be lossless)
- A priority rewrite value that remaps the IEEE 802.1p code point on the FCoE-FC gateway FC interface to the priority used for FCoE traffic on the FCoE network
- An input congestion notification profile (CNP) to enable priority-based flow control (PFC) on the FCoE code point (the code point used as the rewrite value) at the Ethernet ingress interfaces

The ingress and egress configurations must match to achieve lossless behavior. The priority and the forwarding class specified in the BA classifier and in the CNP on the Ethernet ingress interfaces must match the fixed classifier and rewrite value on the FC interfaces. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

For example, if you configure a lossless forwarding class named *my_fcoe_fc* and your Ethernet network uses IEEE 802.1p priority 5 (code point 101) for FCoE traffic, then:

- The forwarding class configuration, the BA classifier, and the fixed classifier all specify *my_fcoe_fc* as the forwarding class
- The BA classifier, the input CNP, and the rewrite value all specify the IEEE 802.1p code point 101

Configuration Rules

The following configuration rules apply when you remap priorities on an FCoE-FC gateway:

- Each native FC interface (NP_Port) supports one IEEE 802.1p priority value. The interface rewrites the IEEE 802.1p code point of all incoming traffic on the interface to the rewrite value. (The FC interface uses either the default value of 3 or the rewrite value for all incoming traffic.)
- Ports in the same FCoE-FC gateway local fc-fabric must use the same rewrite value. For example, if ports fc-0/0/0 and fc-0/0/1 are in the same local FCoE-FC gateway fabric, they must use the same rewrite value. If you attempt to commit a configuration that uses different IEEE 802.1p priority rewrite values, the system returns a commit error.

- Ports in different FCoE-FC gateway local fc-fabrics can use different rewrite values. An example scenario is:

- Interfaces fc-0/0/0 and fc-0/0/1 are in FCoE-FC gateway fc-fabric *my_fc_fab1*.
- Interfaces fc-0/0/4 and fc-0/0/5 are in FCoE-FC gateway fc-fabric *my_fc_fab2*.

In this scenario, interfaces fc-0/0/0 and fc-0/0/1 must use the same rewrite value because they belong to the same local FC fabric on the gateway. Interfaces fc-0/0/4 and fc-0/0/5 also must use the same rewrite value because they belong to the same local FC fabric. However, the rewrite value you use for interfaces fc-0/0/0 and fc-0/0/1 can be different than the rewrite value you use for interfaces fc-0/0/4 and fc-0/0/5 because the interfaces belong to different local FC fabrics.

- You can apply the rewrite value only to native FC interfaces; you cannot apply the rewrite value configuration to Ethernet interfaces.
- The forwarding class specified in the fixed classifier on the native FC interface must be a lossless forwarding class. You cannot apply a fixed classifier to a native FC interface unless the associated forwarding class is lossless. (The forwarding class must be one of the two default lossless forwarding classes, or you must explicitly configure the forwarding class with the *no-loss* drop attribute.)
- The lossless forwarding class and IEEE 802.1p priority configuration must match on the FCoE-FC gateway native FC interfaces and Ethernet interfaces:
 - The same IEEE 802.1p priority (code point) must be enabled for PFC on the Ethernet ingress interfaces, classified to the lossless forwarding class used in the native FC interface fixed classifier, and set as the rewrite value on the native FC interfaces.
 - The same lossless forwarding class must be used in the fixed classifier on the native FC interfaces and in the classifier configuration on the Ethernet interfaces.

Fate Sharing

To ensure that congestion on one interface does not affect the fate of traffic on a native FC interface on which you remap priorities, avoid fate sharing (different traffic flows receiving the same CoS treatment) configurations.

You can avoid fate sharing by ensuring that the remapping priority (code point) on the native FC interface is classified only to the forwarding class used in the fixed classifier on all other interfaces. For example, if you configure a fixed classifier on an FC interface that classifies all of the traffic into lossless forwarding class *myfcoe1* and remaps the priority to priority 5 (IEEE 802.1p code point 101), then in all other classifier configurations on all other interfaces, priority 5 should always be classified to forwarding class *myfcoe1*. If you classify priority 6 on another interface to forwarding class *myfcoe1*, then congestion on priority 6 traffic affects priority 5 traffic unfairly.

Related Documentation

- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 356](#)

- [Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces \(NP_Ports\)](#) on page 423

Understanding CoS Buffer Configuration

Each QFX3500 and QFX3600 switch has 9 MB of Packet Forwarding Engine (PFE) wide common packet buffer memory that is used to store packets on interface queues. Each QFX5100 switch has 12 MB of PFE wide common packet buffer memory. The buffer memory has separate ingress and egress accounting to make accept, drop, or pause decisions. Because the switch has a single pool of memory with separate ingress and egress accounting, the full amount of buffer memory is available from both the ingress and the egress perspective. Packets are accounted for as they enter and leave the switch, but there is no concept of a packet arriving at an ingress buffer and then being moved to an egress buffer.

The buffers are divided into two pools from both an ingress and an egress perspective:

1. *Shared buffers* are a global memory pool that the switch allocates dynamically to ports as needed, so the buffers are shared among the switch ports.
2. *Dedicated buffers* are a memory pool divided equally among the switch ports. Each port receives a minimum guaranteed amount of buffer space, dedicated to each port, not shared among ports.



NOTE: Lossless traffic is traffic on which you enable priority-based flow control (PFC) to ensure lossless transport. Lossless traffic does not refer to best-effort traffic on a link enabled for Ethernet PAUSE (IEEE 802.3x).

The switch reserves nonconfigurable buffer space to ensure that ports and queues receive a minimum memory allocation. You can configure how the system uses the rest of the buffer space to optimize the allocation for your mix of network traffic. You can configure the percentage of available buffer space used as shared buffer space versus dedicated buffer space. You can also configure how shared buffer space is allocated to different types of traffic. You can optimize the buffer settings for the traffic on your network.

The default buffer configuration is designed for networks that have a balance of best-effort and lossless traffic.

The default class-of-service configuration provides two lossless forwarding classes (**fcoe** and **no-loss**), a best-effort unicast forwarding class, a network control traffic forwarding class, and one multidestination (multicast, broadcast, and destination lookup fail) forwarding class. Each default forwarding class maps to a different default output queue. The default configuration allocates the buffers in a manner that supports a moderate amount of lossless traffic while still providing the ability to absorb bursts in best-effort traffic transmission.

Changing the buffer settings changes the abilities of the buffers to absorb traffic bursts and handle lossless traffic. For example, networks with mostly best-effort traffic require

allocating most of the shared buffer space to best-effort buffers. This provides deep, flexible buffers that can absorb traffic bursts with minimal packet loss, at the expense of buffer availability for lossless traffic.

Conversely, networks with mostly lossless traffic require allocating most of the shared buffer space to lossless headroom buffers. This prevents packet loss on lossless flows at the expense of absorbing bursty best-effort traffic efficiently.



CAUTION: Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

This topic describes the buffer architecture and settings:

- [Buffer Pools on page 132](#)
- [Default Buffer Pool Values on page 140](#)
- [Shared Buffer Configuration Recommendations for Different Network Traffic Scenarios on page 143](#)
- [Optimizing Buffer Configuration on page 147](#)
- [General Buffer Configuration Rules and Considerations on page 148](#)

Buffer Pools

From both an ingress and an egress perspective, the PFE buffer is split into two main pools, a shared buffer pool and a dedicated buffer pool that ensures a minimum allocation to each port. You can configure the amount of buffer space allocated to each of the two pools. A portion of the buffer space is reserved so that there is always a minimum amount of shared and dedicated buffer space available to each port.

- **Shared buffer pool**—A global memory space that all of the ports on the switch share dynamically as they need buffers. The shared buffer pool is further partitioned into buffers for best-effort unicast, best-effort multdestination (broadcast, multicast, and destination lookup fail), and PFC (lossless) traffic types. You can allocate global shared memory space to buffer partitions to better support different mixes of network traffic. The larger the shared buffer pool, the better the switch can absorb traffic bursts because more shared memory is available for the traffic.
- **Dedicated buffer pool**—A reserved global memory space allocated equally to each port. The switch reserves a minimum dedicated buffer pool that is not user-configurable. You can divide the dedicated buffer allocation for a port among the port queues on a per-port, per-queue basis. (For example, this enables you to dedicate more buffer space to queues that transport lossless traffic.)

A larger dedicated buffer pool means a larger amount of dedicated buffer space for each port, so congestion on one port is less likely to affect traffic on another port because the traffic does not need to use as much shared buffer space. However, the larger the dedicated buffer pool, the less bursty traffic the switch can handle because there is less dynamic shared buffer memory.

You can configure the way the available unreserved portion of the buffer space is allocated to the global shared buffer pool and to the dedicated shared buffer pool by configuring the ingress and egress shared buffer percentages.

By default, 100 percent of the available unreserved buffer space is allocated to the shared buffer pool. If you change the percentage of space allocated to the shared buffer, the available buffer space that is not allocated to the shared buffer is allocated to the dedicated buffer. For example, if you configure the ingress shared buffer pool as 80 percent, the remaining 20 percent of the available buffer space is allocated to the dedicated buffer pool and divided equally across the ports.



NOTE: When 100 percent of the available (user-configurable) buffers are allocated to the shared buffer pool, the switch still reserves a minimum dedicated buffer pool.

You can separately configure ingress and egress shared buffer pool allocations. You can also partition the ingress and egress shared buffer pool to allocate percentages of the shared buffer pool to specific types of traffic. If you do not use the default configuration or one of the recommended configurations, pay particular attention to the ingress configuration of the lossless and lossless headroom buffers (these buffers handle PFC pause during periods of congestion) and to the egress configuration of the best-effort buffers to handle incast congestion (multiple synchronized sources sending data to the same receiver in parallel).

In addition to the shared buffer pool and the dedicated buffer pool, there is also a small ingress global headroom buffer pool that is reserved and is not configurable.

When contention for buffer space occurs, the switch uses an internal algorithm to ensure that the buffer pools are distributed fairly among competing flows. When traffic for a given flow exceeds the amount of dedicated port buffer reserved for that flow, the flow begins to consume memory from the dynamic shared buffer pool. Competing flows compete for shared buffer memory with other flows that also have exhausted their dedicated buffers. When there is no congestion, there are no competing flows.

- [Buffer Handling of Lossless Flows \(PFC\) Versus Ethernet PAUSE on page 133](#)
- [Shared Buffer Pool and Partitions on page 134](#)
- [Dedicated Port Buffer Pool and Buffer Allocation to Queues on page 135](#)
- [Trade-off Between Shared Buffer Space and Dedicated Buffer Space on page 139](#)
- [Order of Buffer Consumption on page 140](#)

[Buffer Handling of Lossless Flows \(PFC\) Versus Ethernet PAUSE](#)

When we discuss lossless buffers in the following sections, we mean buffers that handle traffic on which you enable PFC to ensure lossless transport. The lossless buffers are not used for best-effort traffic on a link on which you enable Ethernet PAUSE (IEEE 802.3x). The lossless ingress and egress shared buffers, and the ingress lossless headroom shared buffer, are used only for traffic on which you enable PFC.



NOTE: To support lossless flows, you must configure the appropriate data center bridging capabilities (PFC, DCBX, or ETS) and scheduling properties.

Shared Buffer Pool and Partitions

The shared buffer pool is a global memory space that all of the ports on the switch share dynamically as they need buffers. The switch uses the shared buffer pool to absorb traffic bursts after the dedicated buffer pool for a port is exhausted.

You can divide both the ingress shared buffer pool and the egress shared buffer pool into three partitions to allocate percentages of each buffer pool to different types of traffic. When you partition the ingress or egress shared buffer pool:

- If you explicitly configure one ingress shared buffer partition, you must explicitly configure all three ingress shared buffer partitions. (You either explicitly configure all three ingress partitions or you use the default setting for all three ingress partitions.)

If you explicitly configure one egress shared buffer partition, you must explicitly configure all three egress shared buffer partitions. (You either explicitly configure all three egress partitions or you use the default setting for all three egress partitions.)

The switch returns a commit error if you do not explicitly configure all three partitions when configuring the ingress or egress shared buffer partitions.

- The combined percentages of the three ingress shared buffer partitions must total exactly 100 percent.

The combined percentages of the three egress shared buffer partitions must total exactly 100 percent.

When you explicitly configure ingress or egress shared buffer partitions, the switch returns a commit error if the total percentage of the three partitions does not equal 100 percent.

- If you explicitly partition one set of shared buffers, you do not have to explicitly partition the other set of shared buffers. For example, you can explicitly configure the ingress shared buffer partitions and use the default egress shared buffer partitions. However, if you change the buffer partitions for the ingress buffer pool to match the expected types of traffic flows, you would probably also want to change the buffer partitions for the egress buffer pool to match those traffic flows.

You can configure the percentage of available unreserved buffer space allocated to the shared buffer pool. Space that you do not allocate to the shared buffer pool is added to the dedicated buffer pool and divided equally among the ports. The default configuration allocates 100 percent of the unreserved ingress and egress buffer space to the shared buffers.

Configuring the ingress and egress shared buffer pool partitions enables you to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic.

Ingress Shared Buffer Pool Partitions

You can configure three ingress buffer pool partitions:

- Lossless buffers—Shared buffer pool for all lossless ingress traffic. The recommended minimum value for lossless buffers is 5 percent.
- Lossless headroom buffers—Shared buffer pool for packets received while a pause is asserted. If PFC is enabled on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers for which the recommended value can be less than 5 percent.)
- Lossy buffers—Shared buffer pool for all best-effort ingress traffic (best-effort unicast, multdestination, and strict-high priority traffic). The recommended minimum value for best-effort buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and best-effort buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. If you explicitly configure an ingress shared buffer partition, you must explicitly configure all three ingress buffer partitions, even if the lossless headroom buffer partition has a value of 0 (zero) percent.

Egress Shared Buffer Pool Partitions

You can configure three egress buffer pool partitions:

- Lossless buffers—Shared buffer pool for all lossless egress queues. The recommended minimum value for lossless buffers is 5 percent.
- Lossy buffers—Shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The recommended minimum value for best-effort buffers is 5 percent.
- Multicast buffers—Shared buffer pool for all multdestination (multicast, broadcast, and destination lookup fail) egress queues. The recommended minimum value for multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and should have a value of at least 5 percent. If you explicitly configure an egress shared buffer partition, you must explicitly configure all three egress buffer partitions, and each partition should have a value of at least 5 percent.

Dedicated Port Buffer Pool and Buffer Allocation to Queues

The global dedicated buffer pool is memory that is allocated equally to each port, so each port receives a guaranteed minimum amount of buffer space. Dedicated buffers

are not shared among ports. Each port receives an equal proportion of the dedicated buffer pool.

The amount of dedicated buffer space is not user-configurable and depends on the percentage of available nonreserved buffers allocated to the shared buffers. (The dedicated buffer space is equal to the minimum reserved port buffers plus the remainder of the available nonreserved buffers that are not allocated to the shared buffer pool.)

When traffic enters and exits the switch, the switch ports use their dedicated buffers to store packets. If the dedicated buffers are not sufficient to handle the traffic, the switch uses shared buffers. The only way to increase the dedicated buffer pool is to decrease the shared buffer pool from its default value of 100 percent of available unreserved buffers.



NOTE: If 100 percent of the available unreserved buffers are allocated to the shared buffer pool, the switch still reserves a minimum dedicated buffer pool.

The larger the shared buffer pool, the better the burst absorption across the ports. The larger the dedicated buffer pool, the larger the amount of dedicated buffer space for each port. The greater the dedicated buffer space, the less likely that congestion on one port can affect traffic on another port, because the traffic does not need to use as much shared buffer space.

Allocating Dedicated Port Buffers to Queues

You can divide the dedicated buffer allocation for an egress port among the port queues by including the **buffer-size** statement in the scheduler configuration. This enables you to control the egress port dedicated buffer allocation on a per-port, per-queue basis. (For example, this enables you to dedicate more buffer space to queues that transport lossless traffic, or to stop the port from reserving buffers for queues that do not carry traffic.) Egress dedicated port buffer allocation is a hierarchical structure that allocates a global dedicated buffer pool evenly among ports, and then divides the allocation for each port among the port queues.

By default, ports divide their allocation of dedicated buffers among their egress queues in the same proportion as the default scheduler sets the minimum guaranteed transmission rates (the **transmit-rate** option) for traffic. Only the queues included in the default scheduler receive bandwidth and dedicated buffers, in the proportions shown in [Table 50 on page 136](#):

Table 50: Default Dedicated Buffer Allocation to Egress Queues (Based on Default Scheduler)

Forwarding Class	Queue	Minimum Guaranteed Bandwidth (transmit-rate)	Proportion of Reserved Dedicated Port Buffers
best-effort	0	5%	5%
fcoe	3	35%	35%

Table 50: Default Dedicated Buffer Allocation to Egress Queues (Based on Default Scheduler) (*continued*)

Forwarding Class	Queue	Minimum Guaranteed Bandwidth (transmit-rate)	Proportion of Reserved Dedicated Port Buffers
no-loss	4	35%	35%
network-control	7	5%	5%
mcast	8	20%	20%

In the default configuration, no egress queues other than the ones shown in [Table 50 on page 136](#) receive an allocation of dedicated port buffers.



NOTE: The switch uses hierarchical scheduling to control port and queue bandwidth allocation, as described in “[Understanding CoS Hierarchical Port Scheduling \(ETS\)](#)” on page 81 and shown in “[Example: Configuring CoS Hierarchical Port Scheduling \(ETS\)](#)” on page 220. For egress queue buffer size configuration, when you attach a traffic control profile (includes the queue scheduler information) to a port, the dedicated egress buffers on the port are divided among the queues as configured in the scheduler.

If you do not want to use the default allocation of dedicated port buffers to queues, use the **buffer-size** option in the scheduler that is attached to the port to configure the queue allocation. You can configure the dedicated buffer allocation to queues in two ways:

- As a percentage—The queue receives the specified percentage of dedicated port buffers when the queue is mapped to the scheduler and the scheduler is attached to a port.
- As a remainder—After the port services the queues that have an explicit percentage buffer size configuration, the remaining dedicated port buffer space is divided equally among the other queues to which a scheduler is attached. (No default or explicit scheduler for a queue means no dedicated buffer allocation for that queue.) If you configure a scheduler and you do not specify a buffer size as a percentage, *remainder* is the default setting.



NOTE: The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.

Configuring Dedicated Port Buffer Allocation to Queues

In a port configuration that includes multiple forwarding class sets, with multiple forwarding classes mapped to multiple schedulers, the allocation of port dedicated buffers to queues depends on the mix of queues with buffer sizes configured as explicit percentages and queues configured with (or defaulted to) the **remainder** option.

The best way to demonstrate how using the percentage and remainder options affects dedicated port buffer allocation to queues is by showing an example of queue buffer allocation, and then showing how the queue buffer allocation changes when you add another forwarding class (queue) to the port.

[Table 51 on page 138](#) shows an initial configuration that includes four forwarding class sets, the five default forwarding classes (mapped to the five default queues for those forwarding classes), the **buffer-size** option configuration, and the resulting buffer allocation for each queue. [Table 52 on page 139](#) shows the same configuration after we add another forwarding class (best-effort-2, mapped to queue 1) to the best-effort forwarding class set. Comparing the buffer allocations in each table shows you how adding another queue affects buffer allocation when you use remainders and explicit percentages to configure the buffer allocation for different queues.

Table 51: Egress Queue Dedicated Buffer Allocation (Example 1)

Forwarding Class Set (Priority Group)	Forwarding Class	Queue	Scheduler Buffer Size Configuration	Buffer Allocation per Queue (Percentage)
fc-set-be	best-effort	0	10%	10%
fc-set-lossless	fcoe	3	20%	20%
	no-loss	4	40%	40%
fc-set-strict-high	network-control	7	remainder	15%
fc-set-mcast	mcast	8	remainder	15%

In this first example, 70 percent of the egress port dedicated buffer pool is explicitly allocated to the best-effort, fcoe, and no-loss queues. The remaining 30 percent of the port dedicated buffer pool is split between the two queues that use the **remainder** option (network-control and mcast), so each queue receives 15 percent of the dedicated buffer pool.

Now we add another forwarding class (queue) to the best-effort priority group (fc-set-be) and configure it with a buffer size of *remainder* instead of configuring a specific percentage. Because a third queue now shares the remaining dedicated buffers, the queues that share the remainder receive fewer dedicated buffers, as shown in [Table 52 on page 139](#). The queues with explicitly configured percentages receive the configured percentage of dedicated buffers.

Table 52: Egress Queue Dedicated Buffer Allocation with Another Remainder Queue (Example 2)

Priority Group (fc-set)	Forwarding Class	Queue	Scheduler Buffer Size Configuration	Buffer Allocation per Queue (Percentage)
fc-set-be	best-effort	0	10%	10%
	best-effort-2	1	remainder	10%
fc-set-lossless	fcoe	3	20%	20%
	no-loss	4	40%	40%
fc-set-strict-high	network-control	7	remainder	10%
fc-set-mcast	mcast	8	remainder	10%

The two tables show how the port divides the dedicated buffer space that remains after servicing the queues that have an explicitly configured percentage of dedicated buffer space.

Trade-off Between Shared Buffer Space and Dedicated Buffer Space

The trade-off between shared buffer space and dedicated buffer space is:

- Shared buffers provide better absorption of traffic bursts because there is a larger pool of dynamic buffers that ports can use as needed to handle the bursts. However, all flows that exhaust their dedicated buffer space compete for the shared buffer pool. A larger shared buffer pool means a smaller dedicated buffer pool, and therefore more competition for the shared buffer pool because more flows exhaust their dedicated buffer allocation. Too much shared buffer space results in no single flow receiving very much shared buffer space, to maintain fairness when many flows contend for that space.
- Dedicated buffers provide guaranteed buffer space to each port. The larger the dedicated buffer pool, the less likely that congestion on one port affects traffic on another port, because the traffic does not need to use as much shared buffer space. However, less shared buffer space means less ability to dynamically absorb traffic bursts.

For optimal burst absorption, the switch needs enough dedicated buffer space to avoid persistent competition for the shared buffer space. When fewer flows compete for the shared buffers, the flows that need shared buffer space to absorb bursts receive more of the shared buffer because fewer flows exhaust their dedicated buffer space.

The default configuration and all of the recommended configurations allocate 100 percent of the user-configurable memory space to the global shared buffer pool because the amount of space reserved for dedicated buffers provides enough space to avoid persistent competition for dynamic shared buffers. This results in fewer flows competing for the shared buffers, so the competing flows receive more of the buffer space.

Order of Buffer Consumption

The total buffer pool is divided into ingress and egress shared buffer pools and dedicated buffer pools. When traffic flows through the switch, the buffer space is used in a particular order that depends on the type of traffic.

On ingress, the order of buffer consumption is:

- Best-effort unicast traffic:
 1. Dedicated buffers
 2. Shared buffers
 3. Global headroom buffers (very small)
- Lossless unicast traffic:
 1. Dedicated buffers
 2. Shared buffers
 3. Lossless headroom buffers
 4. Global headroom buffers (very small)
- Multidestination traffic:
 1. Dedicated buffers
 2. Shared buffers
 3. Global headroom buffers (very small)

On egress, the order of buffer consumption is the same for unicast best-effort, lossless unicast, and multidestination traffic:

- Dedicated buffers
- Shared buffers

In all cases on all ports, the switch uses the dedicated buffer pool first and the shared buffer pool only after the dedicated buffer pool for the port or queue is exhausted. This reserves the maximum amount of dynamic shared buffer space to absorb traffic bursts.

Default Buffer Pool Values

You can view the default or configured ingress and egress buffer pool values in KB units using the **show class-of-service shared-buffer** operational command. You can view the configured shared buffer pool values in percent units using the **show configuration class-of-service shared-buffer** operational command.

This section provides the default total buffer, shared buffer, and dedicated buffer values for QFX Series switches.

- [Total Buffer Pool Size on page 141](#)
- [Shared Buffer Pool Default Values on page 141](#)
- [Dedicated Buffer Pool Default Values on page 143](#)

Total Buffer Pool Size

The total buffer pool is common memory that has separate ingress and egress accounting, so the full buffer pool is available from both the ingress and egress perspective. The total buffer pool consists of the dedicated buffer space and the shared buffer space. The size of the total buffer pool is not user-configurable, but the allocation of buffer space to the dedicated and shared buffer pools is user-configurable.

On QFX3500 and QFX3600 switches, the combined total size of the ingress and egress buffer pools is approximately 9 MB (exactly 9360 KB).

On QFX5100 switches, the combined total size of the ingress and egress buffer pools is approximately 12 MB (exactly 12480 KB).

Shared Buffer Pool Default Values

The QFX5100 switch has a larger shared buffer pool (12 MB) than QFX3500 and QFX3600 switches (9 MB). However, the allocation of shared buffer space to the individual ingress and egress buffer pools is the same on a percentage basis, even though the absolute values are different. For example, the default ingress lossless buffer is 9 percent of the total shared ingress buffer space on QFX5100, QFX3500, and QFX3600 switches, even though the default absolute value of the ingress lossless buffer is 861.05KB on QFX5100 switches and 648.18KB on QFX3500 and QFX3600 switches.

This section describes the default values in percent and in KB for the shared ingress and shared egress buffers.

- [Shared Ingress Buffer Default Values on page 141](#)
- [Shared Egress Buffer Default Values on page 142](#)

Shared Ingress Buffer Default Values

The QFX5100 switch has a larger shared ingress buffer than the QFX3500 and QFX3600 switches. [Table 53 on page 141](#) shows the default ingress shared buffer allocation values in KB units for QFX5100 switches.

Table 53: QFX5100 Switch Default Shared Ingress Buffer Values (KB)

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
9567.19 KB	861.05 KB	4305.23 KB	4400.91 KB

[Table 54 on page 142](#) shows the default ingress shared buffer allocation values in KB units for QFX3500 and QFX3600 switches.

Table 54: QFX3500 and QFX3600 Switch Default Shared Ingress Buffer Values (KB)

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
7202 KB	648.18 KB	3240.9 KB	3312.92 KB

Table 55 on page 142 shows the default ingress shared buffer allocation values as percentages for QFX5100, QFX3500, and QFX3600 switches. (If you change the default shared buffer allocation, you configure the change as a percentage.)

Table 55: Default Shared Ingress Buffer Values (Percentage)

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
100%	9%	45%	46%

Shared Egress Buffer Default Values

The QFX5100 switch has a larger shared egress buffer than the QFX3500 and QFX3600 switches. Table 56 on page 142 shows the default egress shared buffer allocation values in KB units for QFX5100 switches.

Table 56: QFX5100 Switch Default Shared Egress Buffer Values (KB)

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
8736 KB	4368 KB	2708.16 KB	1659.84 KB

Table 57 on page 142 shows the default egress shared buffer allocation values in KB units.

Table 57: QFX3500 and QFX3600 Switch Default Shared Egress Buffer Values (KB)

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
6656 KB	3328 KB	2063.36 KB	1264.64 KB

Table 58 on page 142 shows the default egress shared buffer allocation values as percentages.

Table 58: Default Shared Egress Buffer Values (Percentage)

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	50%	31%	19%

Dedicated Buffer Pool Default Values

The system reserves ingress and egress dedicated buffer pools that are divided equally among the switch ports. By default, the system allocates 100 percent of the available unreserved buffer space to the shared buffer pool. If you reduce the percentage of available unreserved buffer space allocated to the shared buffer pool, the remaining unreserved buffer space is added to the dedicated buffer pool allocation. You configure the amount of dedicated buffer pool space by reducing (or increasing) the percentage of buffer space allocated to the shared buffer pool. You do not directly configure the dedicated buffer pool allocation.

The default dedicated buffer pool values for QFX3500 and QFX3600 switches in KB units are:

- Ingress dedicated buffer—2158 KB
- Egress dedicated buffer—2704.0 KB

The default dedicated buffer pool values for QFX5100 switches in KB units are:

- Ingress dedicated buffer—2912.81 KB
- Egress dedicated buffer—3744 KB

Shared Buffer Configuration Recommendations for Different Network Traffic Scenarios

The way you configure the shared buffer pool depends on the mix of traffic on your network. This section provides shared buffer configuration recommendations for five basic network traffic scenarios:

- Balanced traffic—The network carries a balanced mix of unicast best-effort, lossless, and multicast traffic. (This is the default configuration.)
- Best-effort unicast traffic—The network carries mostly unicast best-effort traffic.
- Best-effort traffic with Ethernet PAUSE (IEEE 802.3X) enabled—The network carries mostly best-effort traffic with Ethernet PAUSE enabled on the links.
- Best-effort multicast traffic—The network carries mostly multicast best-effort traffic.
- Lossless traffic—The network carries mostly lossless traffic (traffic on which PFC is enabled).



NOTE: Lossless traffic is defined as traffic on which you enable PFC to ensure lossless transport. Lossless traffic does not refer to best-effort traffic on a link on which you enable Ethernet PAUSE. Start with the recommended profiles for each network traffic scenario, and adjust them if necessary for your network traffic conditions.



CAUTION: Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete. This includes changing the default configuration to one of the recommended configurations.

Because you configure buffer allocations in percentages, the recommended allocations for each network traffic scenario are valid for all QFX Series switches. Use one of the following recommended shared buffer configurations for your network traffic conditions. Start with a recommended configuration, then make small adjustments to the buffer allocations to fine-tune the buffers if necessary as described in [“Optimizing Buffer Configuration” on page 147](#).

- [Balanced Traffic \(Default Configuration\) on page 144](#)
- [Best-Effort Unicast Traffic on page 144](#)
- [Ethernet PAUSE Traffic on page 145](#)
- [Best-Effort Multicast \(Multidestination\) Traffic on page 146](#)
- [Lossless Traffic on page 146](#)

Balanced Traffic (Default Configuration)

The default shared buffer configuration is optimized for networks that carry a balanced mix of best-effort unicast, lossless, and multidestination (multicast, broadcast, and destination lookup fail) traffic. The default class-of-service (CoS) configuration is also optimized for networks that carry a balanced mix of traffic.

We recommend that you use the default shared buffer configuration for networks that carry a balanced mix of traffic, especially if you are using the default CoS settings. [Table 59 on page 144](#) shows the default ingress shared buffer allocations:

Table 59: Default Ingress Shared Buffer Configuration

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
100%	9%	45%	46%

[Table 60 on page 144](#) shows the default egress shared buffer allocations:

Table 60: Default Egress Shared Buffer Configuration

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	50%	31%	19%

Best-Effort Unicast Traffic

If your network carries mostly best-effort (lossy) unicast traffic, then the default shared buffer configuration allocates too much buffer space to support lossless transport. Instead of wasting those buffers, we recommend that you use the following ingress

shared buffer settings (see [Table 61 on page 145](#)) and egress shared buffer settings (see [Table 62 on page 145](#)):

Table 61: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Unicast Traffic

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
100%	5%	0%	95%

Table 62: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Unicast Traffic

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	5%	75%	20%

See “[Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic](#)” on page 365 for an example that shows you how to configure the recommended buffer settings shown in [Table 61 on page 145](#) and [Table 62 on page 145](#).

Ethernet PAUSE Traffic

If your network carries mostly best-effort (lossy) traffic *and* enables Ethernet PAUSE on links, then the default shared buffer configuration allocates too much buffer space to the shared ingress buffer (Ethernet PAUSE traffic uses the dedicated buffers instead of shared buffers) and not enough space to the lossless-headroom buffers. We recommend that you use the following ingress shared buffer settings (see [Table 63 on page 145](#)) and egress shared buffer settings (see [Table 64 on page 145](#)):

Table 63: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Traffic and Ethernet PAUSE Enabled

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
70%	5%	80%	15%

Table 64: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Traffic and Ethernet PAUSE Enabled

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	5%	75%	20%

See “[Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled](#)” on page 371 for

an example that shows you how to configure the recommended buffer settings shown in [Table 61 on page 145](#) and [Table 62 on page 145](#).

Best-Effort Multicast (Multidestination) Traffic

If your network carries mostly best-effort (lossy) multicast traffic, then the default shared buffer configuration allocates too much buffer space to support lossless transport. Instead of wasting those buffers, we recommend that you use the following ingress shared buffer settings (see [Table 65 on page 146](#)) and egress shared buffer settings (see [Table 66 on page 146](#)):

Table 65: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Best-Effort Multicast Traffic

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
100%	5%	0%	95%

Table 66: Recommended Egress Shared Buffer Configuration for Networks with Mostly Best-Effort Multicast Traffic

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	5%	20%	75%

See “[Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic](#)” on [page 377](#) for an example that shows you how to configure the recommended buffer settings shown in [Table 65 on page 146](#) and [Table 66 on page 146](#).

Lossless Traffic

If your network carries mostly lossless traffic, then the default shared buffer configuration allocates too much buffer space to support best-effort traffic. Instead of wasting those buffers, we recommend that you use the following ingress shared buffer settings (see [Table 67 on page 146](#)) and egress shared buffer settings (see [Table 68 on page 146](#)):

Table 67: Recommended Ingress Shared Buffer Configuration for Networks with Mostly Lossless Traffic

Total Shared Ingress Buffer	Lossless Buffer	Lossless-Headroom Buffer	Lossy Buffer
100%	15%	80%	5%

Table 68: Recommended Egress Shared Buffer Configuration for Networks with Mostly Lossless Traffic

Total Shared Egress Buffer	Lossless Buffer	Lossy Buffer	Multicast Buffer
100%	90%	5%	5%

See “[Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic](#)” on page 383 for an example that shows you how to configure the recommended buffer settings shown in [Table 67 on page 146](#) and [Table 68 on page 146](#).

Optimizing Buffer Configuration

Starting from the default configuration or from a recommended buffer configuration, you can further optimize the buffer allocation to best support the mix of traffic on your network. Adjust the settings gradually to fine-tune the shared buffer allocation. Use caution when adjusting the shared buffer configuration, not just when you fine-tune the ingress and egress buffer partitions, but also when you fine-tune the total ingress and egress shared buffer percentage. (Remember that if you allocate less than 100 percent of the available buffers to the shared buffers, the remaining buffers are added to the dedicated buffers). Tuning the buffers incorrectly can cause problems such as ingress port congestion.



CAUTION: Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

The relationship between the sizes of the ingress buffer pool and the egress buffer pool affects when and where packets are dropped. The buffer pool sizes include the shared buffers and the dedicated buffers. In general, if there are more ingress buffers than egress buffers, the switch can experience ingress port congestion because egress queues fill before ingress queues can empty.

Use the [show class-of-service shared-buffer](#) operational command to see the sizes in kilobytes (KB) of the dedicated and shared buffers and of the shared buffer partitions.

For best-effort traffic (unicast and multdestination), the combined ingress lossy shared buffer partition and ingress dedicated buffers must be *less than* the combined egress lossy and multicast shared buffer partitions plus the egress dedicated buffers. This prevents ingress port congestion by ensuring that egress best-effort buffers are deeper than ingress best-effort buffers, and ensures that if packets are dropped, they are dropped at the egress queues. (Packets dropping at the ingress prevents the egress schedulers from working properly.)

For lossless traffic (traffic on which you enable PFC), the combined ingress lossless shared buffer partition and a reasonable portion of the ingress headroom buffer partition, plus the dedicated buffers, must be *less than* the total egress lossless shared buffer partition and dedicated buffers. (A reasonable portion of the ingress headroom buffer is approximately 20 to 25 percent of the buffer space, but this varies depending on how much buffer headroom is required to support the lossless traffic.) When these conditions are met, if there is ingress port congestion, the ingress port congestion triggers PFC on the ingress port to prevent packet loss. If the total lossless ingress buffers exceed the total lossless egress buffers, packets could be dropped at the egress instead of PFC being applied at the ingress to prevent packet loss.



NOTE: If you commit a buffer configuration for which the switch does not have sufficient resources, the switch might log an error instead of returning a commit error. After you commit a buffer configuration, check the syslog messages to ensure that the new buffer configuration did not fail to commit.

If the buffer configuration commits but you receive a syslog message that indicates the configuration cannot be implemented, you can:

- Reconfigure the buffers or reconfigure other parameters (for example, the PFC configuration, which affects the need for lossless headroom buffers and lossless buffers—the more priorities you pause, the more lossless and lossless headroom buffer space you need), then attempt the commit operation again.
- Roll back the switch to the last successful configuration.

If you receive a syslog message that says the buffer configuration cannot be implemented, you must take corrective action. If you do not fix the configuration or roll back to a previous successful configuration, the system behavior is unpredictable.

General Buffer Configuration Rules and Considerations

Keep the following rules and considerations in mind when you configure the buffers:

- Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.
- If you configure the ingress or egress shared buffer percentages as less than 100 percent, the remaining percentage of buffer space is added to the dedicated buffer pool.
- The sum of all of the ingress shared buffer partitions must equal 100 percent. Each partition must be configured with a value of at least 5 percent except the lossless headroom buffer, which can have a value of 0 percent.
- The sum of all of the egress shared buffer partitions must equal 100 percent. Each partition must be configured with a value of at least 5 percent.
- Lossless and lossless headroom shared buffers serve traffic on which you enable PFC, and do not serve traffic subject to Ethernet PAUSE.
- The switch uses the dedicated buffer pool first and the shared buffer pool only after the dedicated buffer pool for a port or queue is exhausted.
- Too little dedicated buffer space results in too much competition for shared buffer space.
- Too much dedicated buffer space results in poorer burst absorption because there is less available shared buffer space.
- Always check the syslog messages after you commit a new buffer configuration.
- The optimal buffer configuration for your network depends on the types of traffic on the network. If your network carries less traffic of a certain type (for example, lossless

traffic), then you can reduce the size of the buffers allocated to that type of traffic (for example, you can reduce the sizes of the lossless and lossless headroom buffers).

Related Documentation

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 365](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 371](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 377](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 383](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 425](#)

Understanding CoS WRED Drop Profiles

When the number of packets queued is greater than the ability of the switch to empty an output queue, the queue requires a method for determining which packets to drop to relieve the congestion. Weighted random early detection (WRED) drop profiles define the drop probability of packets as the output queue fills. During periods of congestion, as the output queue fills, the switch drops incoming packets as determined by a drop profile until the output queue becomes less congested.

Depending on the drop probabilities, a drop profile can drop many packets long before the buffer becomes full, or it can drop only a few packets even if the buffer is almost full.

You configure drop profiles in the drop profile section of the class-of-service (CoS) configuration hierarchy. You apply drop profiles using a drop profile map in each scheduler configuration. For each scheduler, you can configure separate drop profiles for each combination of loss priority (low, medium-high, and high) and protocol.

Drop profiles define the meaning of each of the loss priorities by setting the values for when to drop packets and the probability that packets will drop.

You can configure a maximum of 32 drop profiles.



NOTE: You cannot apply drop profiles to multidestination (multicast) queues.

Do not apply drop profiles to lossless flows such as FCoE traffic, because the corresponding queues require lossless behavior. Use priority-based flow control (PFC) to prevent packet drop.

- [Drop Profile Parameters on page 150](#)
- [Default Drop Profile on page 151](#)
- [Packet Drop Method on page 151](#)

- [Drop Profile Maps on page 152](#)
- [Congestion Prevention on page 152](#)
- [Configuring a WRED Drop Profile and Applying it to an Output Queue on page 153](#)

Drop Profile Parameters

Drop profiles specify two values:


- **Fill level**—The queue fullness value, which represents a percentage of the memory used to store packets in relation to the total amount of memory allocated to the queue.
- **Drop probability**—The percentage value that corresponds to the likelihood that an individual packet is dropped.

You set two queue fill levels and two drop probabilities in each drop profile. The two fill levels and the two drop probabilities create two pairs of values. The first fill level and the first drop probability create one value pair and the second fill level and the second drop probability create the second value pair.

The first fill level value specifies the percentage of queue fullness at which packets begin to drop, known as the drop start point. Until the queue reaches this level of fullness, no packets are dropped. The second fill level value specifies the percentage of queue fullness at which all packets are dropped, known as the drop end point.

The first drop probability value is always **0** (zero). This pairs with the drop start point and specifies that until the queue fullness level reaches the first fill level, no packets drop. When the queue fullness exceeds the drop start point, packets begin to drop until the queue exceeds the second fill level, when all packets drop. The second drop probability value, known as the maximum drop rate, specifies the likelihood of dropping packets when the queue fullness reaches the drop end point. As the queue fills from the drop start point to the drop end point, packets drop in a smooth, linear pattern (called an interpolated graph) as shown in [Figure 9 on page 150](#). After the drop end point, all packets drop.

Figure 9: WRED-Drop Profile Packet Drop



The thick line in [Figure 9 on page 150](#) shows the packet drop characteristics for a sample WRED profile. At the drop start point, the queue reaches a fill level of 30 percent. At the drop end point, the queue fill level reaches 50 percent, and the maximum drop rate is 80 percent.

No packets drop until the queue fill level reaches the drop start point of 30 percent. When the queue reaches the 30 percent fill level, packets begin to drop. As the queue fills, the percentage of packets dropped increases in a linear fashion. When the queue fills to the drop end point of 50 percent, the rate of packet drop has increased to the maximum drop rate of 80 percent. When the queue fill level exceeds the drop end point of 50 percent, all of the packets drop until the queue fill level drops below 50 percent.

Default Drop Profile

If you do not configure default profiles and apply them to queue schedulers, the switch uses the default drop profile for lossy traffic classes. In the default drop profile, when the fill level is 0 percent, the drop probability is 0 percent. When the fill level is 100 percent, the drop probability is 100 percent. As soon as packets arrive on a queue, the default profile might begin to drop packets.

Packet Drop Method

When a packet reaches the head of the queue, the switch generates a random number between 0 and 100. The switch plots the random number against the drop profile using the current fullness of the queue. When the random number falls above the graph line, the packet is transmitted. When the number falls below the graph line, the packet is dropped.

To create the linear drop pattern from the drop start point to the drop end point, the drop probabilities are derived using a linear approximation with eight sections, or steps, from the minimum queue fill level to the maximum queue fill level. The fill levels are divided into the eight sections equally, starting at the minimum fill level and ending at the maximum fill level. As the queue fills, the percentage of dropped packets increases. The percentage of packets dropped is based on the maximum drop rate.

For example, the default drop profile (which specifies a maximum drop rate of 100 percent) has the following drop probabilities at each section, or step, in the eight-section linear drop pattern:

- First section—The minimum drop probability is 6.25 percent of the maximum drop rate. The maximum drop probability is 12.5 percent of the maximum drop rate.
- Second section—The minimum drop probability is 18.75 percent of the maximum drop rate. The maximum drop probability is 25 percent of the maximum drop rate.
- Third section—The minimum drop probability is 30.25 percent of the maximum drop rate. The maximum drop probability is 37.5 percent of the maximum drop rate.
- Fourth section—The minimum drop probability is 43.75 percent of the maximum drop rate. The maximum drop probability is 50 percent of the maximum drop rate.
- Fifth section—The minimum drop probability is 56.25 percent of the maximum drop rate. The maximum drop probability is 62 percent of the maximum drop rate.

- Sixth section—The minimum drop probability is 68.75 percent of the maximum drop rate. The maximum drop probability is 75.5 percent of the maximum drop rate.
- Seventh section—The minimum drop probability is 81.25 percent of the maximum drop rate. The maximum drop probability is 87.5 percent of the maximum drop rate.
- Eighth section—The minimum drop probability is 92.75 percent of the maximum drop rate. The maximum drop probability is 100 percent of the maximum drop rate.

Packets drop even when there is no congestion, because packet drops begin at the drop start point regardless of whether congestion exists on the port. The default drop profile example represents the worst-case scenario, because the drop start point fill level is 0 percent, so packet drop begins when the queue starts to receive packets.

You can specify when packets begin to drop by configuring a drop start point at a fill level greater than 0 percent. For example, if you configure a drop profile that has a drop start point of 30 percent, packets do not drop until the queue is 30 percent full. We recommend that you configure drop profiles that are appropriate to your network traffic conditions.

The smaller the gap between the minimum drop rate (which is always 0) and the maximum drop rate, the smaller the gap between the minimum drop probability and the maximum drop probability at each section (step) of the linear drop pattern. The default drop profile, which has the maximum gap between the minimum drop rate (0 percent) and the maximum drop rate (100 percent), has the highest gap between the minimum drop probability and the maximum drop probability at each step. Configuring a lower maximum drop rate for a drop profile reduces the gap between the minimum drop probability and the maximum drop probability.

Drop Profile Maps

Drop profile maps are part of scheduler configuration. A drop profile map maps a drop profile to a loss priority and a protocol. Specifying the drop profile map in a scheduler associates the drop profile with the queues (forwarding classes) that you map to the scheduler in a scheduler map.

You configure loss priority for a queue in the classifier section of the CoS configuration hierarchy, and the loss priority is applied to the queue at the ingress interface.

Each scheduler can have multiple drop profile maps, one for each combination of loss priority and protocol.

Congestion Prevention

Configuring drop profiles on output queues prevents them from impacting other queues on the egress ports. If you do not configure drop profiles and map them to output queues, output queues without drop profiles can impact output queues on other egress ports, even if those queues are not experiencing congestion.

For example, if an ingress port forwards traffic to more than one egress port, and at least one of the egress ports experiences congestion, that can cause ingress port congestion. Ingress port congestion (ingress buffer exceeds its resource allocation) can cause frames to drop at the ingress port instead of at the egress port. Ingress port frame drop affects

all of the egress ports to which the congested ingress port forwards traffic, not just the congested egress port.



NOTE: Do not configure drop profiles for the `fcoe` and `no-loss` forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

Configuring a WRED Drop Profile and Applying it to an Output Queue

To configure a WRED packet drop profile and apply it to an output queue (using hierarchical scheduling):

1. Configure a drop profile using the statement **`set class-of-service drop-profiles profile-name interpolate fill-level drop-start-point fill-level drop-end-point drop-probability 0 drop-probability percentage`**.
2. Map the drop profile to a queue scheduler using the statement **`set class-of-service schedulers scheduler-name drop-profile-map loss-priority (low | medium-high | high) protocol any drop-profile profile-name`**. The name of the drop-profile is the name of the WRED profile configured in step 1.
3. Map the scheduler, which step 2 associates with the drop profile, to the output queue using the statement **`set class-of-service scheduler-maps map-name forwarding-class forwarding-class-name scheduler scheduler-name`**. The forwarding class identifies the output queue. Forwarding classes are mapped to output queues by default, and can be remapped to different queues by explicit user configuration. The scheduler name is the scheduler configured in step 2.
4. Associate the scheduler map with a traffic control profile using the statement **`set class-of-service traffic-control-profiles tcp-name scheduler-map map-name`**. The scheduler map name is the name configured in step 3.
5. Associate the traffic control profile with an interface using the statement **`set class-of-service interface interface-name forwarding-class-set forwarding-class-set-name output-traffic-control-profile tcp-name`**. The output traffic control profile name is the name of the traffic control profile configured in step 4.

The interface uses the scheduler map in the traffic control profile to apply the drop profile (and other attributes) to the output queue (forwarding class) on that interface. Because you can use different traffic control profiles to map different schedulers to different interfaces, the same queue number on different interfaces can handle traffic in different ways.

Related Documentation

- [Understanding Junos CoS Components on page 30](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring WRED Drop Profiles on page 279](#)
- [Example: Configuring Drop Profile Maps on page 281](#)
- [Example: Configuring Unicast Classifiers on page 273](#)

- [Configuring CoS WRED Drop Profiles on page 407](#)
- [Configuring CoS Drop Profile Maps on page 408](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 404](#)

Understanding CoS Rewrite Rules

As packets enter or exit a network, edge switches might be required to alter the class-of-service (CoS) settings of the packets. Rewrite rules set the value of the CoS bits within the header of the outgoing packet. Each rewrite rule reads the current forwarding class and loss priority associated with the packet, locates the chosen CoS value from a table, and writes this CoS value into the packet header, replacing the old CoS value. Rewrite rules must be assigned to an interface for rewrites to be activated.

You can apply (bind) one DSCP or DSCP IPv6 rewrite rule and one IEEE 802.1p rewrite rule to each interface as described in [“Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces” on page 57](#). You can also bind EXP rewrite rules to **family mpls** logical interfaces to rewrite the CoS bits of MPLS traffic.

You cannot apply both a DSCP and a DSCP IPv6 rewrite rule to the same interface. Each interface supports only one DSCP rewrite rule. Both IP and IPv6 packets use the same DSCP rewrite rule, regardless if the configured rewrite rule is DSCP or DSCP IPv6. You can apply an EXP rewrite rule on an interface that has DSCP or IEEE rewrite rules. Only MPLS traffic on **family mpls** interfaces uses the EXP rewrite rule.



NOTE: There are no default rewrite rules.

You can look at behavior aggregate (BA) classifiers and rewrite rules as two sides of the same coin. A BA classifier reads the CoS bits of incoming packets and classifies the packets into forwarding classes, then the system applies the CoS configured for the forwarding class to those packets. Rewrite rules change (rewrite) the CoS bits just before the packets leave the system so that the next switch or router can apply the appropriate level of CoS to the packets. When you apply a rewrite rule to an interface, the rewrite rule is the last CoS action performed on the packet before it is forwarded.

Rewrite rules alter CoS values in outgoing packets on the outbound interfaces of an edge switch to accommodate the policies of a targeted peer. This allows the downstream switch in a neighboring network to classify each packet into the appropriate service group.



NOTE: On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.

QFX Series does not have default rewrite rules. If you want to apply a rewrite rule to outgoing packets, you must explicitly configure the rewrite rule.



NOTE: Rewrite rules are applied *before* the egress filter is matched to traffic. Because the code point rewrite occurs before the egress filter is matched to traffic, the egress filter match is based on the rewrite value, not on the original code point value in the packet.

For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rule rewrites only the outer VLAN tag.

MPLS EXP rewrite rules apply only to **family mpls** logical interfaces. You cannot apply to an EXP rewrite rule to a physical interface. You can configure as many EXP rewrite rules as you want, but you can only use 16 EXP rewrite rules at any time on the switch. On a given logical interface, all pushed MPLS labels have the same EXP rewrite rule applied to them. You can apply different EXP rewrite rules to different logical interfaces on the same physical interface.



NOTE: If the switch is performing penultimate hop popping (PHP), EXP rewrite rules do not take effect. If both an EXP classifier and an EXP rewrite rule are configured on the switch, then the EXP value from the last popped label is copied into the inner label. If either an EXP classifier or an EXP rewrite rule (but not both) is configured on the switch, then the inner label EXP value is sent unchanged.

Related Documentation

- [Understanding Junos CoS Components on page 30](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules](#)
- [Defining CoS Rewrite Rules on page 427](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers](#)
- [Configuring a Global MPLS EXP Classifier](#)

Understanding DCB Features and Requirements

Data center bridging (DCB) is a set of enhancements to the IEEE 802.1 bridge specifications. DCB modifies and extends Ethernet behavior to support I/O convergence in the data center. I/O convergence includes but is not limited to the transport of Ethernet LAN traffic and Fibre Channel (FC) storage area network (SAN) traffic on the same physical Ethernet network infrastructure.



Video: [What is Data Center Bridging?](#)

A converged architecture saves cost by reducing the number of networks and switches required to support both types of traffic, reducing the number of interfaces required, reducing cable complexity, and reducing administration activities.

The Juniper Networks QFX Series supports the DCB features required to transport converged Ethernet and FC traffic while providing the class-of-service (CoS) and other characteristics FC requires for transmitting storage traffic. To accommodate FC traffic, DCB specifications provide:

- A flow control mechanism called priority-based flow control (PFC, described in IEEE 802.1Qbb) to help provide lossless transport.
- A discovery and exchange protocol for conveying configuration and capabilities among neighbors to ensure consistent configuration across the network, called Data Center Bridging Capability Exchange protocol (DCBX), which is an extension of Link Layer Data Protocol (LLDP, described in IEEE 802.1AB).
- A bandwidth management mechanism called enhanced transmission selection (ETS, described in IEEE 802.1Qaz).
- A congestion management mechanism called quantized congestion notification (QCN, described in IEEE 802.1Qau).

The switch supports the PFC, DCBX, and ETS standards but does not support QCN. The switch also provides the high-bandwidth interfaces (10-Gbps minimum) required to support DCB and converged traffic.

This topic describes the DCB standards and requirements the switch supports:

- [Lossless Transport on page 156](#)
- [ETS on page 157](#)
- [DCBX on page 158](#)

Lossless Transport

FC traffic requires lossless transport (defined as no frames dropped because of congestion). Standard Ethernet does not support lossless transport, but the DCB extensions to Ethernet along with proper buffer management enable an Ethernet network to provide the level of class of service (CoS) necessary to transport FC frames encapsulated in Ethernet over an Ethernet network.

This section describes these factors in creating lossless transport over Ethernet:

- [PFC on page 157](#)
- [Buffer Management on page 157](#)
- [Physical Interfaces on page 157](#)

PFC

PFC is a link-level flow control mechanism similar to Ethernet PAUSE (described in IEEE 802.3x). Ethernet PAUSE stops all traffic on a link for a period of time. PFC enables you to divide traffic on a link into eight priorities and stop the traffic of a selected priority without stopping the traffic assigned to other priorities on the link.

Pausing the traffic of a selected priority enables you to provide lossless transport for traffic assigned that priority and at the same time use standard lossy Ethernet transport for the rest of the link traffic.

Buffer Management

Buffer management is critical to the proper functioning of PFC, because if buffers are allowed to overflow, frames are dropped and transport is not lossless.

For each lossless flow priority, the switch requires sufficient buffer space to:

- Store frames sent during the time it takes to send the PFC pause frame across the cable between devices.
- Store the frames that are already on the wire when the sender receives the PFC pause frame.

The propagation delay due to cable length and speed, as well as processing speed, determines the amount of buffer space needed to prevent frame loss due to congestion.

The switch automatically sets the threshold for sending PFC pause frames to accommodate delay from cables as long as 150 meters (492 feet) and to accommodate large frames that might be on the wire when the switch sends the pause frame. This ensures that the switch sends pause frames early enough to allow the sender to stop transmitting before the receive buffers on the switch overflow.

Physical Interfaces

The switch supports 10-Gbps, full-duplex interfaces. The switch enables DCB capability only on 10-Gbps (or faster) Ethernet interfaces.

ETS

PFC divides traffic into up to eight separate streams (priorities, configured on the switch as forwarding classes) on a physical link. ETS enables you to manage the link bandwidth by:

- Grouping the priorities into priority groups (configured on the switch as forwarding class sets).

- Specifying the bandwidth available to each of the priority groups as a percentage of the total available link bandwidth.
- Allocating the bandwidth to the individual priorities in the priority group.

The available link bandwidth is the bandwidth remaining after servicing strict-high priority flows. We recommend that you always configure a shaping rate to limit the amount of bandwidth a strict-high priority flow can consume by including the [shaping-rate](#) statement in the **[edit class-of-service schedulers]** hierarchy on the strict-high priority scheduler. This prevents a strict-high priority from starving other queues on the port.

Managing link bandwidth with ETS provides several advantages:

- There is uniform management of all types of traffic on the link, both congestion-managed traffic and standard Ethernet traffic.
- When a priority group does not use all of its allocated bandwidth, other priority groups on the link can use that bandwidth as needed.

When a priority in a priority group does not use all of its allocated bandwidth, other priorities in the group can use that bandwidth.

The result is better bandwidth utilization, because priorities that consist of bursty traffic can share bandwidth during periods of low traffic transmission instead of consuming their entire bandwidth allocation when traffic loads are light.

- You can assign traffic types with different service needs to different priorities so that each traffic type receives appropriate treatment.
- Strict priority traffic retains its allocated bandwidth.

DCBX

DCB devices use DCBX to exchange configuration information with directly connected peers (switches and endpoints such as servers). DCBX is an extension of LLDP. If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit fails.

DCBX can:

- Discover the DCB capabilities of peers.
- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.

You can configure DCBX operation for PFC, ETS, and for Layer 2 and Layer 4 applications such as FCoE and iSCSI. DCBX is enabled or disabled on a per-interface basis.

Related Documentation

- *Overview of Fibre Channel on the QFX Series*
- *Understanding FCoE*
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)

- [Understanding DCBX on page 170](#)
- [*Understanding Fibre Channel Terminology*](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 241](#)

Understanding CoS Flow Control (Ethernet PAUSE and PFC)

Flow control supports lossless transmission by regulating traffic flows to avoid dropping frames during periods of congestion. Flow control stops and resumes the transmission of network traffic between two connected peer nodes on a full-duplex Ethernet physical link. Controlling the flow by pausing and restarting it prevents buffers on the nodes from overflowing and dropping frames. You configure flow control on a per-interface basis.

The QFX Series supports two methods of flow control:

- IEEE 802.3X Ethernet PAUSE
- IEEE 802.1Qbb priority-based flow control (PFC)

Ethernet PAUSE and PFC are link-level flow control mechanisms.

Ethernet PAUSE pauses transmission of all traffic on a physical Ethernet link.

PFC decouples the pause function from the physical Ethernet link and enables you to divide traffic on one link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that are mapped to forwarding classes and output queues. Each priority is mapped to a 3-bit IEEE 802.1p CoS code point flag in the VLAN header. You can enable PFC on one or more priorities (IEEE 802.1p code points) on a link. When PFC-enabled traffic is paused on a link, traffic that is not PFC-enabled continues to flow (or is dropped if congestion is severe enough).



Video: [Why Use PFC in a Data Center Network?](#)

Use Ethernet PAUSE when you want to prevent packet loss on all of the traffic on a link. Use PFC to prevent traffic loss only on specified types of traffic (for example, Fibre Channel over Ethernet traffic).



NOTE: Depending on the amount of traffic on a link or assigned to a priority, pausing traffic can cause ingress port congestion and spread congestion through the network.

Attempting to configure both Ethernet PAUSE and PFC on a link causes a commit error. Ethernet PAUSE and PFC are mutually exclusive configurations on an interface.

By default, all forms of flow control are disabled. You must explicitly enable flow control on interfaces to pause traffic.

- [Ethernet PAUSE on page 161](#)
- [PFC on page 165](#)
- [Lossless Transport Support Summary on page 168](#)

Ethernet PAUSE

Ethernet PAUSE is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends Ethernet PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to Ethernet PAUSE messages it receives from the connected peer to stop sending traffic. Ethernet PAUSE also works on aggregated Ethernet interfaces. For example, if the connected peer interfaces are called Node A and Node B:

- When the receive buffers on interface Node A reach a certain level of fullness, the interface generates and sends an Ethernet PAUSE message to the connected peer (interface Node B) to tell the peer to stop sending frames. The Node B buffers store frames until the time period specified in the Ethernet PAUSE frame elapses; then Node B resumes sending frames to Node A.
- When interface Node A receives an Ethernet PAUSE message from interface Node B, interface Node A stops transmitting frames until the time period specified in the Ethernet PAUSE frame elapses; then Node A resumes transmission. (The Node A transmit buffers store frames until Node A resumes sending frames to Node B.)

In this scenario, if Node B sends an Ethernet PAUSE frame with a time value of 0 to Node A, the 0 time value indicates to Node A that it can resume transmission. This happens when the Node B buffer empties to below a certain threshold and the buffer can once again accept traffic.

Symmetric flow control means an interface has the same Ethernet PAUSE configuration in both directions. The Ethernet PAUSE generation and Ethernet PAUSE response functions are both configured as enabled, or they are both disabled. You configure symmetric flow control by including the **flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

Asymmetric flow control allows you to configure the Ethernet PAUSE functionality in each direction independently on an interface. The configuration for generating Ethernet PAUSE messages and for responding to Ethernet PAUSE messages does not have to be the same. It can be enabled in both directions, disabled in both directions, or enabled in one direction and disabled in the other direction. You configure asymmetric flow control by including the **configured-flow-control** statement at the **[edit interfaces interface-name ether-options]** hierarchy level.

On any particular interface, symmetric and asymmetric flow control are mutually exclusive. Asymmetric flow control overrides and disables symmetric flow control. (If PFC is configured on an interface, the PFC configuration overrides Ethernet PAUSE flow control.) The QFX Series supports both symmetric and asymmetric flow control.

- [Symmetric Flow Control on page 162](#)
- [Asymmetric Flow Control on page 162](#)

Symmetric Flow Control

Symmetric flow control configures both the receive and transmit buffers in the same state. The interface can both send Ethernet PAUSE messages and respond to them (flow control is enabled), or the interface cannot send Ethernet PAUSE messages or respond to them (flow control is disabled).

When you enable symmetric flow control on an interface, the Ethernet PAUSE behavior depends on the configuration of the connected peer. With symmetric flow control enabled, the interface can perform any Ethernet PAUSE functions that the connected peer can perform. (When symmetric flow control is disabled, the interface does not send or respond to Ethernet PAUSE messages.)

Asymmetric Flow Control

Asymmetric flow control enables you to specify independently whether or not the interface receive buffer generates and sends Ethernet PAUSE messages to stop the connected peer from transmitting traffic, and whether or not the interface transmit buffer responds to Ethernet PAUSE messages it receives from the connected peer and stops transmitting traffic. The receive buffer configuration determines if the interface transmits Ethernet PAUSE messages, and the transmit buffer configuration determines if the interface receives and responds to Ethernet PAUSE messages:

- Receive buffers on—Enable Ethernet PAUSE transmission (generate and send Ethernet PAUSE frames)
- Transmit buffers on—Enable Ethernet PAUSE reception (respond to received Ethernet PAUSE frames)

You must explicitly set the flow control for both the receive buffer and the transmit buffer (**on** or **off**) to configure asymmetric Ethernet PAUSE. [Table 69 on page 162](#) describes the configured flow control state when you set the receive (Rx) and transmit (Tx) buffers on an interface:

Table 69: Asymmetric Ethernet PAUSE Flow Control Configuration

Receive (Rx) Buffer	Transmit (Tx) Buffer	Configured Flow Control State
On	Off	Interface generates and sends Ethernet PAUSE messages. Interface does not respond to Ethernet PAUSE messages (interface continues to transmit even if peer requests that the interface stop sending traffic).
Off	On	Interface responds to Ethernet PAUSE messages received from the connected peer, but does not generate or send Ethernet PAUSE messages. (The interface does not request that the connected peer stop sending traffic.)
On	On	Same functionality as symmetric Ethernet PAUSE. Interface generates and sends Ethernet PAUSE messages and responds to received Ethernet PAUSE messages.
Off	Off	Ethernet PAUSE flow control is disabled.

The configured flow control is the Ethernet PAUSE state configured on the interface.

On 1-Gigabit Ethernet interfaces, the QFX Series supports autonegotiation of Ethernet PAUSE with the connected peer. (The QFX Series does not support autonegotiation on 10-Gigabit Ethernet interfaces.) Autonegotiation enables the interface to exchange state advertisements with the connected peer so that the two devices can agree on the Ethernet PAUSE configuration. Each interface advertises its flow control state to the connected peer using a combination of the Ethernet PAUSE and ASM_DIR bits, as described in [Table 70 on page 163](#):

Table 70: Flow Control State Advertised to the Connected Peer (Autonegotiation)

Rx Buffer State	Tx Buffer State	PAUSE Bit	ASM_DIR Bit	Description
Off	Off	0	0	The interface advertises no Ethernet PAUSE capability. This is equivalent to disabling flow control on an interface.
On	On	1	0	The interface advertises symmetric flow control (both the transmission of Ethernet PAUSE messages and the ability to receive and respond to Ethernet PAUSE messages).
On	Off	0	1	The interface advertises asymmetric flow control (the transmission of Ethernet PAUSE messages, but not the ability to receive and respond to Ethernet PAUSE messages).
Off	On	1	1	The interface advertises both symmetric and asymmetric flow control. Although the interface does not generate and send Ethernet PAUSE requests to the peer, the interface supports both symmetric and asymmetric Ethernet PAUSE configuration on the peer because the peer is not affected if the peer does not receive Ethernet PAUSE requests. (If the interface responds to the peer's Ethernet PAUSE requests, that is sufficient to support either symmetric or asymmetric flow control on the peer.)

The flow control configuration on each switch interface interacts with the flow control configuration of the connected peer. Each peer advertises its state to the other peer. The interaction of the flow control configuration of the peers determines the flow control behavior (resolution) between them, as shown in [Table 71 on page 164](#). The first four columns show the Ethernet PAUSE configuration on the local QFX Series and on the

connected peer (also known as the link partner). The last two columns show the Ethernet PAUSE resolution that results from the local and peer configurations on each interface. This illustrates how the Ethernet PAUSE configuration of each interface affects the Ethernet PAUSE behavior on the other interface.



NOTE: In the Resolution columns of the table, disabling Ethernet PAUSE transmit means that the interface receive buffers do not generate and send Ethernet PAUSE messages to the peer. Disabling Ethernet PAUSE receive means that the interface transmit buffers do not respond to Ethernet PAUSE messages received from the peer.

Table 71: Asymmetric Ethernet PAUSE Behavior on Local and Peer Interfaces

Local Interface (QFX Series)		Peer Interface		Local Resolution	Peer Resolution
PAUSE Bit	ASM_DIR Bit	PAUSE Bit	ASM_DIR Bit		
0	0	Don't care	Don't care	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
0	1	0	Don't care	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
0	1	1	0	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
0	1	1	1	Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive	Disable Ethernet PAUSE transmit and enable Ethernet PAUSE receive
1	0	0	Don't care	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
1	0	1	Don't care	Enable Ethernet PAUSE transmit and receive	Enable Ethernet PAUSE transmit and receive
1	1	0	0	Disable Ethernet PAUSE transmit and receive	Disable Ethernet PAUSE transmit and receive
1	1	0	1	Enable Ethernet PAUSE receive and disable Ethernet PAUSE transmit	Enable Ethernet PAUSE transmit and disable Ethernet PAUSE receive
1	1	Don't care	Don't care	Enable Ethernet PAUSE transmit and receive	Enable Ethernet PAUSE transmit and receive



NOTE: For your convenience, [Table 71 on page 164](#) replicates Table 28B-3 of Section 2 of the IEEE 802.X specification.

PFC

PFC is a lossless transport and congestion relief feature that works by providing granular link-level flow control for each IEEE 802.1p code point (priority) on a full-duplex Ethernet link. When the receive buffer on a switch interface fills to a threshold, the switch transmits a pause frame to the sender (the connected peer) to temporarily stop the sender from transmitting more frames. The buffer threshold must be low enough so that the sender has time to stop transmitting frames and the receiver can accept the frames already on the wire before the buffer overflows. The switch automatically sets queue buffer thresholds to prevent frame loss.

When congestion forces one priority on a link to pause, all of the other priorities on the link continue to send frames. Only frames of the paused priority are not transmitted. When the receive buffer empties below another threshold, the switch sends a message that starts the flow again.

You configure PFC using a congestion notification profile (CNP). A CNP has two parts:

- **Input**—Specify the code point (or code points) on which to enable PFC, and optionally specify the maximum receive unit (MRU) and the cable length between the interface and the connected peer interface.
- **Output**—Specify the output queue or output queues that respond to pause messages from the connected peer.

You apply a PFC configuration by configuring a CNP on one or more interfaces. Each interface that uses a particular CNP is enabled to pause traffic with the priorities (code points) specified in that CNP. You can configure one CNP on an interface, and you can configure different CNPs on different interfaces. When you configure a CNP on an interface, ingress traffic that is mapped to a priority that the CNP enables for PFC is paused whenever the queue buffer fills to the pause threshold. (The pause threshold is not user-configurable.)

Configure PFC for a priority end to end along the entire data path to create a lossless lane of traffic on the network. You can selectively pause the traffic in any queue without pausing the traffic for other queues on the same link. You can create lossless lanes for traffic such as Fibre Channel over Ethernet (FCoE), LAN backup, or management, while using standard frame-drop congestion management for IP traffic on the same link.

Potential consequences of link-level flow control are:

- Ingress port congestion (configuring too many lossless flows can cause ingress port congestion)
- A paused priority that causes upstream devices to pause the same priority, thus spreading congestion back through the network

By definition, PFC supports symmetric pause only (as opposed to Ethernet PAUSE, which supports symmetric and asymmetric pause). With symmetric pause, a device can:

- Transmit pause frames to pause incoming traffic. (You configure this using the input stanza of a congestion notification profile.)

- Receive pause frames and stop sending traffic to a device whose buffer is too full to accept more frames. (You configure this using the output stanza of a congestion notification profile.)

Receiving a PFC frame from a connected peer pauses traffic on egress queues based on the IEEE 802.1p priorities that the PFC pause frame identifies. The priorities are 0 through 7. By default, the priorities map to queue numbers 0 through 7, respectively, and to specific forwarding classes, as shown in [Table 72 on page 166](#):

Table 72: Default PFC Priority to Queue and Forwarding Class Mapping

IEEE 802.1p Priority (Code Point)	Queue	Forwarding Class
0 (000)	0	best-effort
1 (001)	1	best-effort
2 (010)	2	best-effort
3 (011)	3	fcoe
4 (100)	4	no-loss
5 (101)	5	best-effort
6 (110)	6	network-control
7 (111)	7	network-control

For example, a received PFC pause frame that pauses priority 3 pauses output queue 3. If you do not want to use the default configuration, you can configure customized mapping of priorities to queues and forwarding classes.



NOTE: By convention, deployments with converged server access typically use IEEE 802.1p priority 3 for FCoE traffic. The default forwarding class configuration sets the fcoe forwarding class as a lossless forwarding class that is mapped to queue 3. The default classifier maps incoming priority 3 traffic to the fcoe forwarding class. *However, you must apply PFC to the entire FCoE data path to configure the end-to-end lossless behavior that FCoE traffic requires.*

If your network uses priority 3 for FCoE traffic, we recommend that you use the default configuration. If your network uses a priority other than 3 for FCoE traffic, you can configure lossless FCoE transport on any IEEE 802.1p priority as described in [“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 109](#) and [“Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway” on page 128](#).

You enable PFC on a priority by:

1. Specifying the IEEE 802.1p code point to pause in the input stanza of a CNP
2. Applying the CNP to the ingress interfaces on which you want to pause the traffic



CAUTION: Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

A change to the PFC configuration means any change to a CNP, including changing the input portion of the CNP (enabling or disabling PFC on a priority, or changing the MRU or cable-length values) or changing the output portion of the CNP that enables or disables output flow control on a queue. A PFC configuration change only affects ports that use the changed CNP.

The following actions change the PFC configuration:

- Deleting or disabling a PFC configuration (input or output) in a CNP that is in use on one or more interfaces. For example:
 1. An existing CNP with an input stanza that enables PFC on priorities 3, 5, and 6 is configured on interfaces xe-0/0/20 and xe-0/0/21.
 2. We disable the PFC configuration for priority 6 in the input CNP, and then commit the configuration.
 3. The PFC configuration change causes all traffic on interfaces xe-0/0/20 and xe-0/0/21 to stop until the PFC change has been implemented. When the PFC change has been implemented, traffic resumes.
- Configuring a CNP on an interface. (This changes the PFC state by enabling PFC on one or more priorities.)
- Deleting a CNP from an interface. (This changes the PFC state by disabling PFC on one or more priorities.)

When you associate the CNP with an interface, the interface uses PFC to send pause requests when the output queue buffer for the lossless traffic fills to the pause threshold.

Although unicast traffic and multideestination (multicast, broadcast, and destination lookup fail) traffic must use different classifiers, you can map a unicast queue (queue 0 through 7) and a multideestination queue (queue 8, 9, 10, or 11) to the same PFC priority so that both unicast and multicast traffic use that priority. Do not map multideestination traffic to lossless priorities. Starting with Junos OS Release 12.3, you can map one priority to multiple output queues.



NOTE: You can attach a maximum of one CNP to an interface, but you can create an unlimited number of CNPs that explicitly configure only the input stanza and use the default output stanza.

The output stanza of the CNP maps to a profile that interfaces use to respond to pause messages received from the connected peer. On standalone QFX3500 switches and QFX3600 switches, you can create two CNPs with an explicitly configured output stanza.

When a QFX3500 switch or a QFX3600 switch is a Node device in a QFabric system, you can create one CNP with an explicitly configured output stanza. (One fewer profile is available on QFabric systems because the system needs a default profile for fabric interfaces, which are not used as fabric interfaces when the switches are not part of a QFabric system. “[Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows](#)” on page 109 describes configuring output flow control.

Lossless Transport Support Summary

The QFX Series supports up to six lossless forwarding classes. For lossless transport, you must enable PFC on the IEEE 802.1p priorities (code points) mapped to lossless forwarding classes.



CAUTION: Any change to the PFC configuration on a port temporarily blocks the entire port (not just the priorities affected by the PFC change) so that the port can implement the change, then unblocks the port. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

The following limitation applies to support lossless transport on QFabric systems only:

- The internal fiber cable length from the QFabric system Node device to the QFabric system Interconnect device cannot exceed 150 meters.

The default CoS configuration provides two lossless forwarding classes, *fcoe* and *no-loss*. If you explicitly configure lossless forwarding classes, you must include the **no-loss** packet drop attribute to enable lossless behavior, or the traffic is not lossless. For both default and explicit lossless forwarding class configuration, you must configure CNP input stanzas to enable PFC on the priority of the lossless traffic and apply the CNPs to ingress interfaces.



NOTE: Junos OS Release 12.2 introduced changes to the way the QFX Series handles lossless forwarding classes (including the default `fcoe` and `no-loss` forwarding classes).

In Junos OS Release 12.1, either explicitly configuring the `fcoe` and `no-loss` forwarding classes or using the default configuration for these forwarding classes resulted in the same lossless behavior for traffic mapped to those forwarding classes.

However, in Junos OS Release 12.2, if you explicitly configure the `fcoe` or the `no-loss` forwarding class, that forwarding class is no longer treated as a lossless forwarding class. Traffic mapped to these forwarding classes is treated as lossy (best-effort) traffic. This is true even if the explicit configuration is exactly the same as the default configuration.

If your CoS configuration from Junos OS Release 12.1 or earlier includes the explicit configuration of the `fcoe` or the `no-loss` forwarding class, then when you upgrade to Junos OS Release 12.2, those forwarding classes are not lossless. To preserve the lossless treatment of these forwarding classes, delete the the explicit `fcoe` and `no-loss` forwarding class configuration before you upgrade to Junos OS Release 12.2.

See [“Overview of CoS Changes Introduced in Junos OS Release 12.2” on page 21](#) for detailed information about this change and how to delete an existing lossless configuration.

In Junos OS Release 12.3, the default behavior of the `fcoe` and `no-loss` forwarding classes is the same as in Junos OS Release 12.2. However, in Junos OS Release 12.3, you can configure up to six lossless forwarding classes. All explicitly configured lossless forwarding classes must include the new `no-loss` packet drop attribute or the forwarding class is lossy.

[“Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows” on page 109](#) provides detailed information about the explicit configuration of lossless priorities and about the default configuration of lossless priorities, including the input and output stanzas of the CNP.



NOTE: PFC and Ethernet PAUSE are used only on Ethernet interfaces. Fabric (fte) ports on QFabric systems (Node device fabric ports and Interconnect device fabric ports) use link-layer flow control (LLFC) to ensure the appropriate treatment of lossless traffic.

Related Documentation

- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 21](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 128](#)

- [Understanding DCB Features and Requirements on page 156](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 420](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 421](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 241](#)

Understanding DCBX

Data Center Bridging Capability Exchange protocol (DCBX) is an extension of Link Layer Data Protocol (LLDP). If you disable LLDP on an interface, that interface cannot run DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails. Data center bridging (DCB) devices use DCBX to exchange configuration information with directly connected peers.



Video: [What is DCBX Protocol?](#)

This topic describes:

- [DCBX Basics on page 170](#)
- [DCBX Modes and Support on page 171](#)
- [DCBX Attribute Types on page 174](#)
- [DCBX Application Protocol TLV Exchange on page 175](#)
- [DCBX and PFC on page 176](#)
- [DCBX and ETS on page 177](#)

DCBX Basics

DCBX can:

- Discover the DCB capabilities of peers.
- Detect DCB feature misconfiguration or mismatches between peers.
- Configure DCB features on peers.

You can configure DCBX operation for priority-based flow control (PFC), Layer 2 and Layer 4 applications such as FCoE and iSCSI, and ETS. DCBX is enabled or disabled on a per-interface basis.

By default, for PFC and ETS, DCBX automatically negotiates administrative state and configuration with each interface's connected peer. To enable DCBX negotiation for applications, you must configure the applications, map them to IEEE 802.1p code points in an application map, and apply the application map to interfaces.

The FCoE application only needs to be included in an application map when you want an interface to exchange type, length, and values (TLVs) for other applications in addition

to FCoE. If FCoE is the only application you want an interface to advertise, then you do not need to use an application map. For ETS, DCBX pushes the switch configuration to peers if they are set to learn the configuration from the switch (unless you disable sending the ETS recommendation TLV on interfaces in IEEE DCBX mode).

You can override the default behavior for PFC, for ETS, or for all applications mapped to an interface by turning off autonegotiation to force an interface to enable or disable that feature. You can also disable DCBX autonegotiation for applications on an interface by excluding those applications from the application map you apply to that interface or by deleting the application map from the interface.

The default autonegotiation behavior for applications that are mapped to an interface is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

During negotiation of capabilities, the switch can push the PFC configuration to an attached peer if the peer is configured as “willing” to learn the PFC configuration from other peers. The Juniper Networks switch does not support self autoprovisioning and does not change its configuration during autonegotiation to match the peer configuration. (The Juniper switch is not “willing” to learn the PFC configuration from peers.)



NOTE: When a port with DCBX enabled begins to exchange type, length, and value (TLV) entries, optional LLDP TLVs on that port are not advertised to neighbors, so that the switch can interoperate with a wider variety of converged network adapters (CNAs) and Layer 2 switches that support DCBX.

DCBX Modes and Support

This section describes DCBX support on the QFX Series:

- [DCBX Modes \(Versions\) on page 171](#)
- [Autonegotiation on page 173](#)
- [CNA Support for DCBX Modes on page 174](#)
- [Interface Support for DCBX on page 174](#)

DCBX Modes (Versions)

The QFX Series supports the two most common DCBX modes:

- IEEE DCBX—The newest DCBX version. Different TLVs have different subtypes (for example, the subtype for the ETS configuration TLV is 9); the IEEE DCBX Organizationally Unique Identifier (OUI) is 0x0080c2.
- DCBX version 1.01—The Converged Enhanced Ethernet (CEE) version of DCBX. It has a subtype of 2 and an OUI of 0x001b21.

IEEE DCBX and DCBX version 1.01 differ mainly in frame format. DCBX version 1.01 uses one TLV that includes all DCBX attribute information, which is sent as sub-TLVs. IEEE DCBX uses a unique TLV for each DCB attribute.



NOTE: The QFX Series does not support pre-CEE (pre-DCB) DCBX versions. Unsupported older versions of DCBX have a subtype of 1 and an OUI of 0x001b21. The QFX Series drops LLDP frames that contain pre-CEE DCBX TLVs.

Table 73 on page 172 summarizes the differences between IEEE DCBX and DCBX version 1.01, including show command output:

Table 73: Summary of Differences Between IEEE DCBX and DCBX Version 1.01

Characteristic	IEEE DCBX	DCBX Version 1.01
OUI	0x0080c2	0x001b21
Frame Format	Sends a separate, unique TLV for each DCBX attribute. For example, IEEE DCBX uses separate TLVs for ETS, PFC, and each application. Configuration and Recommendation information is sent in different TLVs	Sends one TLV that includes all DCBX attribute information organized in sub-TLVs. The “willing” bit determines whether or not an interface can change its configuration to match the connected peer.
Symmetric/asymmetric configuration with peer	Asymmetric or symmetric	Symmetric only
Differences in the show dcbx interface interface-name operational command	<ul style="list-style-type: none"> Synchronization information is not shown because symmetric configuration is not required. Operational state information is not shown because the operational states do not have to be symmetric. TLV type is shown because unique TLVs are sent for each DCBX attribute. ETS peer Configuration TLV and Recommendation TLV information is shown separately because they are different TLVs. 	<ul style="list-style-type: none"> Synchronization information is shown because symmetric configuration is required. Operational state information is shown because the operational states do have to be symmetric. TLV type is not shown because one TLV is used for all attribute information. Recommendation TLV is not sent (DCBX Version 1.01 uses the “willing” bit to determine whether or not an interface uses the peer interface configuration).

For more information about how each DCBX mode exchanges TLVs, see the following specifications:

- For DCBX version 1.01—
<http://www.ieee802.org/1/files/public/docs2008/az-wadkar-dcbx-capability-exchange-discovery-protocol-T108-v1.01.pdf>

- For IEEE DCBX—<http://www.ieee802.org/1/files/private/az-drafts/d2/802-1az-d2-4.pdf>



NOTE: As of Junos OS Release 12.2, this document is located in a private area of the IEEE website, and access requires a password from the IEEE organization. If you are not an IEEE member, you might not be able to access this document until it moves to the public area of the IEEE website.

You can configure interfaces to use the following DCBX modes:

- IEEE DCBX—The interface uses IEEE DCBX regardless of the configuration on the connected peer.
- DCBX version 1.01—The interface uses DCBX version 1.01 regardless of the configuration on the connected peer.
- Autonegotiation—The interface automatically negotiates with the connected peer to determine the DCBX version the peers use. Autonegotiation is the default DCBX mode.

If you configure a DCBX mode on an interface, the interface ignores DCBX protocol data units (PDUs) it receives from the connected peer if the PDUs do not match the DCBX version configured on the interface. For example, if you configure an interface to use IEEE DCBX and the connected peer sends DCBX version 1.01 LLDP PDUs, the interface ignores the version 1.01 PDUs. If you configure an interface to use DCBX version 1.01 and the peer sends IEEE DCBX LLDP PDUs, the interface ignores the IEEE DCBX PDUs.



NOTE: On interfaces that use the IEEE DCBX mode, the `show dcbx neighbors interface interface-name` operational command does not include application, PFC, or ETS operational state in the output.

Autonegotiation

Autonegotiation is the default DCBX mode. Each interface automatically negotiates with its connected peer to determine the DCBX version that both interfaces use to exchange DCBX information.

When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives one IEEE DCBX PDU from the peer, the interface sets the DCBX mode as IEEE DCBX. If the interface receives three DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.

Autonegotiation works slightly differently on QFX3500 switches and QFabric systems:

- QFX3500 switch—When an interface connects to its peer interface, the interface advertises IEEE DCBX TLVs to the peer. If the interface receives an IEEE DCBX TLV from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface receives three consecutive DCBX version 1.01 TLVs from the peer, the interface sets DCBX version 1.01 as the DCBX mode.

- QFabric system—When an interface connects to its peer interface, the interface advertises DCBX version 1.01 TLVs to the peer. If the interface receives an IEEE DCBX TLVs from the peer, the interface sets IEEE DCBX as the DCBX mode. If the interface receives three consecutive DCBX version 1.01 TLVs from the peer, the interface retains DCBX version 1.01 as the DCBX mode.



NOTE: If the link flaps or the LLDP process restarts, the interface starts the autonegotiation process again. The interface does not use the last received DCBX communication mode.

CNA Support for DCBX Modes

Different CNA vendors support different versions and capabilities of DCBX. The DCBX configuration you use on QFX Series interfaces depends on the DCBX features that the CNAs in your network support.

Interface Support for DCBX

You can configure DCBX on 10-Gigabit Ethernet interfaces and on link aggregation group (LAG) interfaces whose member interfaces are all 10-Gigabit Ethernet interfaces.

DCBX Attribute Types

DCBX has three attribute types:

- Informational—These attributes are exchanged using LLDP, but do not affect DCBX state or operation; they only communicate information to the peer. For example, application priority TLVs are informational TLVs.
- Asymmetric—The values for these types of attributes do not have to be the same on the connected peer interfaces. Peers exchange asymmetric attributes when the attribute values can differ on each peer interface. The peer interface configurations might match or they might differ. For example, ETS Configuration and Recommendation TLVs are asymmetric TLVs.
- Symmetric—The intention is that the values for these types of attributes should be the same on both of the connected peer interfaces. Peer interfaces exchange symmetric attributes to ensure symmetric DCBX configuration for those attributes. For example, PFC Configuration TLVs are symmetric TLVs.

The following sections describe asymmetric and symmetric DCBX attributes:

- [Asymmetric Attributes on page 174](#)
- [Symmetric Attributes on page 175](#)

Asymmetric Attributes

DCBX passes asymmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features). The resulting configuration for an attribute might be different on each peer, so the parameters configured on one interface might not match the parameters on the connected peer interface.

There are two types of asymmetric attribute TLVs:

- **Configuration TLV**—Configuration TLVs communicate the current operational state and the state of the “willing” bit. The “willing” bit communicates whether or not the interface is willing to accept and use the configuration from the peer interface. If an interface is “willing,” the interface uses the configuration it receives from the peer interface. (The peer interface configuration can override the configuration on the “willing” interface.) If an interface is “not willing,” the configuration on the interface cannot be overridden by the peer interface configuration.
- **Recommendation TLV**—Recommendation TLVs communicate the parameters the interface recommends that the connected peer interface should use. When an interface sends a Recommendation TLV, if the connected peer is “willing,” the connected peer changes its configuration to match the parameters in the Recommendation TLV.

Symmetric Attributes

DCBX passes symmetric attributes between connected peer interfaces to communicate parameter information about those attributes (features), with the objective that both interfaces should use the same configuration. The intent is that the parameters configured on one interface should match the parameters on the connected peer interface.

There is one type of symmetric attribute TLV, the Configuration TLV. As with asymmetric attributes, symmetric attribute Configuration TLVs communicate the current operational state and the state of the “willing” bit. “Willing” interfaces use the peer interface parameter values for the attribute. (The attribute configuration of the peer overrides the configuration on the “willing” interface.)

DCBX Application Protocol TLV Exchange

DCBX advertises the switch's capabilities for Layer 2 applications such as FCoE and Layer 4 applications such as iSCSI:

- [Application Protocol TLV Exchange on page 175](#)
- [FCoE Application Protocol TLV Exchange on page 175](#)
- [Disabling Application Protocol TLV Exchange on page 176](#)

Application Protocol TLV Exchange

For all applications, DCBX advertises the application's state and IEEE 802.1p code points on the interfaces to which the application is mapped. If an application is not mapped to an interface, that interface does not advertise the application's TLVs. There is an exception for FCoE application protocol TLV exchange when FCoE is the only application you want DCBX to advertise on an interface.

FCoE Application Protocol TLV Exchange

Protocol TLV exchange for the FCoE application depends on whether FCoE is the only application you want the interface to advertise or whether you want the interface to exchange other application TLVs in addition to FCoE TLVs.

If FCoE is the only application you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)
- Does *not* have an application map



NOTE: If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

If you want DCBX to advertise FCoE and other applications on an interface, you must specify all of the applications, including FCoE, in an application map, and apply the application map to the desired interfaces.



NOTE: If an application map is applied to an interface, the FCoE application must be explicitly configured in the application map, or the interface does not exchange FCoE TLVs.

When DCBX advertises the FCoE application, it advertises the FCoE state and IEEE 802.1p code points. If a peer device connected to a switch interface does not support FCoE, DCBX uses autonegotiation to mark the interface as “FCoE down,” and FCoE is disabled on that interface.

Disabling Application Protocol TLV Exchange

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

DCBX and PFC

After you enable PFC on a switch interface, DCBX uses autonegotiation to control the operational state of the PFC functionality.

If the peer device connected to the interface supports PFC and is provisioned compatibly with the switch, DCBX sets the PFC operational state to enabled. If the peer device connected to the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled. (PFC must be symmetrical.)

If the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state.

You can manually override DCBX control of the PFC operational state on a per-interface basis by disabling autonegotiation. If you disable autonegotiation on an interface on which you have configured PFC, then PFC is enabled on that interface regardless of the peer configuration. To disable PFC on an interface, do not configure PFC on that interface.

DCBX and ETS

This section describes:

- [Default DCBX ETS Advertisement on page 177](#)
- [ETS Advertisement and Peer Configuration on page 177](#)
- [ETS Recommendation TLV on page 178](#)

Default DCBX ETS Advertisement

If you do not configure ETS on an interface, the switch automatically creates a default priority group that contains all of the priorities (forwarding classes, which represent output queues) and assigns 100 percent of the port output bandwidth to that priority group. The default priority group is transparent. It does not appear in the configuration and is used for DCBX advertisement. DCBX advertises the default priority group, its priorities, and the assigned bandwidth.

If you configure ETS on an interface, DCBX advertises:

- Each priority group on the interface
- The priorities in each priority group
- The bandwidth properties of each priority group and priority

Any priority on that interface that is not part of an explicitly configured priority group (forwarding class set) is assigned to the automatically generated default priority group and receives no bandwidth. If you configure ETS on an interface, every forwarding class (priority) on that interface for which you want to forward traffic must belong to a forwarding class set (priority group).

ETS Advertisement and Peer Configuration

DCBX does not control the switch’s ETS (hierarchical scheduling) operational state. If the connected peer is configured as “willing,” DCBX pushes the switch’s ETS configuration to the switch’s peers if the ETS Recommendation TLV is enabled (it is enabled by default). If the peer does not support ETS or is not consistently provisioned with the switch, DCBX does not change the ETS operational state on the switch. The ETS operational state remains enabled or disabled based only on the switch hierarchical scheduling configuration and is enabled by default.

When ETS is configured, DCBX advertises the priority groups, the priorities in the priority groups, and the bandwidth configuration for the priority groups and priorities. Any priority

(essentially a forwarding class or queue) that is not part of a priority group has no scheduling properties and receives no bandwidth.

You can manually override whether DCBX advertises the ETS state to the peer on a per-interface basis by disabling autonegotiation. This does not affect the ETS state on the switch or on the peer, but it does prevent the switch from sending the Recommendation TLV or the Configuration TLV to the connected peer. To disable ETS on an interface, do not configure priority groups (forwarding class sets) on the interface.

ETS Recommendation TLV

The ETS Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is “willing,” it changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV by including the **no-recommendation-tlv** statement at the **[edit protocols dcbx interface *interface-name* enhanced-transmission-selection]** hierarchy level.



NOTE: You can disable the ETS Recommendation TLV only when the DCBX mode on the interface is IEEE DCBX. Disabling the ETS Recommendation TLV has no effect if the DCBX mode on the interface is DCBX version 1.01. (IEEE DCBX uses separate application attribute TLVs, but DCBX version 1.01 sends all application attributes in the same TLV and uses sub-TLVs to separate the information.)

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

For example, if you want a CNA connected to a switch interface to have different bandwidth allocations than the switch ETS configuration, you can disable the ETS Recommendation TLV and configure the CNA for the desired bandwidth. The switch interface and the CNA exchange configuration parameters, but the CNA does not change its configuration to match the switch interface configuration.

Related Documentation

- [Understanding DCBX Application Protocol TLV Exchange on page 179](#)
- [Understanding DCB Features and Requirements on page 156](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)

- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Understanding FCoE](#)
- [Configuring the DCBX Mode on page 430](#)
- [Configuring DCBX Autonegotiation on page 431](#)
- [Disabling the ETS Recommendation TLV on page 434](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)

Understanding DCBX Application Protocol TLV Exchange

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers. DCBX also advertises the capabilities of applications on interfaces by exchanging application protocol information through application type, length, and value (TLV) elements. DCBX is an extension of Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



NOTE: LLDP and DCBX are enabled by default on all interfaces.

Setting up application protocol exchange consists of:

- Defining applications
- Mapping the applications to IEEE 802.1p code points in an *application map*
- Configuring classifiers to prioritize incoming traffic and map the incoming traffic to the application by the traffic code points
- Applying the application maps and classifiers to interfaces

You need to explicitly define the applications that you want an interface to advertise. The FCoE application is a special case (see [“Applications” on page 180](#)) and only needs to be defined on an interface if you want DCBX to exchange application protocol TLVs for other applications in addition to FCoE on that interface.

You also need to explicitly map all defined applications that you want an interface to advertise to IEEE 802.1p code points in an application map. The FCoE application is a special case (see [“Application Maps” on page 180](#)) and only requires inclusion in an application map when you want an interface to use DCBX for other applications in addition to FCoE, as described later in this topic.

This topic describes:

- [Applications on page 180](#)
- [Application Maps on page 180](#)
- [Classifying and Prioritizing Application Traffic on page 181](#)

- [Enabling Interfaces to Exchange Application Protocol Information on page 182](#)
- [Disabling DCBX Application Protocol Exchange on page 182](#)

Applications

Before an interface can exchange application protocol information, you need to define the applications that you want to advertise, except FCoE if FCoE is the only application that you want the interface to advertise.



NOTE: If FCoE is the only application that you want DCBX to advertise on an interface, DCBX exchanges FCoE application protocol TLVs by default if the interface:

- Carries FCoE traffic (traffic mapped by CoS configuration to the FCoE forwarding class and applied to the interface)
- Has a congestion notification profile with PFC enabled on the FCoE priority (IEEE 802.1p code point)
- Does *not* have an application map

If you apply an application map to an interface, then all applications that you want DCBX to advertise must be defined and configured in the application map, including the FCoE application.

If no CoS configuration for FCoE is mapped to an interface, that interface does not exchange FCoE application protocol TLVs.

You can define:

- Layer 2 applications by EtherType
- Layer 4 applications by a combination of protocol (TCP or UDP) and destination port number

The EtherType is a two-octet field in the Ethernet frame that denotes the protocol encapsulated in the frame. For a list of common EtherTypes, see <http://standards.ieee.org/develop/regauth/ethertype/eth.txt> on the IEEE standards organization website. For a list of port numbers and protocols, see the *Service Name and Transport Protocol Port Number Registry* at <http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml> on the Internet Assigned Numbers Authority (IANA) website.

You must explicitly define each application that you want to advertise, except FCoE. The FCoE application is defined by default (EtherType 0x8906).

Application Maps

An application map maps defined applications to one or more IEEE 802.1p code points. Each application map contains one or more applications. DCBX includes the configured application code points in the protocol TLVs exchanged with the connected peer.

To exchange protocol TLVs for an application, you must include the application in an application map. The FCoE application is a special case:

- If you want DCBX to exchange application protocol TLVs for more than one application on a particular interface, you must configure the applications, define an application map to map the applications to code points, and apply the application map to the interface. In this case, you must also define the FCoE application and add it to the application map.

This is the same process and treatment required for all other applications. In addition, for DCBX to exchange FCoE application TLVs, you must enable priority-based flow control (PFC) on the FCoE priority (the FCoE IEEE 802.1p code point) on the interface.

- If FCoE is the only application that you want DCBX to advertise on an interface, then you do not need to configure an application map and apply it to the interface. By default, when an interface has no application map, and the interface carries traffic mapped to the FCoE forwarding class, and PFC is enabled on the FCoE priority, the interface advertises FCoE TLVs (autonegotiation mode). DCBX exchanges FCoE application protocol TLVs by default until you apply an application map to the interface, remove the FCoE traffic from the interface (you can do this by removing the or editing the classifier for FCoE traffic), or disable PFC on the FCoE priority.

If you apply an application map to an interface that did not have an application map and was exchanging FCoE application TLVs, and you do not include the FCoE application in the application map, the interface stops exchanging FCoE TLVs. Every interface that has an application map must have FCoE included in the application map (and PFC enabled on the FCoE priority) in order for DCBX to exchange FCoE TLVs.

Mapping an application to code points does two things:

- Maps incoming traffic with the same code points to that application
- Allows you to configure classifiers that map incoming application traffic, by code point, to a forwarding class and a loss priority, in order to apply class of service (CoS) to application traffic and prioritize application traffic

You apply an application map to an interface to enable DCBX application protocol exchange on that interface for each application specified in the application map. All of the applications that you want an interface to advertise must be configured in the application map that you apply to the interface, with the previously noted exception for the FCoE application when FCoE is the only application for which you want DCBX to exchange protocol TLVs on an interface.

Classifying and Prioritizing Application Traffic

When traffic arrives at an interface, the interface classifies the incoming traffic based on its code points. Classifiers map code points to loss priorities and forwarding classes. The loss priority prioritizes the traffic. The forwarding class determines the traffic output queue and CoS service level.

When you map an application to an IEEE 802.1p code point in an application map and apply the application map to an interface, incoming traffic on the interface that matches the application code points is mapped to the appropriate application. The application

receives the loss priority and the CoS associated with the forwarding class for those code points, and is placed in the output queue associated with the forwarding class.

You can use the default classifier or you can configure a classifier to map the application code points defined in the application map to forwarding classes and loss priorities.

Enabling Interfaces to Exchange Application Protocol Information

Each interface with the **fcoe** forwarding class and PFC enabled on the FCoE code point is enabled for FCoE application protocol exchange by default until you apply an application map to the interface. If you apply an application map to an interface and you want that interface to exchange FCoE application protocol TLVs, you must include the FCoE application in the application map. (In all cases, to achieve lossless transport, you must also enable PFC on the FCoE code point or code points.)

Except when FCoE is the only protocol you want DCBX to advertise on an interface, interfaces on which you want to exchange application protocol TLVs must include the following two items:

- The application map that contains the application(s)
- A classifier



NOTE: You must also enable PFC on the code point of any traffic for which you want to achieve lossless transport.

Disabling DCBX Application Protocol Exchange

To disable DCBX application protocol exchange for all applications on an interface, issue the **set protocols dcbx interface *interface-name* applications no-auto-negotiation** command.

You can also disable DCBX application protocol exchange for applications on an interface by deleting the application map from the interface, or by deleting a particular application from the application map. However, when you delete an application from an application map, the application protocol is no longer exchanged on any interface which uses that application map.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable sending the enhanced transmission selection (ETS) Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers.

Related Documentation

- [Understanding DCBX on page 170](#)
- [Understanding CoS Classifiers on page 51](#)
- [Configuring DCBX Autonegotiation on page 431](#)
- [Disabling the ETS Recommendation TLV on page 434](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 434](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 436](#)

- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 437](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)
- [Example: Configuring Unicast Classifiers on page 273](#)

CHAPTER 3

QFabric-Specific CoS Overview

- [Understanding CoS Fabric Forwarding Class Sets on page 186](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 199](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 203](#)

Understanding CoS Fabric Forwarding Class Sets

Fabric forwarding class sets (fabric fc-sets) are similar to the fc-sets (priority groups) you configure on Node devices. The major differences are:

1. Fabric fc-sets group traffic for transport across the QFX3008-I or QFX3600-I Interconnect device (the fabric). Node device fc-sets group traffic on a Node device for transport across that Node device.
2. Fabric fc-sets are global. They apply to the entire fabric. Node device fc-sets apply only to the Node device on which they are configured.
3. You can configure class of service (CoS) scheduling for Node device fc-sets, but you cannot configure CoS for fabric fc-sets.
4. Fabric fc-sets map to Interconnect device fabric output queues statically—you cannot configure the mapping of fabric fc-sets to fabric output queues. All traffic in a fabric fc-set maps to the same output queue.

Node device fc-sets include forwarding classes that map to Node device output queues, and you can configure the mapping of forwarding classes to output queues (or you can use the default mapping). Because output queues are mapped to forwarding classes, different classes of traffic in a Node device fc-set can be mapped to different output queues.

Node device fc-sets consist of forwarding classes containing traffic that requires similar CoS treatment. (Forwarding classes are default forwarding classes or user-defined forwarding classes.) You can configure CoS for each fc-set to determine how the traffic of its forwarding classes is scheduled on a Node device.

When traffic exits a Node device interface and enters an Interconnect device fabric interface, the Interconnect device uses the same forwarding classes to group traffic. The forwarding classes are mapped to global fabric fc-sets for transport across the fabric. Like fc-sets on a Node device, fabric fc-sets also contain traffic that requires similar CoS treatment. Unlike fc-sets on a Node device, you cannot configure CoS on fabric fc-sets.

Fabric fc-sets reside on the Interconnect device and are global to the QFabric system. Fabric fc-sets apply to all traffic that traverses the fabric. The mapping of forwarding classes to fabric fc-sets is global and applies to all forwarding classes with traffic that traverses the fabric from all connected Node devices. You can change the mapping of forwarding classes to fabric fc-sets. All mapping changes you make are global. For example, if you change the fabric fc-set to forwarding class mapping of the default best-effort forwarding class, then every Node device's best-effort forwarding class traffic that traverses the fabric is mapped to that fabric fc-set.

This topic describes:

- [Default Fabric Forwarding Class Sets on page 187](#)
- [Fabric Forwarding Class Set Configuration and Implementation on page 190](#)
- [Fabric Forwarding Class Set Scheduling \(CoS\) on page 192](#)
- [Support for Flow Control and Lossless Transport Across the Fabric on page 194](#)

- [Viewing Fabric Forwarding Class Set Information on page 196](#)
- [Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences on page 198](#)

Default Fabric Forwarding Class Sets

Interconnect devices have 12 default fabric fc-sets, including five visible default fabric fc-sets, four for unicast traffic and one for multdestination (multicast, broadcast, and destination lookup failure) traffic.

There are also seven hidden default fabric fc-sets. There are three hidden default fabric fc-sets for multdestination traffic that you can use if you want to map different multdestination forwarding classes to different multdestination fabric fc-sets. There are four hidden default fabric fc-sets for lossless traffic that you can use to map different lossless forwarding classes (priorities) to different lossless fabric fc-sets.

[Table 74 on page 187](#) shows the default fabric fc-sets:

Table 74: Default Fabric Forwarding Class Sets

Fabric Forwarding Class Set Name	Characteristics
fabric_fcset_be	Transports best-effort unicast traffic across the fabric.
fabric_fcset_strict_high	Transports unicast traffic that has been configured with strict-high priority and in the network-control forwarding class across the fabric. This fabric fc-set receives as much bandwidth across the fabric as it needs to service the traffic in the group up to the entire fabric interface bandwidth. For this reason, exercise caution when mapping traffic to this fabric fc-set to avoid starving other traffic.
fabric_fcset_noloss1	Transports unicast traffic in the default fc0e forwarding class across the fabric.
fabric_fcset_noloss2	Transports unicast traffic in the default no-loss forwarding class across the fabric.
fabric_fcset_noloss3	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.
fabric_fcset_noloss4	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.
fabric_fcset_noloss5	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.

Table 74: Default Fabric Forwarding Class Sets (*continued*)

Fabric Forwarding Class Set Name	Characteristics
fabric_fcset_noloss6	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for lossless forwarding classes.
fabric_fcset_multicast1	Transports multdestination traffic in the mcast forwarding class across the fabric. This fabric fc-set is valid only for multdestination forwarding classes.
fabric_fcset_multicast2	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes.
fabric_fcset_multicast3	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes.
fabric_fcset_multicast4	(Hidden) No traffic is assigned by default to this fabric fc-set. Unless traffic is mapped to this fabric fc-set, this fabric fc-set remains hidden. This fabric fc-set is valid only for multdestination forwarding classes.

The five default forwarding classes (**best-effort**, **fcoe**, **no-loss**, **network-control**, and **mcast**) are mapped to the fabric fc-sets by default as shown in [Table 75 on page 188](#).

Table 75: Default Forwarding Class to Fabric Forwarding Class Set Mapping

Forwarding Class	Fabric Forwarding Class Set	Fabric Output Queue	Maximum MTU Supported for Lossless Operation
best-effort	fabric_fcset_be	0	NA
network-control	fabric_fcset_strict_high	7	NA
fcoe	fabric_fcset_noloss1	1	9K
no-loss	fabric_fcset_noloss2	2	9K
mcast	fabric_fcset_multicast1	8	NA
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss3	3	9k

Table 75: Default Forwarding Class to Fabric Forwarding Class Set Mapping (*continued*)

Forwarding Class	Fabric Forwarding Class Set	Fabric Output Queue	Maximum MTU Supported for Lossless Operation
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss4	4	9k
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss5	5	9k
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_noloss6	6	9k
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_multicast2	9	NA
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_multicast3	10	NA
No forwarding classes are mapped by default to this hidden fabric fc-set.	fabric_fcset_multicast4	11	NA

The maximum fiber cable length between the QFabric system Node device and the QFabric system Interconnect device is 150 meters.



TIP: If you explicitly configure lossless forwarding classes, we recommend that you map each user-configured lossless forwarding class to an unused fabric fc-set (fabric_fcset_noloss3 through fabric_fcset_noloss6) on a one-to-one basis: one lossless forwarding class mapped to one lossless fabric fc-set.

The reason for one-to-one mapping is to avoid fate sharing of lossless flows. Because each fabric fc-set is mapped statically to an output queue, when you map more than one forwarding class to a fabric fc-set, all of the traffic in all of the forwarding classes that belong to the fabric fc-set uses the same output queue. If that output queue becomes congested due to congestion caused by one of the flows, the other flows are also affected. (They share fate because the flow that congests the output queue affects flows that are not experiencing congestion.)

However, it is important to understand that fabric_fcset_noloss1 and fabric_fcset_noloss2 have a scheduling weight of 35, while the other fabric fc-sets have a scheduling weight of 1. The scheduling weights mean that fabric_fcset_noloss1 and fabric_fcset_noloss2 receive most of the bandwidth available to lossless fabric fc-sets if the amount of traffic on fabric_fcset_noloss1 and fabric_fcset_noloss2 requires the bandwidth.

If you believe that the traffic on `fabric_fcset_noloss1` and `fabric_fcset_noloss2` will consume most of that bandwidth, then you should place all lossless traffic on `fabric_fcset_noloss1` and `fabric_fcset_noloss2`. If you believe that the traffic on `fabric_fcset_noloss1` and `fabric_fcset_noloss2` will **<emphasis>not</emphasis>** consume most of that bandwidth, then you can map lossless forwarding classes in a one-to-one manner to lossless fabric fc-sets to avoid fate sharing.

If you want to map different multidestination forwarding classes to different multidestination fabric fc-sets, use one or more of the hidden multidestination fabric fc-sets.



NOTE: The global mapping of forwarding classes to fabric fc-sets is independent of the mapping of forwarding classes to Node device fc-sets. Global mapping of forwarding classes to fabric fc-sets occurs only on the Interconnect device. The Node device mapping of forwarding classes to fc-sets does not affect the global mapping of forwarding classes to fabric fc-sets on the Interconnect device, and vice versa.

When you define new forwarding classes on a Node device, you explicitly map those forwarding classes to Node device fc-sets. However, new (user-created) forwarding classes are mapped by default to fabric fc-sets. (You can override the default mapping if you want to configure the forwarding class to fabric fc-set mapping explicitly, as described in the next section.)

By default:

- All best-effort traffic forwarding classes that you create are mapped to the **`fabric_fcset_be`** fabric fc-set.
- All lossless traffic forwarding classes that you create are mapped to the **`fabric_fcset_noloss1`** or **`fabric_fcset_noloss2`** fabric fc-set.
- All multidestination traffic forwarding classes that you create are mapped to the **`fabric_fcset_multicast1`** fabric fc-set.
- All **strict-high** priority traffic and **network-control** forwarding classes that you create are mapped to the **`fabric_fcset_strict_high`** fabric fc-set.

Fabric Forwarding Class Set Configuration and Implementation

You can map forwarding classes to fabric fc-sets, but no other attributes of fabric fc-sets are user-configurable, including CoS. This section describes:

- [Mapping Forwarding Classes to Fabric Forwarding Class Sets on page 191](#)
- [Fabric Forwarding Class Set Implementation on page 191](#)

Mapping Forwarding Classes to Fabric Forwarding Class Sets

If you do not want to use the default mapping of forwarding classes to fabric fc-sets, you can map forwarding classes to fabric fc-sets the same way as you map forwarding classes to Node device fc-sets. To do this, use exactly the same statement that you use to map forwarding classes to fc-sets, but instead of specifying a Node device fc-set name, specify a fabric fc-set name.



NOTE: The global mapping of forwarding classes to fabric fc-sets does not affect the mapping of forwarding classes to Node device fc-sets. The global forwarding class mapping to fabric fc-sets pertains to the traffic only when it enters, traverses, and exits the fabric. The forwarding class mapping to fc-sets on a Node device is valid within that Node device.

Mapping forwarding classes to fabric fc-sets does not affect the scheduling configuration of the forwarding classes or fc-sets on Node devices. Fabric fc-set scheduling (which is not user-configurable) pertains to traffic only when it enters, traverses, and exits the Interconnect device fabric.

If you change the mapping of a forwarding class to a fabric fc-set, the new mapping is global and applies to all traffic in that forwarding class, regardless of which Node device forwards the traffic to the Interconnect device.

- To assign one or more forwarding classes to a fabric fc-set:

```
[edit class-of-service]
```

```
user@switch# set forwarding-class-sets fabric-forwarding-class-set-name class  
forwarding-class-name
```

For example, to map a user-defined forwarding class named **best-effort-2** to the fabric fc-set **fabric_fcset_be**:

```
[edit class-of-service]
```

```
user@switch# set forwarding-class-sets fabric_fcset_be class best-effort-2
```



NOTE: Because fabric fc-set configuration is global, in this example all forwarding classes with the name **best-effort-2** on all of the Node devices attached to the fabric use the **fabric_fcset_be** fabric fc-set to transport traffic across the fabric.

Fabric Forwarding Class Set Implementation

The following rules apply to fabric fc-sets:

- You cannot create new fabric fc-sets. Only the twelve default fabric fc-sets are available.
- You cannot delete a default fabric fc-set.
- You cannot attach a fabric fc-set to a Node device interface. Fabric fc-sets are used only on the Interconnect device fabric, not on Node devices.

- You can map only multdestination forwarding classes to multdestination fabric fc-sets.
- You cannot map multdestination forwarding classes to unicast fabric fc-sets.
- You cannot map unicast forwarding classes to multdestination fabric fc-sets.
- You cannot configure CoS for fabric fc-sets. (However, default CoS scheduling properties are applied to traffic on the fabric, and the fabric interfaces use link layer flow control (LLFC) for flow control.)

Fabric Forwarding Class Set Scheduling (CoS)

Although fabric fc-set CoS is not user-configurable, CoS is applied to traffic on the fabric. (In addition, fabric interfaces use LLFC to ensure lossless transport for lossless traffic flows.) This section describes how the fabric applies CoS scheduling to traffic:

- [Class Groups for Fabric Forwarding Class Sets on page 192](#)
- [Class Group Scheduling on page 192](#)
- [QFabric System CoS on page 194](#)

Class Groups for Fabric Forwarding Class Sets

To transport traffic across the fabric, the QFabric system organizes the fabric fc-sets into three classes called *class groups*. The three class groups are:

- **Strict-high priority**—All traffic in the fabric fc-set **fabric_fcset_strict_high**. This class group includes the traffic in **strict-high** priority and **network-control** forwarding classes and in any forwarding classes you create on a Node device that consist of **strict-high** priority or **network-control** forwarding class traffic.
- **Unicast**—All traffic in the fabric fc-sets **fabric_fcset_be**, **fabric_fcset_noloss1**, and **fabric_fcset_noloss2**. This class group includes the traffic in the **best-effort**, **fcoe**, and **no-loss** forwarding classes and in any forwarding classes you create on a Node device that consist of best-effort or lossless traffic. If you use any of the hidden no loss fabric fc-sets (**fabric_fcset_noloss3**, **fabric_fcset_noloss4**, **fabric_fcset_noloss5**, or **fabric_fcset_noloss6**), that traffic is part of this class group.
- **Multidestination**—All traffic in the fabric fc-set **fabric_fcset_multicast1**. This class group includes the traffic in the **mcast** forwarding class and in any forwarding classes you create on a Node device that consist of multidestination traffic. If you use any of the hidden multidestination fabric fc-sets (**fabric_fcset_multicast2**, **fabric_fcset_multicast3**, or **fabric_fcset_multicast4**), that traffic is also classified as part of this class group.

Class Group Scheduling

You cannot configure CoS for class groups or for fabric fc-sets (that is, you cannot attach a traffic control profile to a fabric fc-set—you attach traffic control profiles to Node device fc-sets to apply scheduling to the traffic that belongs to the Node device fc-set). By default, the fabric uses weighted round-robin (WRR) scheduling in which each class group receives a portion of the total available fabric bandwidth based on its type of traffic, as shown in [Table 76 on page 193](#):

Table 76: Class Group Scheduling Properties and Membership

Class Group	Fabric fc-sets	Forwarding Classes (Default Mapping)	Class Group Scheduling Properties (Weight)
Strict-high priority	fabric_fcset_strict_high	<ul style="list-style-type: none"> All strict-high priority forwarding classes network-control 	Traffic in the strict-high priority class group is served first. This class group receives all of the bandwidth it needs to empty its queues and therefore can starve other types of traffic during periods of high-volume strict-high priority traffic. Plan carefully and use caution when determining how much traffic to configure as strict-high priority traffic.
Unicast	<ul style="list-style-type: none"> fabric_fcset_be fabric_fcset_noloss1 fabric_fcset_noloss2 <p>Includes the hidden lossless fabric fc-sets if used:</p> <ul style="list-style-type: none"> fabric_fcset_noloss3 fabric_fcset_noloss4 fabric_fcset_noloss5 fabric_fcset_noloss6 	<ul style="list-style-type: none"> best-effort fcoe no-loss <p>NOTE: No forwarding classes are mapped to the hidden lossless fabric_fcsets by default.</p>	Traffic in the unicast class group receives an 80% weight in the WRR calculations. After the strict-high priority class group has been served, the unicast class group receives 80% of the remaining fabric bandwidth. (If more bandwidth is available, the unicast class group can use more bandwidth.)
Multidestination	fabric_fcset_multicast1 <p>Includes the hidden multidestination fabric fc-sets if used:</p> <ul style="list-style-type: none"> fabric_fcset_multicast2 fabric_fcset_multicast3 fabric_fcset_multicast4 	<ul style="list-style-type: none"> mcast <p>NOTE: No forwarding classes are mapped to the hidden multidestination fabric_fcsets by default.</p>	Traffic in the multidestination class group receives a 20% weight in the WRR calculations. After the strict-high priority class group has been served, the multidestination class group receives 20% of the remaining fabric bandwidth. (If more bandwidth is available, the multidestination class group can use more bandwidth.)

The fabric fc-sets within each class group are weighted equally and receive bandwidth using round-robin scheduling. For example:

- If the unicast class group has three member fabric fc-sets, **fabric_fcset_be**, **fabric_fcset_noloss1**, and **fabric_fcset_noloss2**, then each of the three fabric fc-sets receives one-third of the bandwidth available to the unicast class group.
- If the multidestination class group has one member fc-set, **fabric_fcset_multicast1**, then that fc-set receives all of the bandwidth available to the multidestination class group.
- If the multidestination class group has two member fc-sets, **fabric_fcset_multicast1** and **fabric_fcset_multicast2**, then each of the two fabric fc-sets receives one-half of the bandwidth available to the multidestination class group.

QFabric System CoS

When traffic enters and exits the same Node device, CoS works the same as it works on a standalone QFX3500 switch.

However, when traffic enters a Node device, crosses the Interconnect device, and then exits a different Node device, CoS is applied differently:

1. Traffic entering the ingress Node device receives the CoS configured at the Node ingress (packet classification, congestion notification profile for PFC).
2. When traffic goes from the ingress Node device to the Interconnect device, the fabric fc-set CoS is applied as described in the discussion of fabric forwarding class set scheduling.
3. When traffic goes from the Interconnect device to the egress Node device, the egress Node device applies CoS at the egress port (egress queue scheduling, WRED, IEEE 802.1p or DSCP code-point rewrite).

Support for Flow Control and Lossless Transport Across the Fabric

The Interconnect device incorporates flow control mechanisms to support lossless transport during periods of congestion on the fabric. To support the priority-based flow control (PFC) feature on the Node devices, the fabric interfaces use LLFC to support lossless transport for up to six IEEE 802.1p priorities when the following two configuration constraints are met:

1. The IEEE 802.1p priority used for the traffic that requires lossless transport is mapped to a lossless forwarding class on the Node devices.
2. The lossless forwarding class must be mapped to a lossless fabric fc-set on the Interconnect device (**fabric_fcset_noloss1**, **fabric_fcset_noloss2**, **fabric_fcset_noloss3**, **fabric_fcset_noloss4**, **fabric_fcset_noloss5**, or **fabric_fcset_noloss6**).

When traffic meets the two configuration constraints, the fabric propagates the back pressure from the egress Node device across the fabric to the ingress Node device during periods of congestion. However, to achieve end-to-end lossless transport across the switch, you must also configure a congestion notification profile to enable PFC on the Node device ingress ports.

For all other combinations of IEEE 802.1p priority to forwarding class mapping and all other combinations of forwarding class to fabric fc-set mapping, the congestion control mechanism is normal packet drop. For example:

- **Case 1**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric_fcset_noloss1** fabric fc-set, then the congestion control mechanism is PFC.
- **Case 2**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.

- **Case 3**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric_fcset_noloss2** fabric fc-set, then the congestion control mechanism is PFC.
- **Case 4**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.
- **Case 5**—If the IEEE 802.1p priority 5 is mapped to the **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.
- **Case 6**—If the IEEE 802.1p priority 5 is mapped to the **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric_fcset_noloss1** fabric fc-set, then the congestion control mechanism is packet drop.



NOTE: Lossless transport across the fabric also must meet the following two conditions:

1. The maximum cable length between the Node device and the Interconnect device is a 150 meters of fiber cable.
2. The maximum frame size is 9216 bytes.

If the MTU is 9216 KB, in some cases the QFabric system supports only five lossless forwarding classes instead of six lossless forwarding classes because of headroom buffer limitations.

The number of IEEE 802.1p priorities (forwarding classes) the QFabric system can support for lossless transport across the Interconnect device fabric depends on several factors:

- **Approximate fiber cable length**—The longer the fiber cable that connects Node device fabric (FTE) ports to the Interconnect device fabric ports, the more data the connected ports need to buffer when a pause is asserted. (The longer the fiber cable, the more frames are traversing the cable when a pause is asserted. Each port must be able to store all of the “in transit” frames in the buffer to preserve lossless behavior and avoid dropping frames.)
- **MTU size**—The larger the maximum frame sizes the buffer must hold, the fewer frames the buffer can hold. The larger the MTU size, the more buffer space each frame consumes.
- **Total number of Node device fabric ports connected to the Interconnect device**—The higher the number of connected fabric ports, the more headroom buffer space the Node device needs on those fabric ports to support the lossless flows that traverse the Interconnect device. Because more buffer space is used on the Node device fabric ports, less buffer space is available for the Node device access ports, and a lower total number of lossless flows are supported.

The QFabric system supports six lossless priorities (forwarding classes) under most conditions. The priority group headroom that remains after allocating headroom to lossless flows is sufficient to support best-effort and multideestination traffic.

Table 77 on page 196 shows how many lossless priorities the QFabric system supports under different conditions (fiber cable lengths and MTUs) in cases when the QFabric system supports fewer than six lossless priorities. The number of lossless priorities is the same regardless of how many Node device FTE ports are connected to the Interconnect device. However, the higher the number of FTE ports connected to the Interconnect device, the lower the number of total lossless flows supported. In all cases that are not shown in Table 77 on page 196, the QFabric system supports six lossless priorities.



NOTE: The system does not perform a configuration commit check that compares available system resources with the number of lossless forwarding classes configured. If you commit a configuration with more lossless forwarding classes than the system resources can support, frames in lossless forwarding classes might be dropped.

Table 77: Lossless Priority (Forwarding Class) Support for QFX3500 and QFX3600 Node Devices When Fewer than Six Lossless Priorities Are Supported

MTU in Bytes	Fiber Cable Length in Meters (Approximate)	Maximum Number of Lossless Priorities (Forwarding Classes) on the Node Device
9216 (9K)	100	5
9216 (9K)	150	5



NOTE: The total number of lossless flows decreases as resource consumption increases. For a Node device, the higher the number of FTE ports connected to the Interconnect device, the larger the MTU, and the longer the fiber cable length, the fewer total lossless flows the QFabric system can support.

Viewing Fabric Forwarding Class Set Information

You can display information about fabric fc-sets using the same CLI command you use to display information about Node device fc-sets:

```
user@switch> show class-of-service forwarding-class-set
Forwarding class set: fabric_fcset_be, Type: fabric-type, Forwarding class set
index: 1
  Forwarding class      Index
  best-effort           0

Forwarding class set: fabric_fcset_mcast1, Type: fabric-type, Forwarding class
set index: 5
  Forwarding class      Index
```



```

mcast                                     8

Forwarding class set: fabric_fcset_mcast2, Type: fabric-type, Forwarding class
set index: 6

Forwarding class set: fabric_fcset_mcast3, Type: fabric-type, Forwarding class
set index: 7

Forwarding class set: fabric_fcset_mcast4, Type: fabric-type, Forwarding class
set index: 8

Forwarding class set: fabric_fcset_noloss1, Type: fabric-type, Forwarding class
set index: 2
  Forwarding class                        Index
  fcoe                                    1

Forwarding class set: fabric_fcset_noloss2, Type: fabric-type, Forwarding class
set index: 3
  Forwarding class                        Index
  no-loss                                2

Forwarding class set: fabric_fcset_noloss3, Type: fabric-type, Forwarding class
set index: 9

Forwarding class set: fabric_fcset_noloss4, Type: fabric-type, Forwarding class
set index: 10

Forwarding class set: fabric_fcset_noloss5, Type: fabric-type, Forwarding class
set index: 11

Forwarding class set: fabric_fcset_noloss6, Type: fabric-type, Forwarding class
set index: 12

Forwarding class set: fabric_fcset_strict_high, Type: fabric-type, Forwarding
class set index: 4
  Forwarding class                        Index
  network-control                         3

```

[Table 78 on page 197](#) describes the meaning of the **show class-of-service forwarding-class-set** output fields when you display fabric fc-set information.

Table 78: show class-of-service forwarding-class-set Command Output Fields

Field Name	Field Description
Forwarding class set	Name of the fabric forwarding class set.
Type	Type of forwarding class set: <ul style="list-style-type: none"> Fabric-type—Fabric fc-set Normal-type—Node device fc-set
Forwarding class set index	Index of this forwarding class set.
Forwarding class	Name of a forwarding class.
Index	Index of the forwarding class.

Summary of Fabric Forwarding Class Set and Node Device Forwarding Class Set Differences

Table 79 on page 198 summarizes the differences between fabric fc-sets and Node device fc-sets:

Table 79: Summary of Differences Between Fabric fc-sets and Node Device fc-sets

Characteristic	Fabric fc-set	Node device fc-set
Location	QFX3008-I or QFX3600-I Interconnect device (the fabric).	QFabric system Node device.
Global or Node-device specific	Global, valid for the entire fabric.	Local to the Node device on which the fc-set is configured.
Ability to create (define) a new fc-set	No. Use the 12 default fabric fc-sets provided.	Yes.
Ability to configure CoS	Default CoS settings only. CoS is not user-configurable.	User-configurable using traffic control profiles.
Ability to map forwarding classes to an fc-set	Yes. Mapping is global and applies to all forwarding classes across Interconnect device fabric (traffic from all connected Node devices).	Yes. Mapping is local to a Node device and applies only to the forwarding classes on the Node device.

Related Documentation

- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 73](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 199](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 203](#)
- [Defining CoS Forwarding Class Sets on page 410](#)
- [Example: Configuring Forwarding Class Sets on page 286](#)
- [show class-of-service forwarding-class-set on page 570](#)

Understanding CoS Scheduling on QFabric System Node Device Fabric (fte) Ports

Beginning with Junos OS Release 13.1, you can configure two-tier hierarchical scheduling (enhanced transmission selection, IEEE 802.1Qaz) on the fabric (fte) ports of QFabric system Node devices. Configuring CoS on Node device fabric interfaces provides increased control over traffic scheduling and helps to ensure predictable bandwidth consumption.

You can configure CoS on the following QFabric system interface types:

- Node device access interfaces (xe interfaces)—Schedule traffic on the output queues of the 10-Gigabit Ethernet access ports using standard Node device CoS scheduling configuration components, as described elsewhere in the QFX Series documentation. You can configure different scheduling for different ports and output queues.
- Node device fabric interfaces (fte interfaces)—Schedule traffic on the output queues of the 40-Gbps fabric interfaces that connect a Node device to a QFX3008-I or a QFX3600-I Interconnect device using standard Node device CoS scheduling configuration components. You can configure different scheduling for different interfaces and output queues.

This topic describes:

- [Hierarchical Scheduling Architecture on QFabric System Node Devices on page 199](#)
- [Default Scheduling on Node Device Fabric Interfaces on page 200](#)
- [Configuring Scheduling on Node Device Fabric Interfaces on page 201](#)

Hierarchical Scheduling Architecture on QFabric System Node Devices

CoS architecture on Node device access interfaces is the same as CoS architecture on standalone switch access interfaces. CoS architecture on Node device fabric interfaces is also the same as the CoS architecture on the access interfaces. You apply schedulers to queues (priorities), fc-sets (priority groups), and interfaces in the same hierarchical manner as described in [“Understanding CoS Hierarchical Port Scheduling \(ETS\)” on page 81](#).

You configure scheduling on Node device fabric interfaces (fte interfaces) using the same statements and configuration constructs that you use to configure scheduling on Node device access interfaces (xe interfaces). For example, on Node device fabric interfaces you can:

- Define up to four fc-sets (three unicast, one multdestination)



NOTE: If the fabric interface handles strict-high priority traffic, you must define a separate fc-set (priority group) for strict-high priority traffic. Strict-high priority traffic cannot be mixed with traffic of other priorities in an fc-set. For example, you might choose to create different fc-sets for best-effort, lossless, strict-high priority, and multdestination traffic.

- Map forwarding classes to fc-sets

- Configure scheduling for each forwarding class (scheduler)
- Configure scheduling for each fc-set (traffic control profile)

The differences in configuring CoS on Node device fabric interfaces compared to configuring CoS on Node device access interfaces are:

- You specify a Node device *fabric* interface instead of a Node device *access* interface when you apply CoS to an interface.
- You cannot attach classifiers, congestion notification profiles, or rewrite rules to fabric interfaces. Also, you cannot configure buffer settings on fabric interfaces. You can only attach fc-sets and traffic control profiles.

Default Scheduling on Node Device Fabric Interfaces

Default scheduling on Node device fabric interfaces is the same as default scheduling on Node device access interfaces. Only the default forwarding classes (best-effort, network-control, fcoe, no-loss, and multideestination) receive port bandwidth, based on the default minimum guaranteed bandwidth (transmit rate) scheduler settings for each default forwarding class.

To transport traffic on Node device fabric interfaces, the system organizes the default forwarding classes into three *class groups*. Class groups are not user-configurable. The three class groups are:

- **Unicast**—All traffic in the default forwarding classes **best-effort**, **network-control**, **fcoe**, and **no-loss** belong to this default class group.
- **Multideestination**—All traffic in the default forwarding class **mcast** belongs to this default class group.
- **Strict-high priority**—There is no default strict-high priority forwarding class, so there is no default strict-high priority class group and there is no default configuration for strict-high priority traffic.



NOTE: If you configure strict-high priority forwarding classes, you must also configure an fc-set (priority group) for strict-high priority traffic, map the strict-high priority forwarding classes to the strict-high priority fc-set, create a scheduler for the strict-high priority traffic and map it to the strict-high priority forwarding classes, create a traffic control profile for the strict-high priority traffic, and apply the strict-high priority fc-set and traffic control profile to the appropriate fabric interfaces.

The default forwarding classes receive port bandwidth based on their default transmit rate settings (weights). Forwarding classes that are not default forwarding classes receive no default bandwidth.

Default class group scheduling uses weighted round-robin (WRR) scheduling, in which each class group receives a portion of the total available fabric interface bandwidth based on the class group traffic type, as shown in [Table 80 on page 201](#). Within each class

group, the scheduler bandwidth allocation for individual forwarding classes is based on the default transmit rate for each forwarding class.

Table 80: Class Group Default Scheduling Properties and Membership on Node Device Fabric Interfaces

Class Group	Forwarding Class Mapping and Bandwidth Allocation (Default Transmit Rate)	Class Group Scheduling Properties (Weight)
Unicast	<ul style="list-style-type: none"> • best-effort (5%) • fcoe (35%) • no-loss (35%) • network-control (5%) 	Traffic in the unicast class group receives an 80% weight in the weighted round-robin (WRR) calculations. After the strict-high priority class group has been served, the unicast class group receives 80% of the remaining fabric bandwidth. (If more bandwidth is available, the unicast class group can use more bandwidth.)
Multidestination	<ul style="list-style-type: none"> • mcast (20%) 	Traffic in the multidestination class group receives a 20% weight in the WRR calculations. After the strict-high priority class group has been served, the multidestination class group receives 20% of the remaining fabric bandwidth. (If more bandwidth is available, the multidestination class group can use more bandwidth.)



NOTE: Strict-high priority traffic is served first, before any other traffic is served. Strict-high priority traffic receives all of the bandwidth it needs to empty its queues and therefore can starve other types of traffic during periods of high-volume strict-high priority traffic. Plan carefully and use caution when determining how much traffic to configure as strict-high priority traffic. We recommend that you always configure a shaping rate in the strict-high priority scheduler to set a maximum bandwidth limit for strict-high priority traffic.

Configuring Scheduling on Node Device Fabric Interfaces

If you do not want to use default CoS scheduling on Node device fabric interfaces, you can configure two-tier hierarchical scheduling (ETS) the same way that you configure ETS on Node device access interfaces.

- [Similarities Between Node Device Fabric Interface and Access Interface Scheduling on page 202](#)
- [Differences Between Node Device Fabric Interface and Access Interface Scheduling on page 202](#)

[Similarities Between Node Device Fabric Interface and Access Interface Scheduling](#)

Configuring scheduling on a Node device fabric interface is similar to configuring scheduling on an access interface in many ways. In both cases, you configure:

- Schedulers to specify the output scheduling for forwarding class traffic
- Scheduler maps to map schedulers to forwarding classes
- Forwarding classes (or use the default forwarding classes)
- Forwarding class sets (groups of forwarding classes that require similar CoS treatment)
- A separate fc-set for strict-high priority traffic (an fc-set cannot contain a mix of strict-high priority traffic and traffic with a different priority)
- Traffic control profiles to specify the output scheduling for fc-sets
- Traffic control profile and fc-set mapping to interfaces

On Node device fabric interfaces, you configure ETS in the same way, and ETS works the same way as on Node device access interfaces

In addition, strict-high priority queues are served first, and then the remaining port bandwidth is allocated to other traffic. Unless you configure a shaping rate in the scheduler for strict-high priority traffic, a strict-high priority queue can consume all of the port bandwidth and starve other queues, so we recommend that you always configure a shaping rate on strict-high priority traffic.

[Differences Between Node Device Fabric Interface and Access Interface Scheduling](#)

Configuring scheduling on a Node device fabric interface differs from configuring scheduling on an access interface in several ways. On fabric interfaces:

- You cannot attach classifiers.
- You cannot attach congestion notification profiles (flow control is applied automatically to lossless forwarding classes).
- You cannot attach rewrite rules.
- You cannot configure buffer settings.
- You specify a Node device fabric interface name instead of a Node device access interface name when you apply CoS to an interface.

Related Documentation

- [Understanding CoS Fabric Forwarding Class Sets on page 186](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Understanding Default CoS Settings on page 37](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 203](#)

- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring WRED Drop Profiles on page 279](#)
- [Example: Configuring Drop Profile Maps on page 281](#)

Understanding Default CoS Scheduling on QFabric System Interconnect Devices (Junos OS Release 13.1 and Later Releases)

The default class-of-service (CoS) properties on the QFabric system Interconnect device interfaces are optimized to best utilize the fabric resources. You cannot configure CoS properties on QFabric System Interconnect device interfaces.

- [Hierarchical CoS Architecture Across a QFabric System Interconnect Device on page 203](#)
- [Default CoS on Interconnect Device Fabric Interfaces on page 205](#)

Hierarchical CoS Architecture Across a QFabric System Interconnect Device

Because Interconnect devices support traffic from multiple Node devices that have multiple CoS configurations, CoS on Interconnect device fabric interfaces differs from CoS on Node device access and fabric interfaces.

The hierarchical CoS scheduling structure on the Interconnect device interfaces consists of two tiers:

1. Fabric forwarding class sets—Similar to fc-sets on Node devices, fabric fc-sets group traffic for transport across the Interconnect device fabric. Fabric fc-sets are global and apply to all traffic that crosses the fabric from all Node devices. See [“Understanding CoS Fabric Forwarding Class Sets” on page 186](#) for a detailed description of fabric fc-sets.
2. Class groups—Fabric fc-sets are grouped into class groups for transport across the Interconnect device.

Node devices and Interconnect devices each have a two-tier hierarchical CoS scheduling architecture. The architectures are slightly different, but each tier of the scheduling hierarchy performs analogous functions, as shown in [Table 81 on page 203](#).

Table 81: Hierarchical Scheduler Architecture on Node Devices and Interconnect Devices

Bandwidth Pool	Bandwidth Configuration on Node Devices	Bandwidth Configuration on Interconnect Devices
Port—Entire amount of bandwidth available to traffic on a port.	Access (xe) or fabric (fte) interfaces	Fabric (fte) or Clos fabric (bfte) interfaces

Table 81: Hierarchical Scheduler Architecture on Node Devices and Interconnect Devices (*continued*)

Bandwidth Pool	Bandwidth Configuration on Node Devices	Bandwidth Configuration on Interconnect Devices
Priority group—Group of traffic types that requires similar CoS treatment. Each priority group receives a portion of the total available port bandwidth.	Forwarding class set (fc-set)	Class group
Priority—Most granular tier of bandwidth allocation. Each priority receives a portion of the total available priority group bandwidth.	Forwarding class (mapped to output queue)	Fabric fc-set (mapped to output queue)

Fabric FC-Sets

Fabric fc-sets are groups of forwarding classes that receive similar CoS treatment across the Interconnect device. Fabric fc-sets are global to the QFabric system and apply to all traffic that traverses the fabric, from all connected Node devices. The CoS on a fabric fc-set applies to all the traffic that belongs to that fabric fc-set.

For example, a fabric fc-set that includes the **best-effort** forwarding class handles all of the **best-effort** traffic from all of the connected Node devices that traverses the Interconnect device fabric.

There are 12 default fabric fc-sets, including 5 visible fabric fc-sets and 7 hidden fabric fc-sets. The five visible fabric fc-sets have forwarding classes mapped to them by default. By default, the seven hidden fabric fc-sets do not carry traffic, but you can map forwarding classes to the hidden fabric fc-sets if you want to use them.

You can configure the forwarding class membership of each fabric fc-set. However, you cannot create new fabric fc-sets, and you cannot delete the 12 default fabric fc-sets.

Each fabric fc-set is mapped to an output queue. Each fabric interface has 12 output queues, one for each of the 12 fabric fc-sets. The traffic from all of the forwarding classes mapped to a fabric fc-set uses that fabric fc-set's output queue.

Fabric fc-sets are grouped into class groups for transport across the Interconnect device.

Class Groups for Fabric FC-Sets

To transport traffic across the fabric, the fabric organizes the fabric fc-sets into three classes called *class groups*. Class groups are not user-configurable. The three class groups are:

- **Strict-high priority**—All traffic in the fabric fc-set **fabric_fcset_strict_high**. This class group includes the traffic in **strict-high** priority and **network-control** forwarding classes, and in any forwarding classes you create on a Node device that consist of **strict-high** priority traffic.

- **Unicast**—All traffic in the fabric fc-sets **fabric_fcset_be**, **fabric_fcset_noloss1**, and **fabric_fcset_noloss2**. This class group includes the traffic in the **best-effort**, **fcoe**, and **no-loss** forwarding classes, and the traffic in any forwarding classes you create on a Node device that consist of best-effort or lossless unicast traffic. If you use any of the hidden no loss fabric fc-sets (**fabric_fcset_noloss3**, **fabric_fcset_noloss4**, **fabric_fcset_noloss5**, or **fabric_fcset_noloss6**), that traffic is part of this class group.
- **Multidestination**—All traffic in the fabric fc-set **fabric_fcset_multicast1**. This class group includes the traffic in the **mcast** forwarding class and in any forwarding classes you create on a Node device that consist of multidestination traffic. If you use any of the hidden multidestination fabric fc-sets (**fabric_fcset_multicast2**, **fabric_fcset_multicast3**, or **fabric_fcset_multicast4**), that traffic is part of this class group.

Default CoS on Interconnect Device Fabric Interfaces

The Interconnect device interfaces use the default CoS configuration as described in these sections:

- [Default Class Group Scheduling on page 205](#)
- [Default Fabric FC-Set Scheduling on page 206](#)
- [Default Class Group and Fabric FC-Set Scheduling Example on page 209](#)
- [Default PFC and Lossless Transport Across the Interconnect Device on page 210](#)

Default Class Group Scheduling

Default class group bandwidth scheduling is analogous to default fc-set (priority group) scheduling on a Node device. Default class group scheduling uses weighted round-robin (WRR) scheduling, in which each class group receives a portion of the total available fabric interface bandwidth, based on the class group's traffic type, as shown in

[Table 76 on page 193](#):

Table 82: Class Group Default Scheduling Properties and Membership

Class Group	Fabric fc-sets	Forwarding Classes (Default Mapping)	Class Group Scheduling Properties (Weight)
Strict-high priority	fabric_fcset_strict_high	<ul style="list-style-type: none"> • All strict-high priority forwarding classes • network-control 	Traffic in the strict-high priority class group is served first. This class group receives all of the bandwidth it needs to empty its queues and therefore can starve other types of traffic during periods of high-volume strict priority traffic. Plan carefully and use caution when determining how much traffic to configure as strict-high priority traffic.

Table 82: Class Group Default Scheduling Properties and Membership (*continued*)

Class Group	Fabric fc-sets	Forwarding Classes (Default Mapping)	Class Group Scheduling Properties (Weight)
Unicast	<ul style="list-style-type: none"> • fabric_fcset_be • fabric_fcset_noloss1 • fabric_fcset_noloss2 <p>Includes the hidden lossless fabric fc-sets if used:</p> <ul style="list-style-type: none"> • fabric_fcset_noloss3 • fabric_fcset_noloss4 • fabric_fcset_noloss5 • fabric_fcset_noloss6 	<ul style="list-style-type: none"> • best-effort • fcoe • no-loss <p>NOTE: No forwarding classes are mapped to the hidden lossless fabric_fcsets by default.</p>	Traffic in the unicast class group receives an 80% weight in the weighted round-robin (WRR) calculations. After the strict-high priority class group has been served, the unicast class group receives 80% of the remaining fabric bandwidth. (If more bandwidth is available, the unicast class group can use more bandwidth.)
Multidestination	<p>fabric_fcset_multicast1</p> <p>Includes the hidden multidestination fabric fc-sets if used:</p> <ul style="list-style-type: none"> • fabric_fcset_multicast2 • fabric_fcset_multicast3 • fabric_fcset_multicast4 	<ul style="list-style-type: none"> • mcast <p>NOTE: No forwarding classes are mapped to the hidden multidestination fabric_fcsets by default.</p>	Traffic in the multidestination class group receives a 20% weight in the WRR calculations. After the strict-high priority class group has been served, the multidestination class group receives 20% of the remaining fabric bandwidth. (If more bandwidth is available, the multidestination class group can use more bandwidth.)

Only the five visible fabric fc-sets have traffic mapped to them by default. The fabric fc-sets within each class group are weighted by their transmit rates (guaranteed minimum bandwidth), and they receive bandwidth from the class group's total bandwidth using weighted round-robin (WRR) scheduling.

Default Fabric FC-Set Scheduling

Default fabric fc-set bandwidth scheduling is analogous to default forwarding class (priority) scheduling on a Node device. Each fabric fc-set receives a guaranteed minimum percentage of the port bandwidth that the class group receives. The guaranteed minimum percentage is called the *transmit rate*.

Table 83 on page 206 shows the default transmit rate for each of the default fabric fc-sets.

Table 83: Default Fabric FC-Set Scheduler Configuration

Default Fabric FC-Set	Transmit Rate (Percentage of Class Group Bandwidth)
fabric_fcset_strict_high	N/A
	Strict-high priority traffic is served first, before any other traffic is served. Strict-high priority traffic receives all of the bandwidth it needs to empty its queues and therefore can starve other types of traffic during periods of high-volume strict priority traffic. Plan carefully and use caution when determining how much traffic to configure as strict-high priority traffic.

Table 83: Default Fabric FC-Set Scheduler Configuration (*continued*)

Default Fabric FC-Set	Transmit Rate (Percentage of Class Group Bandwidth)
fabric_fcset_noloss1	35%
fabric_fcset_noloss2	35%
fabric_fcset_be	10%
fabric_fcset_multicast1	20%

Each fabric fc-set belongs to a class group. Each class group receives a portion of the total available port bandwidth. Each fabric fc-set in a class group receives a portion of the total available class group bandwidth based on the transmit rate (weight) of the fabric fc-set.

Traffic in fabric_fcset_strict_high does not have a default transmit rate because fabric_fcset_strict_high receives all of the bandwidth needed to empty its queue before other queues are served. Traffic in the remaining fabric fc-sets receives bandwidth in a ratio proportional to the default transmit rate of each fabric fc-set.

Each of the following hidden fabric fc-sets receives a default scheduling weight of 1:

- fabric_fcset_noloss3
- fabric_fcset_noloss4
- fabric_fcset_noloss5
- fabric_fcset_noloss6
- fabric_fcset_multicast2
- fabric_fcset_multicast3
- fabric_fcset_multicast4

You must explicitly map forwarding classes to hidden fabric fc-sets if you want to use the hidden fabric fc-sets.



CAUTION: Bandwidth is allocated to fabric fc-sets based on scheduling weight. The scheduling weights of the visible (default) fabric fc-sets are the same as their transmit rates, so in the unicast class group, fabric_fcset_noloss1 and fabric_fcset_noloss2 each have a weight of 35 and fabric_fcset_be has a weight of 10. In the multidestination class group, the default fabric_fcset_multicast1 has a weight of 20. The hidden multicast and noloss fabric fc-sets each have a scheduling weight of 1.

The scheduling weights mean that when the visible fabric fc-sets are fully utilizing their allocated bandwidth:

- The hidden noloss fc-sets (fabric_fcset_noloss3, fabric_fcset_noloss4, fabric_fcset_noloss5, and fabric_fcset_noloss6) receive bandwidth at a proportional rate of 1:35 compared to the default noloss fc-sets.
- The hidden multicast fc-sets (fabric_fcset_multicast2, fabric_fcset_multicast3, and fabric_fcset_multicast4) receive bandwidth at a proportional rate of 1:20 compared to the default multicast fc-sets.

If you map traffic to a hidden fabric fc-set, that fabric fc-set receives the proportional amount of class group bandwidth that corresponds to its scheduling weight (1). The amount of bandwidth allocated to a hidden fabric fc-set depends on how much bandwidth the other fc-sets in the same class group consume. When the visible fabric fc-sets fully utilize their bandwidth, hidden fabric fc-sets receive only their minimum weight in bandwidth. (However, even a low scheduling weight results in a relatively large absolute bandwidth allocation because each fabric port is a 40-Gbps port.)

For example, if fabric_fcset_noloss1 and fabric_fcset_noloss2 each consume all of the 35 percent of bandwidth allocated to them, and fabric_fcset_be consumes all of the 10 percent of bandwidth allocated to it, then fabric_fcset_noloss3, fabric_fcset_noloss4, fabric_fcset_noloss5, and fabric_fcset_noloss6 receive bandwidth at a rate of 1:80 compared to the visible noloss fabric fc-sets. (If the visible fabric fc-sets do not use all of their allocated bandwidth, then the hidden fabric fc-sets receive more bandwidth.)

Another example is if we map lossless traffic to fabric_fcset_noloss3 and to fabric_fcset_noloss4. Fabric_fcset_noloss1 uses 10 percent of its 35 percent allocation of unicast class group bandwidth. Fabric_fcset_noloss2 uses 15 percent of its 35 percent allocation of unicast class group bandwidth. Fabric_fcset_be uses 5 percent of its allocated bandwidth. Fabric_fcset_noloss3 and fabric_fcset_noloss4 can use the remaining unicast class group bandwidth allocated to lossless traffic. However, if the traffic on fabric_fcset_noloss1, fabric_fcset_noloss2, or fabric_fcset_be increases, the bandwidth allocated to the hidden fabric fc-sets decreases.

Similarly, if you map traffic to a hidden multidestination fabric fc-set (fabric_fcset_multicast2, fabric_fcset_multicast3, fabric_fcset_multicast4), that multidestination fabric fc-set receives the proportional amount of class group bandwidth that corresponds to its scheduling weight (1). The amount of bandwidth allocated to a hidden multidestination fabric fc-set depends on how much bandwidth the other fc-sets in the multidestination class group consume. When fabric_fcset_multicast1 (the visible fabric fc-set) fully utilizes its bandwidth, hidden fabric fc-sets receive only their minimum weight in bandwidth. For example, if fabric_fcset_multicast1 uses its full bandwidth allocation, then the hidden multidestination fabric fc-sets receive bandwidth at a rate of 1:20 compared to fabric_fcset_multicast1.

Default Class Group and Fabric FC-Set Scheduling Example

The following example shows how default scheduling allocates the total port bandwidth among the class groups and their fabric fc-sets. In the example, traffic is mapped to each of the forwarding classes in the five visible fabric fc-sets, and the strict-high priority class group consumes an average of 10 percent of the 40-Gbps fabric interface bandwidth (4 gigabits), leaving 90 percent of the fabric interface bandwidth (36 gigabits) for the remaining class groups.

In this scenario, by default, the strict-high priority class group includes one fabric fc-set (fabric_fcset_strict_high), the unicast class group includes three fabric fc-sets (fabric_fcset_be, fabric_fcset_noloss1, and fabric_fcset_noloss2), and the multidestination class group includes one fabric fc-set (fabric_fcset_multicast1). Each individual fabric fc-set receives the following treatment:

- Strict-high priority class group (fabric_fcset_strict_high)—This group is assumed to average 10 percent (4 gigabits) for the purposes of this example. Because the strict-high priority class group is served first and receives all of the bandwidth it requires to empty its queue, in real networks the amount of required bandwidth fluctuates and affects the amount of bandwidth available to the other class groups.



TIP: To prevent strict-high priority traffic from using too much bandwidth, you can set a maximum bandwidth limit by configuring a scheduler shaping rate for the fabric_fcset_strict_high fabric fc-set.

- Unicast class group (fabric_fcset_be, fabric_fcset_noloss1, and fabric_fcset_noloss2)—Each of these fabric fc-sets receives a weighted portion of the 80 percent of the total port bandwidth available after the strict-high traffic has been served. The weight corresponds to the transmit rate of each fabric fc-set. The following calculations show the minimum port bandwidth allocated to each of the unicast class group fabric fc-sets:

- fabric_fcset_be

$10 / (35 + 35 + 10) \%$ of 80% of the available port bandwidth (12.5 percent of 80 percent of port bandwidth)

The 10 that is the numerator in $10 / (35 + 35 + 10)$ is the percentage of bandwidth allocated to the fabric_fcset_be by the transmit rate weight. The $(35 + 35 + 10)$ in the denominator sums the percentage of bandwidth (transmit rate weights) allocated to each of the three fabric fc-sets in the unicast class group.

The 80 percent represents 80 percent of the port bandwidth available after strict-high priority traffic is served (36 gigabits).

The resulting equation is:

$10 / (35 + 35 + 10) \% \times (0.8 \times 36 \text{ gigabits}) = \text{approximately } 3.6 \text{ gigabits}$

- fabric_fcset_noloss1 and fabric_fcset_noloss2

The default minimum bandwidth for the two visible lossless fabric fc-sets is the same because both of these fabric fc-sets have the same transmit rate weight.

$35 / (35 + 35 + 10)\%$ of 80% of the port bandwidth (43.75 percent of 80 percent of port bandwidth)

The 35 that is the numerator in $35 / (35 + 35 + 10)$ is the percentage of bandwidth allocated to each of the no-loss fabric fc-sets by the transmit rate weight. The $(35 + 35 + 10)$ in the denominator sums the percentage of bandwidth (transmit rate weights) allocated to each of the three fabric fc-sets in the unicast class group.

The 80 percent represents 80 percent of the port bandwidth available after strict-high priority traffic is served (36 gigabits).

The resulting equation is:

$35 / (35 + 35 + 10)\% \times (0.8 \times 36 \text{ gigabits}) = \text{approximately } 12.6 \text{ gigabits}$

- Multidestination class group (`fabric_fcset_multicast1`)—Because only one fabric fc-set is configured by default in the multidestination class group, it receives 100 percent of the 20 percent of the total port bandwidth available to the multidestination class group after the strict-high traffic has been served:

$100 / (100)\%$ of 20% of the available port bandwidth (100 percent of 20 percent of available port bandwidth)

The resulting equation is:

$100 / 100\% \times (0.2 \times 36 \text{ gigabits}) = \text{approximately } 7.2 \text{ gigabits}$

Default PFC and Lossless Transport Across the Interconnect Device

The Interconnect device incorporates flow control mechanisms to support lossless transport during periods of congestion on the fabric. To support the priority-based flow control (PFC) feature on the Node devices, the Interconnect device fabric supports lossless transport for up to six IEEE 802.1p priorities when the following two configuration constraints are met:

1. The IEEE 802.1p priority used for the traffic that requires lossless transport is mapped to a lossless forwarding class (a forwarding class configured with the **no-loss** parameter or the default **fcoe** or **no-loss** forwarding class).
2. The lossless forwarding class must be mapped to one of the lossless fabric fc-sets (**fabric_fcset_noloss1**, **fabric_fcset_noloss2**, **fabric_fcset_noloss3**, **fabric_fcset_noloss4**, **fabric_fcset_noloss5**, or **fabric_fcset_noloss6**). If you do not explicitly map lossless forwarding classes to fabric fc-sets, lossless forwarding classes are mapped by default to lossless fabric fc-sets **fabric_fcset_noloss1** and **fabric_fcset_noloss2**.

When traffic meets these two constraints, the fabric propagates back-pressure from egress queues during periods of congestion. However, to achieve end-to-end lossless transport across the QFabric system, you must also configure a congestion notification profile to enable PFC on the Node device ingress interfaces. To achieve end-to-end lossless transport across the network, you must configure PFC on all of the devices in the lossless traffic path.

For all other combinations of IEEE 802.1p priority to forwarding class mapping and all other combinations of forwarding class to fabric fc-set mapping, the default congestion control mechanism is normal packet drop. For example:

- **Case 1**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric_fcset_noloss1** fabric fc-set, then the congestion control mechanism is PFC.
- **Case 2**—If the IEEE 802.1p priority 5 is mapped to the lossless **fcoe** forwarding class, and the **fcoe** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop, and the traffic does not receive lossless treatment.
- **Case 3**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric_fcset_noloss2** fabric fc-set, then the congestion control mechanism is PFC.
- **Case 4**—If the IEEE 802.1p priority 5 is mapped to the lossless **no-loss** forwarding class, and the **no-loss** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop, and the traffic does not receive lossless treatment.
- **Case 5**—If the IEEE 802.1p priority 5 is mapped to the lossy **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric_fcset_be** fabric fc-set, then the congestion control mechanism is packet drop.
- **Case 6**—If the IEEE 802.1p priority 5 is mapped to the lossy **best-effort** forwarding class, and the **best-effort** forwarding class is mapped to the **fabric_fcset_noloss1** fabric fc-set, then the congestion control mechanism is packet drop.



NOTE: Lossless transport across the fabric must also meet the following two conditions:

1. The maximum cable length between the Node device and the Interconnect device is 150 meters of fiber cable.
2. The maximum frame size is 9216 bytes.

If the MTU is 9216 KB, in some cases the QFabric system supports only five lossless forwarding classes instead of six lossless forwarding classes because of headroom buffer limitations.

The number of IEEE 802.1p priorities (forwarding classes) the QFabric system can support for lossless transport across the Interconnect device fabric depends on several factors:

- **Approximate fiber cable length**—The longer the fiber cable that connects Node device fabric (FTE) ports to the Interconnect device fabric ports, the more data the connected ports need to buffer when a pause is asserted. (The longer the fiber cable, the more frames are traversing the cable when a pause is asserted. Each port must be able to store all of the “in transit” frames in the buffer to preserve lossless behavior and avoid dropping frames.)

- MTU size—The larger the maximum frame sizes the buffer must hold, the fewer frames the buffer can hold. The larger the MTU size, the more buffer space each frame consumes.
- Total number of Node device fabric ports connected to the Interconnect device—The higher the number of connected fabric ports, the more headroom buffer space the Node device needs on those fabric ports to support the lossless flows that traverse the Interconnect device. Because more buffer space is used on the Node device fabric ports, less buffer space is available for the Node device access ports, and a lower total number of lossless flows are supported.

The QFabric system supports six lossless priorities (forwarding classes) under most conditions. The priority group headroom that remains after allocating headroom to lossless flows is sufficient to support best-effort and multidestination traffic.

Table 77 on page 196 shows how many lossless priorities the QFabric system supports under different conditions (fiber cable lengths and MTUs) in cases when the QFabric system supports fewer than six lossless priorities. The number of lossless priorities is the same regardless of how many Node device FTE ports are connected to the Interconnect device. However, the higher the number of FTE ports connected to the Interconnect device, the lower the number of total lossless flows supported. In all cases that are not shown in Table 77 on page 196, the QFabric system supports six lossless priorities.



NOTE: The system does not perform a configuration commit check that compares available system resources with the number of lossless forwarding classes configured. If you commit a configuration with more lossless forwarding classes than the system resources can support, frames in lossless forwarding classes might be dropped.

Table 84: Lossless Priority (Forwarding Class) Support for QFX3500 and QFX3600 Node Devices When Fewer than Six Lossless Priorities Are Supported

MTU in Bytes	Fiber Cable Length in Meters (Approximate)	Maximum Number of Lossless Priorities (Forwarding Classes) on the Node Device
9216 (9K)	100	5
9216 (9K)	150	5



NOTE: The total number of lossless flows decreases as resource consumption increases. For a Node device, the higher the number of FTE ports connected to the Interconnect device, the larger the MTU, and the longer the fiber cable length, the fewer total lossless flows the QFabric system can support.

- Related Documentation**
- [Understanding CoS Fabric Forwarding Class Sets on page 186](#)
 - [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 199](#)

CHAPTER 4

Learn About Technology

- [Data Center Technology Overview Videos on page 215](#)

Data Center Technology Overview Videos

Juniper Information Experience (IX) videos provide brief, high-level overviews of data center technologies and concepts. Each video runs approximately one-and-a-half to two minutes in length. This document contains SDN-related videos and links to conceptual documents that contain other data center technology videos:

- [Learn About Video: Why Do We Need an IP Fabric? on page 215](#)
- [Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric? on page 215](#)
- [Learn About Video: Why Use an Overlay Network in a Data Center? on page 215](#)
- [Conceptual Documents That Contain Technology Overview Videos on page 216](#)

Learn About Video: Why Do We Need an IP Fabric?

The video *Why Do We Need an IP Fabric?* presents a brief overview of IP Fabric use cases.



Video: [Why Do We Need an IP Fabric?](#)

Learn About Video: What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?

The video *What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?* presents a brief overview of the arguments for using Border Gateway Protocol (BGP) as the data center IP fabric control plane protocol.



Video: [What is the Best Control Plane Protocol to Use in a Data Center IP Fabric?](#)

Learn About Video: Why Use an Overlay Network in a Data Center?

The video *Why Use an Overlay Network in a Data Center?* presents a brief overview of the advantages of data center overlay networks.



Video: [Why Use an Overlay Network in a Data Center?](#)

Conceptual Documents That Contain Technology Overview Videos

The following conceptual documents include brief video overviews of the technology:

- [Understanding DCB Features and Requirements on page 156](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)
- [Understanding DCBX on page 170](#)
- *Understanding PFC Functionality Across Layer 3 Interfaces*
- *Virtual Chassis Fabric Overview*
- *Understanding In-Service Software Upgrade (ISSU) and In-Service Software Upgrade (ISSU) System Requirements (same video)*

PART 2

Configuration

- [Configuration Examples on page 219](#)
- [Configuration Tasks on page 401](#)
- [Configuration Statements on page 439](#)

CHAPTER 5

Configuration Examples

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 241](#)
- [Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG on page 249](#)
- [Example: Configuring Unicast Classifiers on page 273](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 277](#)
- [Example: Configuring WRED Drop Profiles on page 279](#)
- [Example: Configuring Drop Profile Maps on page 281](#)
- [Example: Configuring Forwarding Classes on page 283](#)
- [Example: Configuring Forwarding Class Sets on page 286](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring Queue Scheduling Priority on page 295](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 300](#)
- [Example: Configuring Maximum Output Bandwidth on page 305](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 309](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 317](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 340](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 356](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 365](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 371](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 377](#)

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 383](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)

Example: Configuring CoS Hierarchical Port Scheduling (ETS)

Hierarchical port scheduling defines the class-of-service (CoS) properties of output queues, which are mapped to forwarding classes (forwarding classes are mapped to IEEE 802.1p priorities, so mapping queues to forwarding classes also maps queues to priorities). Hierarchical port scheduling enables you to group priorities that require similar CoS resources into priority groups. You define the port bandwidth resources for a priority group, and you define the amount of the priority group's resources that each priority in the group can use.

Hierarchical port scheduling is the Junos OS implementation of enhanced transmission selection (ETS, described in IEEE 802.1Qaz). One major benefit of hierarchical port scheduling is greater port bandwidth utilization. If a priority group on a port does not use all of its allocated bandwidth, other priority groups on that port can use that bandwidth. Also, if a priority within a priority group does not use its allocated bandwidth, other priorities within that priority group can use that bandwidth.

Configuring hierarchical scheduling is a multistep procedure that includes:

- Mapping forwarding classes to queues
- Defining forwarding class sets (priority groups)
- Defining behavior aggregate classifiers
- Configuring priority-based flow control (PFC) for lossless priorities (queues)
- Applying classifiers and PFC configuration to ingress interfaces
- Defining drop profiles
- Defining schedulers
- Mapping forwarding classes to schedulers
- Defining traffic control profiles
- Assigning priority groups and traffic control profiles to egress ports

This example describes how to configure hierarchical scheduling:

- [Requirements on page 220](#)
- [Overview on page 221](#)
- [Configuration on page 224](#)
- [Verification on page 232](#)

Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch

- Junos OS Release 11.1 or later for the QFX Series

Overview

Keep the following considerations in mind when you plan the port bandwidth allocation for priority groups and for individual priorities:

- How much traffic and what types of traffic you expect to traverse the system.
- How you want to divide different types of traffic into priorities (forwarding classes, also called queues) to apply different CoS treatment to the traffic. Dividing traffic into priorities includes:
 - Mapping the code points of ingress traffic to forwarding classes using behavior aggregate (BA) classifiers. This classifies incoming traffic into the appropriate forwarding class.
 - Mapping forwarding classes to output queues. This defines the output queue for each type of traffic.
 - Attaching the BA classifier to the desired ingress interfaces so that incoming traffic maps to the desired forwarding classes and queues.
- How you want to organize priorities into priority groups (forwarding class sets).

Traffic that requires similar treatment usually belongs in the same priority group. To do this, place forwarding classes that require similar bandwidth, loss, and other characteristics in the same forwarding class set. For example, you can map all types of best-effort traffic forwarding classes into one forwarding class set.

- How much of the port bandwidth you want to allocate to each priority group and to each of the priorities in each priority group. The following considerations apply to bandwidth allocation:
 - Estimate how much traffic you expect in each forwarding class (output queue) and how much traffic you expect in each forwarding class set (the aggregate amount of traffic in the forwarding classes that belong to the forwarding class set).
 - The combined minimum guaranteed bandwidth of the priorities (forwarding classes) in a priority group should not exceed the minimum guaranteed bandwidth of the priority group. The transmit rate scheduler parameter defines the minimum guaranteed bandwidth for forwarding classes. Scheduler maps associate schedulers with forwarding classes.
 - The combined minimum guaranteed bandwidth of the priority groups (forwarding class sets) on a port should not exceed the port's total bandwidth. Traffic control profiles define the minimum bandwidth for a forwarding class set. Associating a scheduler map with a traffic control profile sets the scheduling for the individual forwarding classes in the forwarding class set.

This example creates hierarchical port scheduling by defining priority groups for best effort, guaranteed delivery, and high-performance computing (HPC) traffic. Each priority group includes priorities that need to receive similar CoS treatment. Each priority group and each priority within each priority group receive the CoS resources needed to service

their flows. Lossless priorities use PFC to prevent packet loss when the network experiences congestion.

Topology

Table 85 on page 222 shows the configuration components for this example.

Table 85: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology

Property	Settings
Hardware	QFX3500 switch
Mapping of forwarding classes (priorities) to queues	<p>best-effort to queue 0</p> <p>be to queue 1</p> <p>fcoe (Fibre Channel over Ethernet) to queue 3</p> <p>no-loss to queue 4</p> <p>hpc (high-performance computing) to queue 5</p> <p>network-control to queue 7</p> <p>NOTE: If you are using Junos OS Release 12.2 or later, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the default lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best-effort) traffic and does <i>not</i> receive lossless treatment.</p> <p>In Junos OS Release 12.3 and later, you can include the <i>no-loss</i> packet drop attribute in the explicit forwarding class configuration to configure a lossless forwarding class.</p>
Forwarding class sets (priority groups)	<p>best-effort-pg: contains forwarding classes best-effort, be, and network control</p> <p>guar-delivery-pg: contains forwarding classes fcoe and no-loss</p> <p>hpc-pg: contains forwarding class hpc</p>
Behavior aggregate classifier (maps forwarding classes and loss priorities to incoming packets by IEEE 802.1 code point)	<p>Name—hsclassifier1</p> <p>Code point mapping:</p> <ul style="list-style-type: none"> • 000 to forwarding class best-effort and loss priority low • 001 to forwarding class be and loss priority high • 011 to forwarding class fcoe and loss priority low • 100 to forwarding class no-loss and loss priority low • 101 to forwarding class hpc and loss priority low • 110 to forwarding class network-control and loss priority low
PFC	<p>Congestion notification profile name—gd-cnp</p> <p>PFC enabled on code points: 011 (fcoe priority), 010 (no-loss priority)</p>

Table 85: Components of the Hierarchical Port Scheduling (ETS) Configuration Topology (continued)

Property	Settings
Drop profiles	<p>dp-be-low: drop start point 25, drop end point 50, maximum drop rate 80</p> <p>dp-be-high: drop start point 10, drop end point 40, maximum drop rate 100</p> <p>dp-hpc: drop start point 75, drop end point 90, maximum drop rate 75</p> <p>dp-nc: drop start point 80, drop end point 100, maximum drop rate 100</p>
Queue schedulers	<p>be-sched: minimum bandwidth 3g, maximum bandwidth 100%, priority low, drop profiles dp-be-low and dp-be-high</p> <p>fcoe-sched: minimum bandwidth 2.5g, maximum bandwidth 100%, priority low</p> <p>hpc-sched: minimum bandwidth 2g, maximum bandwidth 100%, priority low, drop profile dp-hpc</p> <p>nc-sched: minimum bandwidth 500m, maximum bandwidth 100%, priority low, drop profile dp-nc</p> <p>nl-sched: minimum bandwidth 2g, maximum bandwidth 100%, priority low</p>
Forwarding class-to-scheduler mapping	<p>Scheduler map be-map: Forwarding class best-effort, scheduler be-sched Forwarding class be, scheduler be-sched Forwarding class network-control, scheduler nc-sched</p> <p>Scheduler map gd-map: Forwarding class fcoe, scheduler fcoe-sched Forwarding class no-loss, scheduler nl-sched</p> <p>Scheduler map hpc-map: Forwarding class hpc, scheduler hpc-sched</p>
Traffic control profiles	<p>be-tcp: scheduler map be-map, minimum bandwidth 3.5g, maximum bandwidth 100%</p> <p>gd-tcp: scheduler map gd-map, minimum bandwidth 4.5g, maximum bandwidth 100%</p> <p>hpc-tcp: scheduler map hpc-map, minimum bandwidth 2g, maximum bandwidth 100%</p>
Interfaces	<p>This example configures hierarchical port scheduling on interfaces xe-0/0/20 and xe-0/0/21. Because traffic is bidirectional, you apply the ingress and egress configuration components to both interfaces:</p> <ul style="list-style-type: none"> Classifier Name—hsclassifier1 Forwarding class sets—best-effort-pg, guar-deliver-pg, hpc-pg Congestion notification profile—gd-cnp

Figure 10 on page 224 shows a block diagram of the configuration components and the configuration flow of the CLI statements used in the example. You can perform the configuration steps in a different sequence if you want.

Figure 10: Hierarchical Port Scheduling Components Block Diagram





Figure 11 on page 224 shows a block diagram of the hierarchical scheduling packet flow from ingress to egress.

Figure 11: Hierarchical Port Scheduling Packet Flow Block Diagram



Configuration

CLI Quick Configuration

To quickly configure hierarchical port scheduling, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
[edit class-of-service]
set forwarding-classes class best-effort queue-num 0
set forwarding-classes class be2 queue-num 1
set forwarding-classes class hpc queue-num 5
set forwarding-classes class network-control queue-num 7
```

```

set forwarding-class-sets best-effort-pg class best-effort
set forwarding-class-sets best-effort-pg class be2
set forwarding-class-sets best-effort-pg class network-control
set forwarding-class-sets guar-delivery-pg class fcoe
set forwarding-class-sets guar-delivery-pg class no-loss
set forwarding-class-sets hpc-pg class hpc
set classifiers ieee-802.1 hsclassifier1 forwarding-class best-effort loss-priority low code-points
000
set classifiers ieee-802.1 hsclassifier1 forwarding-class be2 loss-priority high code-points 001
set classifiers ieee-802.1 hsclassifier1 forwarding-class fcoe loss-priority low code-points 011
set classifiers ieee-802.1 hsclassifier1 forwarding-class no-loss loss-priority low code-points 100
set classifiers ieee-802.1 hsclassifier1 forwarding-class hpc loss-priority low code-points 101
set classifiers ieee-802.1 hsclassifier1 forwarding-class network-control loss-priority low
code-points 110
set congestion-notification-profile gd-cnp input ieee-802.1 code-point 011 pfc
set congestion-notification-profile gd-cnp input ieee-802.1 code-point 100 pfc
set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 hsclassifier1
set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 hsclassifier1
set interfaces xe-0/0/20 congestion-notification-profile gd-cnp
set interfaces xe-0/0/21 congestion-notification-profile gd-cnp
set drop-profiles dp-be-low interpolate fill-level 25 fill-level 50 drop-probability 0 drop-probability
80
set drop-profiles dp-be-high interpolate fill-level 10 fill-level 40 drop-probability 0 drop-probability
100
set drop-profiles dp-nc interpolate fill-level 80 fill-level 100 drop-probability 0 drop-probability
100
set drop-profiles dp-hpc interpolate fill-level 75 fill-level 90 drop-probability 0 drop-probability
75
set schedulers be-sched priority low transmit-rate 3g
set schedulers be-sched shaping-rate percent 100
set schedulers be-sched drop-profile-map loss-priority low protocol any drop-profile dp-be-low
set schedulers be-sched drop-profile-map loss-priority high protocol any drop-profile dp-be-high
set schedulers fcoe-sched priority low transmit-rate 2500m
set schedulers fcoe-sched shaping-rate percent 100
set schedulers hpc-sched priority low transmit-rate 2g
set schedulers hpc-sched shaping-rate percent 100
set schedulers hpc-sched drop-profile-map loss-priority low protocol any drop-profile dp-hpc
set schedulers nc-sched priority low transmit-rate 500m
set schedulers nc-sched shaping-rate percent 100
set schedulers nc-sched drop-profile-map loss-priority low protocol any drop-profile dp-nc
set schedulers nl-sched priority low transmit-rate 2g
set schedulers nl-sched shaping-rate percent 100
set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
set scheduler-maps be-map forwarding-class be2 scheduler be-sched
set scheduler-maps be-map forwarding-class network-control scheduler nc-sched
set scheduler-maps gd-map forwarding-class fcoe scheduler fcoe-sched
set scheduler-maps gd-map forwarding-class no-loss scheduler nl-sched
set scheduler-maps hpc-map forwarding-class hpc scheduler hpc-sched
set traffic-control-profiles be-tcp scheduler-map be-map guaranteed-rate 3500m
set traffic-control-profiles be-tcp shaping-rate percent 100
set traffic-control-profiles gd-tcp scheduler-map gd-map guaranteed-rate 4500m
set traffic-control-profiles gd-tcp shaping-rate percent 100
set traffic-control-profiles hpc-tcp scheduler-map hpc-map guaranteed-rate 2g
set traffic-control-profiles hpc-tcp shaping-rate percent 100
set interfaces xe-0/0/20 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp
set interfaces xe-0/0/20 forwarding-class-set guar-delivery-pg output-traffic-control-profile
gd-tcp
set interfaces xe-0/0/20 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp
set interfaces xe-0/0/21 forwarding-class-set best-effort-pg output-traffic-control-profile be-tcp

```

```
set interfaces xe-0/0/21 forwarding-class-set guar-delivery-pg output-traffic-control-profile
gd-tcp
set interfaces xe-0/0/21 forwarding-class-set hpc-pg output-traffic-control-profile hpc-tcp
```

Step-by-Step Procedure

To perform a step-by-step configuration of the forwarding classes (priorities), forwarding class sets (priority groups), classifiers, queue schedulers, PFC, traffic control profiles, and interfaces to set up hierarchical port scheduling (ETS):

1. Configure the forwarding classes (priorities) and map them to unicast output queues (do not explicitly map the **fcoe** and **no-loss** forwarding classes to output queues; use the default configuration):

```
[edit class-of-service]
user@switch# set forwarding-classes class best-effort queue-num 0
user@switch# set forwarding-classes class be2 queue-num 1
user@switch# set forwarding-classes class hpc queue-num 5
user@switch# set forwarding-classes class network-control queue-num 7
```

2. Configure forwarding class sets (priority groups) to group forwarding classes (priorities) that require similar CoS treatment:

```
[edit class-of-service]
user@switch# set forwarding-class-sets best-effort-pg class best-effort
user@switch# set forwarding-class-sets best-effort-pg class be2
user@switch# set forwarding-class-sets best-effort-pg class network-control
user@switch# set forwarding-class-sets guar-delivery-pg class fcoe
user@switch# set forwarding-class-sets guar-delivery-pg class no-loss
user@switch# set forwarding-class-sets hpc-pg class hpc
```

3. Configure a classifier to set the loss priority and IEEE 802.1 code points assigned to each forwarding class at the ingress:

```
[edit class-of-service]
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class best-effort
loss-priority low code-points 000
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class be2 loss-priority
high code-points 001
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class fcoe loss-priority
low code-points 011
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class no-loss loss-priority
low code-points 100
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class hpc loss-priority low
code-points 101
user@switch# set classifiers ieee-802.1 hsclassifier1 forwarding-class network-control
loss-priority low code-points 110
```

4. Configure a congestion notification profile to enable PFC on the FCoE and no-loss queue IEEE 802.1 code points:

```
[edit class-of-service]
user@switch# set congestion-notification-profile gd-cnp input ieee-802.1 code-point 011
pfc
user@switch# set congestion-notification-profile gd-cnp input ieee-802.1 code-point 100
pfc
```

5. Assign the classifier to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 hsclassifier1
user@switch# set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 hsclassifier1
```

6. Apply the PFC configuration to the interfaces:


```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 congestion-notification-profile gd-cnp
user@switch# set interfaces xe-0/0/21 congestion-notification-profile gd-cnp
```
7. Configure the drop profile for the best-effort low loss-priority queue:


```
[edit class-of-service]
user@switch# set drop-profiles dp-be-low interpolate fill-level 25 fill-level 50
drop-probability 0 drop-probability 80
```
8. Configure the drop profile for the best-effort high loss-priority queue:


```
[edit class-of-service]
user@switch# set drop-profiles dp-be-high interpolate fill-level 10 fill-level 40
drop-probability 0 drop-probability 100
```
9. Configure the drop profile for the network-control queue:


```
[edit class-of-service]
user@switch# set drop-profiles dp-nc interpolate fill-level 80 fill-level 100 drop-probability
0 drop-probability 100
```
10. Configure the drop profile for the high-performance computing queue:


```
[edit class-of-service]
user@switch# set drop-profiles dp-hpc interpolate fill-level 75 fill-level 90 drop-probability
0 drop-probability 75
```
11. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profiles for the best-effort queue:


```
[edit class-of-service]
user@switch# set schedulers be-sched priority low transmit-rate 3g
user@switch# set schedulers be-sched shaping-rate percent 100
user@switch# set schedulers be-sched drop-profile-map loss-priority low protocol any
drop-profile dp-be-low
user@switch# set schedulers be-sched drop-profile-map loss-priority high protocol any
drop-profile dp-be-high
```
12. Define the minimum guaranteed bandwidth, priority, and maximum bandwidth for the FCoE queue:


```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 2500m
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```
13. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profile for the high-performance computing queue:


```
[edit class-of-service]
user@switch# set schedulers hpc-sched priority low transmit-rate 2g
user@switch# set schedulers hpc-sched shaping-rate percent 100
user@switch# set schedulers hpc-sched drop-profile-map loss-priority low protocol any
drop-profile dp-hpc
```
14. Define the minimum guaranteed bandwidth, priority, maximum bandwidth, and drop profile for the network-control queue:


```
[edit class-of-service]
user@switch# set schedulers nc-sched priority low transmit-rate 500m
user@switch# set schedulers nc-sched shaping-rate percent 100
```

- ```

user@switch# set schedulers nc-sched drop-profile-map loss-priority low protocol any
drop-profile dp-nc

```
15. Define the minimum guaranteed bandwidth, priority, and maximum bandwidth for the no-loss queue:
 

```

[edit class-of-service]
user@switch# set schedulers nl-sched priority low transmit-rate 2g
user@switch# set schedulers nl-sched shaping-rate percent 100

```
  16. Map the schedulers to the appropriate forwarding classes (queues):
 

```

[edit class-of-service]
user@switch# set scheduler-maps be-map forwarding-class best-effort scheduler be-sched
user@switch# set scheduler-maps be-map forwarding-class be2 scheduler be-sched
user@switch# set scheduler-maps be-map forwarding-class network-control scheduler
nc-sched
user@switch# set scheduler-maps gd-map forwarding-class fcoe scheduler fcoe-sched
user@switch# set scheduler-maps gd-map forwarding-class no-loss scheduler nl-sched
user@switch# set scheduler-maps hpc-map forwarding-class hpc scheduler hpc-sched

```
  17. Define the traffic control profile for the best-effort priority group (queue scheduler to mapping, minimum guaranteed bandwidth, and maximum bandwidth):
 

```

[edit class-of-service]
user@switch# set traffic-control-profiles be-tcp scheduler-map be-map guaranteed-rate
3500m
user@switch# set traffic-control-profiles be-tcp shaping-rate percent 100

```
  18. Define the traffic control profile for the guaranteed delivery priority group (queue to scheduler mapping, minimum guaranteed bandwidth, and maximum bandwidth):
 

```

[edit class-of-service]
user@switch# set traffic-control-profiles gd-tcp scheduler-map gd-map guaranteed-rate
4500m
user@switch# set traffic-control-profiles gd-tcp shaping-rate percent 100

```
  19. Define the traffic control profile for the high-performance computing priority group (queue to scheduler mapping, minimum guaranteed bandwidth, and maximum bandwidth):
 

```

[edit class-of-service]
user@switch# set traffic-control-profiles hpc-tcp scheduler-map hpc-map guaranteed-rate
2g
user@switch# set traffic-control-profiles hpc-tcp shaping-rate percent 100

```
  20. Apply the three priority groups (forwarding class sets) and the appropriate traffic control profiles to the egress ports:
 

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/20 forwarding-class-set best-effort-pg
output-traffic-control-profile be-tcp
user@switch# set interfaces xe-0/0/20 forwarding-class-set guar-delivery-pg
output-traffic-control-profile gd-tcp
user@switch# set interfaces xe-0/0/20 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set best-effort-pg
output-traffic-control-profile be-tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set guar-delivery-pg
output-traffic-control-profile gd-tcp
user@switch# set interfaces xe-0/0/21 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp

```



## Results

Display the results of the configuration (the system shows only the explicitly configured parameters; it does not show default parameters such as the **fcoe** and **no-loss** lossless forwarding classes):

```
user@switch> show configuration class-of-service
```

```
classifiers {
 ieee-802.1 hsclassifier1 {
 forwarding-class best-effort {
 loss-priority low code-points 000;
 }
 forwarding-class be2 {
 loss-priority high code-points 001;
 }
 forwarding-class fcoe {
 loss-priority low code-points 011;
 }
 forwarding-class no-loss {
 loss-priority low code-points 100;
 }
 forwarding-class hpc {
 loss-priority low code-points 101;
 }
 forwarding-class network-control {
 loss-priority low code-points 110;
 }
 }
}
drop-profiles {
 dp-be-low {
 interpolate {
 fill-level [25 50];
 drop-probability [0 80];
 }
 }
 dp-be-high {
 interpolate {
 fill-level [10 40];
 drop-probability [0 100];
 }
 }
 dp-hpc {
 interpolate {
 fill-level [75 90];
 drop-probability [0 75];
 }
 }
 dp-nc {
 interpolate {
 fill-level [80 100];
 drop-probability [0 100];
 }
 }
}
forwarding-classes {
```

```
class best-effort queue-num 0;
class be2 queue-num 1;
class hpc queue-num 5;
class network-control queue-num 7;
}
traffic-control-profiles {
 be-tcp {
 scheduler-map be-map;
 shaping-rate percent 100;
 guaranteed-rate 3500000000;
 }
 gd-tcp {
 scheduler-map gd-map;
 shaping-rate percent 100;
 guaranteed-rate 4500000000;
 }
 hpc-tcp {
 scheduler-map hpc-map;
 shaping-rate percent 100;
 guaranteed-rate 2g;
 }
}
forwarding-class-sets {
 guar-delivery-pg {
 class fcoe;
 class no-loss;
 }
 best-effort-pg {
 class best-effort;
 class be2;
 class network-control;
 }
 hpc-pg {
 class hpc;
 }
}
congestion-notification-profile {
 gd-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 code-point 100 {
 pfc;
 }
 }
 }
 }
}
interfaces {
 xe-0/0/20 {
 forwarding-class-set {
 best-effort-pg {
 output-traffic-control-profile be-tcp;
 }
 }
 }
}
```

```

 guar-delivery-pg {
 output-traffic-control-profile gd-tcp;
 }
 hpc-pg {
 output-traffic-control-profile hpc-tcp;
 }
}
congestion-notification-profile gd-cnp;
unit 0 {
 classifiers {
 ieee-802.1 hsclassifier1;
 }
}
}
xe-0/0/21 {
 forwarding-class-set {
 best-effort-pg {
 output-traffic-control-profile be-tcp;
 }
 guar-delivery-pg {
 output-traffic-control-profile gd-tcp;
 }
 hpc-pg {
 output-traffic-control-profile hpc-tcp;
 }
 }
 congestion-notification-profile gd-cnp;
 unit 0 {
 classifiers {
 ieee-802.1 hsclassifier1;
 }
 }
}
}
scheduler-maps {
 be-map {
 forwarding-class best-effort scheduler be-sched;
 forwarding-class network-control scheduler nc-sched;
 forwarding-class be2 scheduler be-sched;
 }
 gd-map {
 forwarding-class fcoe scheduler fcoe-sched;
 forwarding-class no-loss scheduler nl-sched;
 }
 hpc-map {
 forwarding-class hpc scheduler hpc-sched;
 }
}
schedulers {
 be-sched {
 transmit-rate 3g;
 shaping-rate percent 100;
 priority low;
 drop-profile-map loss-priority low protocol any drop-profile dp-be-low;
 drop-profile-map loss-priority high protocol any drop-profile dp-be-high;
 }
}

```

```
fcoe-sched {
 transmit-rate 2500000000;
 shaping-rate percent 100;
 priority low;
}
hpc-sched {
 transmit-rate 2g;
 shaping-rate percent 100;
 priority low;
 drop-profile-map loss-priority low protocol any drop-profile dp-hpc;
}
nc-sched {
 transmit-rate 500m;
 shaping-rate percent 100;
 priority low;
 drop-profile-map loss-priority low protocol any drop-profile dp-nc;
}
nl-sched {
 transmit-rate 2g;
 shaping-rate percent 100;
 priority low;
}
}
```



**TIP:** To quickly configure the interfaces, issue the **load merge terminal** command, and then copy the hierarchy and paste it into the switch terminal window.

## Verification

To verify that the hierarchical port scheduling components have been created and are operating properly, perform these tasks:

- [Verifying That the Forwarding Classes \(Priorities\) Have Been Created on page 233](#)
- [Verifying That the Forwarding Class Sets \(Priority Groups\) Have Been Created on page 233](#)
- [Verifying That the Classifier Has Been Created on page 234](#)
- [Verifying That Priority-Based Flow Control Has Been Enabled on page 234](#)
- [Verifying That the Output Queue Schedulers Have Been Created on page 235](#)
- [Verifying That the Drop Profiles Have Been Created on page 238](#)
- [Verifying That the Priority Group Output Schedulers \(Traffic Control Profiles\) Have Been Created on page 239](#)
- [Verifying the Interface Configuration on page 240](#)

### Verifying That the Forwarding Classes (Priorities) Have Been Created

**Purpose** Verify that the forwarding classes have been created and mapped to the correct queues. (The system shows only the explicitly configured forwarding classes. It does not show default forwarding classes such as **fcoe** and **no-loss**.)

**Action** List the forwarding classes using the operational mode command **show class-of-service forwarding-class**:

```
user@switch> show class-of-service forwarding-class
```

| Forwarding class | ID | Queue | Policing priority | No-Loss  |
|------------------|----|-------|-------------------|----------|
| best-effort      | 0  | 0     | normal            | Disabled |
| be2              | 1  | 3     | normal            | Disabled |
| hpc              | 2  | 4     | normal            | Disabled |
| network-control  | 3  | 7     | normal            | Disabled |
| mcast            | 8  | 8     | normal            | Disabled |

**Meaning** The **show class-of-service forwarding-class** command lists all of the configured forwarding classes, the internal identification number of each forwarding class, the queues that are mapped to the forwarding classes, the policing priority, and whether the forwarding class is lossless (no-loss packet drop attribute enabled) or lossy forwarding class (no-loss packet drop attribute disabled). The command output shows that:

- Forwarding class **best-effort** maps to queue **0** and is lossy
- Forwarding class **be2** maps to queue **1** and is lossy
- Forwarding class **hpc** maps to queue **5** and is lossy
- Forwarding class **network-control** maps to queue **7** and is lossy

In addition, the command lists the default multicast (multidestination) forwarding class and the default queue to which it is mapped.

### Verifying That the Forwarding Class Sets (Priority Groups) Have Been Created

**Purpose** Verify that the priority groups have been created and that the correct priorities (forwarding classes) belong to the appropriate priority group.

**Action** List the forwarding class sets using the operational mode command **show class-of-service forwarding-class-set**:

```
user@switch> show class-of-service forwarding-class-set
```

```
Forwarding class set: best-effort-pg, Type: normal-type, Forwarding class set
index: 19907
```

| Forwarding class | Index |
|------------------|-------|
| best-effort      | 0     |
| be2              | 1     |
| network-control  | 5     |

Forwarding class set: guar-delivery-pg, Type: normal-type, Forwarding class set index: 43700

| Forwarding class | Index |
|------------------|-------|
| fcoe             | 2     |
| no-loss          | 3     |

Forwarding class set: hpc-pg, Type: normal-type, Forwarding class set index: 60758

| Forwarding class | Index |
|------------------|-------|
| hpc              | 4     |

**Meaning** The `show class-of-service forwarding-class-set` command lists all of the configured forwarding class sets (priority groups), the forwarding classes (priorities) that belong to each priority group, and the internal index number of each priority group. The command output shows that:

- The forwarding class set **best-effort-pg** includes the forwarding classes **best-effort**, **be2**, and **network-control**.
- The forwarding class set **guar-delivery-pg** includes the forwarding classes **fcoe** and **no-loss**.
- The forwarding class set **hpc-pg** includes the forwarding class **hpc**.

### Verifying That the Classifier Has Been Created

**Purpose** Verify that the classifier maps forwarding classes to the correct IEEE 802.1p code points and packet loss priorities.

**Action** List the classifier configured for hierarchical port scheduling using the operational mode command `show class-of-service classifier name hsclassifier1`:

```
user@switch> show class-of-service classifier name hsclassifier1
Classifier: hsclassifier1, Code point type: ieee-802.1, Index: 43607
Code point Forwarding class Loss priority
000 best-effort low
001 be2 high
011 fcoe low
100 no-loss low
101 hpc low
110 network-control low
```

**Meaning** The `show class-of-service classifier name hsclassifier1` command lists all of the IEEE 802.1p code points and the loss priorities mapped to all of the forwarding classes in the classifier. The command output shows that the forwarding classes **best-effort**, **be2**, **no-loss**, **fcoe**, **hpc**, and **network-control** have been created and mapped to IEEE 802.1p code points and loss priorities.

### Verifying That Priority-Based Flow Control Has Been Enabled

**Purpose** Verify that PFC is enabled on the correct priorities for lossless transport.

**Action** List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Type: Input, Name: gd-cnp, Index: 51687
```

```
Cable Length: 100 m
```

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2500 |
| 100      | Enabled  | 2500 |
| 101      | Disabled |      |
| 110      | Disabled |      |
| 111      | Disabled |      |

```
Type: Output
```

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 000      | 0                   |
| 001      | 0                   |
| 010      | 1                   |
| 011      | 2                   |
| 100      | 3                   |
| 101      | 4                   |
| 110      | 5                   |
| 111      | 6                   |
|          | 7                   |

**Meaning** The **show class-of-service congestion-notification** command lists all of the congestion notification profiles and the IEEE 802.1p code points with PFC enabled. The command output shows that PFC is enabled for code points **011** (**fcoe** priority and queue) and **100** (**no-loss** priority and queue) for the **gd-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

### Verifying That the Output Queue Schedulers Have Been Created

**Purpose** Verify that the output queue schedulers have been created with the correct bandwidth parameters and priorities, mapped to the correct queues, and mapped to the correct drop profiles.

**Action** List the scheduler maps using the operational mode command **show class-of-service scheduler-map**:

```
user@switch> show class-of-service scheduler-map
```

```
Scheduler map: be-map, Index: 64023
```

```
Scheduler: be-sched, Forwarding class: best-effort, Index: 13005
Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
```

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 55387 | dp-be-low              |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 4369  | dp-be-high             |

Scheduler: be-sched, Forwarding class: be2, Index: 13005

Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 55387 | dp-be-low              |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 4369  | dp-be-high             |

Scheduler: nc-sched, Forwarding class: network-control, Index: 45740

Transmit rate: 5000000000 bps, Rate Limit: none, Buffer size: remainder,

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 44207 | dp-nc                  |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 1     | <default-drop-profile> |

Scheduler map: gd-map, Index: 61447

Scheduler: fcoe-sched, Forwarding class: fcoe, Index: 37289

Transmit rate: 2500000000 bps, Rate Limit: none, Buffer size: remainder,

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 44207 | <default-drop-profile> |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 1     | <default-drop-profile> |

Scheduler: nl-sched, Forwarding class: no-loss, Index: 29359

Transmit rate: 2000000000 bps, Rate Limit: none, Buffer size: remainder,

Buffer Limit: none, Priority: low

Excess Priority: unspecified

Shaping rate: 100 percent,

drop-profile-map-set-type: mark

Drop profiles:

| Loss priority | Protocol | Index | Name                   |
|---------------|----------|-------|------------------------|
| Low           | any      | 44207 | <default-drop-profile> |
| Medium high   | any      | 1     | <default-drop-profile> |
| High          | any      | 1     | <default-drop-profile> |

Scheduler map: hpc-map, Index: 56941



```
Scheduler: hpc-sched, Forwarding class: hpc, Index: 55900
Transmit rate: 2000000000 bps, Rate Limit: none, Buffer size: remainder,
Buffer Limit: none, Priority: low
Excess Priority: unspecified
Shaping rate: 100 percent,
drop-profile-map-set-type: mark
Drop profiles:
 Loss priority Protocol Index Name
 Low any 57716 dp-hpc
 Medium high any 1 <default-drop-profile>
 High any 1 <default-drop-profile>
```

**Meaning** The **show class-of-service scheduler-map** command lists all of the configured scheduler maps. For each scheduler map, the command output includes:

- The name of the scheduler map (**scheduler-map** field)
- The name of the scheduler (**scheduler** field)
- The forwarding classes mapped to the scheduler (**forwarding-class** field)
- The minimum guaranteed queue bandwidth (**transmit-rate** field)
- The scheduling priority (**priority** field)
- The maximum bandwidth in the priority group the queue can consume (**shaping-rate** field)
- The drop profile loss priority (**loss priority** field) for each drop profile name (**name** field)

The command output shows that:

- The scheduler map **be-map** has been created and has these properties:
  - There are two schedulers, **be-sched** and **nc-sched**.
  - The scheduler **be-sched** has two forwarding classes, **best-effort** and **be2**.
  - Scheduler **be-sched** forwarding classes **best-effort** and **be2** share a minimum guaranteed bandwidth of **3000000000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and use the drop profile **dp-be-low** for low loss-priority traffic, the default drop profile for medium-high loss-priority traffic, and the drop profile **dp-be-high** for high loss-priority traffic.
  - The scheduler **nc-sched** has one forwarding class, **network-control**.
  - The **network-control** forwarding class has a minimum guaranteed bandwidth of **5000000000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and uses the drop profile **dp-nc** for low loss-priority traffic and the default drop profile for medium-high and high loss priority traffic.
- The scheduler map **gd-map** has been created and has these properties:
  - There are two schedulers, **fcoe-sched** and **nl-sched**.
  - The scheduler **fcoe-sched** has one forwarding class, **fcoe**.

- The **fcoe** forwarding class has a minimum guaranteed bandwidth of **2500000000 bps**, and can consume a maximum of **100 percent** of the priority group bandwidth.
- The scheduler **nl-sched** has one forwarding class, **no-loss**.
- The **no-loss** forwarding class has a minimum guaranteed bandwidth of **2000000000 bps**, and can consume a maximum of **100 percent** of the priority group bandwidth.
- The scheduler map **hpc-map** has been created and has these properties:
  - There is one scheduler, **hpc-sched**.
  - The scheduler **hpc-sched** has one forwarding class, **hpc**.
  - The **hpc** forwarding class has a minimum guaranteed bandwidth of **2000000000 bps**, can consume a maximum of **100 percent** of the priority group bandwidth, and uses the drop profile **dp-hpc** for low loss-priority traffic and the default drop profile for medium-high and high loss-priority traffic.

---

### Verifying That the Drop Profiles Have Been Created

**Purpose** Verify that the drop profiles **dp-be-high**, **dp-be-low**, **dp-hpc**, and **dp-nc** have been created with the correct fill levels and drop probabilities.

**Action** List the drop profiles using the operational mode command **show configuration class-of-service drop-profiles**:

```
user@switch> show configuration class-of-service drop-profiles
dp-be-low {
 interpolate {
 fill-level [25 50];
 drop-probability [0 80];
 }
}
dp-be-high {
 interpolate {
 fill-level [10 40];
 drop-probability [0 100];
 }
}
dp-hpc {
 interpolate {
 fill-level [75 90];
 drop-probability [0 75];
 }
}
dp-nc {
 interpolate {
 fill-level [80 100];
 drop-probability [0 100];
 }
}
```

**Meaning** The **show configuration class-of-service drop-profiles** command lists the drop profiles and their properties. The command output shows that there are four drop profiles configured, **dp-be-high**, **dp-be-low**, **dp-hpc**, and **dp-nc**. The output also shows that:

- For **dp-be-low**, the drop start point (the first fill level) is when the queue is 25 percent filled, the drop end point (the second fill level) occurs when the queue is 50 percent filled, and the drop probability at the drop end point is 80 percent.
- For **dp-be-high**, the drop start point (the first fill level) is when the queue is 10 percent filled, the drop end point (the second fill level) occurs when the queue is 40 percent filled, and the drop probability at the drop end point is 100 percent.
- For **dp-hpc**, the drop start point (the first fill level) is when the queue is 75 percent filled, the drop end point (the second fill level) occurs when the queue is 90 percent filled, and the drop probability at the drop end point is 75 percent.
- For **dp-nc**, the drop start point (the first fill level) is when the queue is 80 percent filled, the drop end point (the second fill level) occurs when the queue is 100 percent filled, and the drop probability at the drop end point is 100 percent.

### Verifying That the Priority Group Output Schedulers (Traffic Control Profiles) Have Been Created

|                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | Verify that the traffic control profiles <b>be-tcp</b> , <b>gd-tcp</b> , and <b>hpc-tcp</b> have been created with the correct bandwidth parameters and scheduler mapping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Action</b>  | <p>List the traffic control profiles using the operational mode command <b>show class-of-service traffic-control-profile</b>:</p> <pre> user@switch&gt; show class-of-service traffic-control-profile Traffic control profile: be-tcp, Index: 40535   Shaping rate: 100 percent   Scheduler map: be-map   Guaranteed rate: 3500000000  Traffic control profile: gd-tcp, Index: 37959   Shaping rate: 100 percent   Scheduler map: gd-map   Guaranteed rate: 4500000000  Traffic control profile: hpc-tcp, Index: 47661   Shaping rate: 100 percent   Scheduler map: hpc-map   Guaranteed rate: 2000000000 </pre>                                    |
| <b>Meaning</b> | <p>The <b>show class-of-service traffic-control-profile</b> command lists all of the configured traffic control profiles. For each traffic control profile, the command output includes:</p> <ul style="list-style-type: none"> <li>• The name of the traffic control profile (<b>traffic-control-profile</b>)</li> <li>• The maximum port bandwidth the priority group can consume (<b>shaping-rate</b>)</li> <li>• The scheduler map associated with the traffic control profile (<b>scheduler-map</b>)</li> <li>• The minimum guaranteed priority group port bandwidth (<b>guaranteed-rate</b>)</li> </ul> <p>The command output shows that:</p> |

- The traffic control profile **be-tcp** can consume a maximum of **100 percent** of the port bandwidth, is associated with the scheduler map **be-map**, and has a minimum guaranteed bandwidth of **3500000000 bps**.
- The traffic control profile **gd-tcp** can consume a maximum of **100 percent** of the port bandwidth, is associated with the scheduler map **gd-map**, and has a minimum guaranteed bandwidth of **4500000000 bps**.
- The traffic control profile **hpc-tcp** can consume a maximum of **100 percent** of the port bandwidth, is associated with the scheduler map **hpc-map**, and has a minimum guaranteed bandwidth of **2000000000 bps**.

### Verifying the Interface Configuration

---

**Purpose** Verify that the classifier, the congestion notification profile, and the forwarding class sets are configured on interfaces **xe-0/0/20** and **xe-0/0/21**.

**Action** List the interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/20** and **show configuration class-of-service interfaces xe-0/0/21**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/20
```

```
forwarding-class-set {
 best-effort-gp {
 output-traffic-control-profile be-tcp;
 }
 guar-delivery-pg {
 output-traffic-control-profile gd-tcp;
 }
 hpc-pg {
 output-traffic-control-profile hpc-tcp;
 }
}
congestion-notification-profile gd_cnp;
unit 0 {
 classifiers {
 ieee-802.1 hscclassifier1;
 }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/21
```

```
forwarding-class-set {
 best-effort-gp {
 output-traffic-control-profile be-tcp;
 }
 guar-delivery-pg {
 output-traffic-control-profile gd-tcp;
 }
 hpc-pg {
 output-traffic-control-profile hpc-tcp;
 }
}
congestion-notification-profile gd_cnp;
unit 0 {
 classifiers {
 ieee-802.1 hscclassifier1;
 }
}
```

```
 }
}
```

**Meaning** The **show configuration class-of-service interfaces *interface-name*** command shows that each interface includes the forwarding class sets **best-effort-pg**, **guar-delivery-pg**, and **hpc-pg**, congestion notification profile **gd-cnp**, and the IEEE 802.1p classifier **hsclassifier1**.

**Related Documentation**

- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 404](#)
- [Benefits of Configuring CoS Hierarchical Port Scheduling](#)
- [Assigning CoS Components to Interfaces on page 429](#)
- [Example: Configuring WRED Drop Profiles on page 279](#)
- [Example: Configuring Drop Profile Maps on page 281](#)
- [Example: Configuring Forwarding Classes on page 283](#)
- [Example: Configuring Forwarding Class Sets on page 286](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring Queue Scheduling Priority on page 295](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 300](#)
- [Example: Configuring Maximum Output Bandwidth on page 305](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 21](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Understanding CoS Scheduling Behavior and Configuration Considerations on page 104](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 199](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 203](#)

## Example: Configuring CoS PFC for FCoE Traffic

Priority-based flow control (PFC, described in IEEE 802.1Qbb) is a link-level flow control mechanism that you apply at ingress interfaces. PFC enables you to divide traffic on one physical link into eight priorities. You can think of the eight priorities as eight “lanes” of traffic that correspond to queues (forwarding classes). Each priority is mapped to a 3-bit IEEE 802.1p CoS flag in the VLAN header.

You can selectively apply PFC to the traffic in any queue without pausing the traffic in other queues on the same link. You must apply PFC to FCoE traffic to ensure lossless transport.

To configure PFC on FCoE traffic, use the default FCoE forwarding-class-to-queue mapping and:

- Configure a classifier that associates the FCoE forwarding class with FCoE traffic.
- Configure a congestion notification profile to apply PFC to the FCoE traffic.
- Apply the classifier and the PFC configuration to ingress interfaces.
- Configure the bandwidth scheduling for the FCoE forwarding class output queue.
- Create a forwarding class set (priority group) that includes the FCoE forwarding class; this is required to configure enhanced transmission selection (ETS) and support data center bridging (DCB).
- Configure the bandwidth scheduling for the FCoE priority group.
- Apply the scheduling to the egress interfaces.



**NOTE:** If you are using Junos OS Release 12.2 or later, use the default forwarding classes for the lossless fcoe forwarding class. If you explicitly configure default lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best-effort) traffic and does *not* receive lossless treatment.

In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in explicit forwarding class configurations to configure a lossless forwarding class.

This example describes how to configure PFC for FCoE traffic:

- [Requirements on page 242](#)
- [Overview on page 242](#)
- [Configuration on page 244](#)
- [Verification on page 247](#)

## Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

## Overview

FCoE traffic requires PFC to ensure lossless packet transport. This example shows you how to:

- Assign FCoE traffic to the FCoE priority at the ingress.
- Create and apply CoS for the FCoE traffic using ETS (hierarchical port scheduling).

- Apply PFC to the FCoE traffic.
- Apply the configuration to ingress and egress interfaces.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

Each interface in this example is configured as both an ingress interface and an egress interface, so the classifier, congestion notification profile, and port scheduling are applied to all of the interfaces.

### Topology

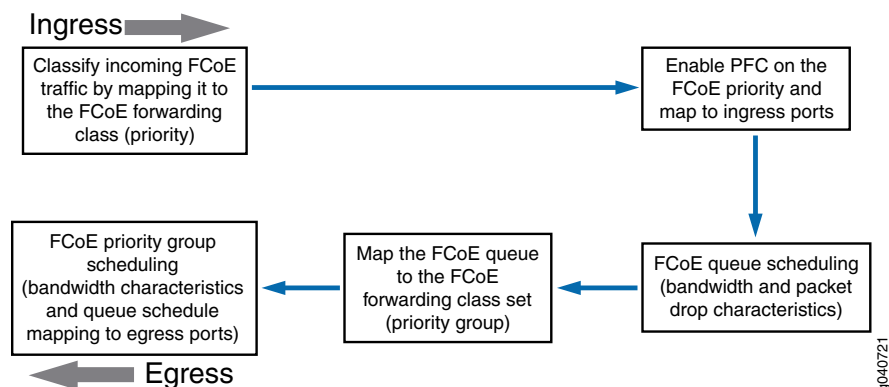
Table 86 on page 243 shows the configuration components for this example.

**Table 86: Components of the PFC for FCoE Traffic Configuration Topology**

| Component                                                                                                   | Settings                                                                                                                                             |
|-------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware                                                                                                    | QFX3500 switch                                                                                                                                       |
| Behavior aggregate classifier (maps the FCoE forwarding class to incoming packets by IEEE 802.1 code point) | Code point 011 to forwarding class <b>fcoe</b> and loss priority <b>low</b><br>Ingress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b> |
| PFC congestion notification profile                                                                         | <b>fcoe-cnp:</b><br>Code point 011<br>Ingress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b>                                          |
| FCoE queue scheduler                                                                                        | <b>fcoe-sched:</b><br>Minimum bandwidth <b>3g</b><br>Maximum bandwidth <b>100%</b><br>Priority <b>low</b>                                            |
| Forwarding class-to-scheduler mapping                                                                       | Scheduler map <b>fcoe-map:</b><br>Forwarding class <b>fcoe</b><br>Scheduler <b>fcoe-sched</b>                                                        |
| Forwarding class set (FCoE priority group)                                                                  | <b>fcoe-pg:</b><br>Forwarding class <b>fcoe</b><br>Egress interfaces: <b>xe-0/0/31, xe-0/0/32, xe-0/0/33, xe-0/0/34</b>                              |
| Traffic control profile                                                                                     | <b>fcoe-tcp:</b><br>Scheduler map <b>fcoe-map</b><br>Minimum bandwidth <b>3g</b><br>Maximum bandwidth <b>100%</b>                                    |

Figure 12 on page 244 shows a block diagram of the configuration components and the configuration flow of the CLI statements used in the example.

Figure 12: PFC for FCoE Traffic Configuration Components Block Diagram



## Configuration

### CLI Quick Configuration

To quickly configure PFC for FCoE traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the [edit] hierarchy level:

```
[edit class-of-service]
set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority low code-points 011
set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier
set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp
set schedulers fcoe-sched priority low transmit-rate 3g
set schedulers fcoe-sched shaping-rate percent 100
set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set forwarding-class-sets fcoe-pg class fcoe
set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate 3g
set traffic-control-profiles fcoe-tcp shaping-rate percent 100
set interfaces xe-0/0/31 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/32 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/33 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
set interfaces xe-0/0/34 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
```

### Step-by-Step Procedure

To configure the FCoE forwarding class (priority), ingress classifier, output queue scheduling, forwarding class set (priority group) and its output port scheduling, PFC application, and interfaces to set up PFC for FCoE traffic:

1. Configure a classifier to set the loss priority and IEEE 802.1 code point assigned to the FCoE forwarding class at the ingress:
 

```
[edit class-of-service]
user@switch# set classifiers ieee-802.1 fcoe-classifier forwarding-class fcoe loss-priority low code-points 011
```
2. Configure PFC on the FCoE queue by applying FCoE to the IEEE 802.1 code point 011:



```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```

3. Apply the PFC configuration to the ingress interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
user@switch# set interfaces xe-0/0/34 congestion-notification-profile fcoe-cnp
```

4. Assign the classifier to the ingress interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe-classifier
user@switch# set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 fcoe-classifier
```

5. Configure output scheduling for the FCoE queue:

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```

6. Map the FCoE forwarding class to the FCoE scheduler:

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```

7. Configure the forwarding class set for the FCoE traffic:

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```

8. Define the traffic control profile for the FCoE forwarding class set:

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```

9. Apply the FCoE forwarding class set and traffic control profile to the egress ports:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
user@switch# set interfaces xe-0/0/34 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
```

## Results

Display the results of the configuration (the system shows only the explicitly configured parameters; it does not show default parameters such as the **fcoe** lossless forwarding class):

```
user@switch> show configuration class-of-service
```

```
classifiers {
 ieee-802.1 fcoe-classifier {
 forwarding-class fcoe {
 loss-priority low code-points 011;
 }
 }
}
traffic-control-profiles {
 fcoe-tcp {
 scheduler-map fcoe-map;
 shaping-rate percent 100;
 guaranteed-rate 3000000000;
 }
}
forwarding-class-sets {
 fcoe-pg {
 class fcoe;
 }
}
congestion-notification-profile {
 fcoe-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 }
 }
 }
}
}
interfaces {
 xe-0/0/31 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
 }
 xe-0/0/32 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
 }
}
```

```

xe-0/0/33 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
}
xe-0/0/34 {
 congestion-notification-profile fcoe-cnp;
 forwarding-class-set {
 fcoe-pg {
 output-traffic-control-profile fcoe-tcp;
 }
 }
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
 }
}
}
scheduler-maps {
 fcoe-map {
 forwarding-class fcoe scheduler fcoe-sched;
 }
}
schedulers {
 fcoe-sched {
 transmit-rate 3000000000;
 shaping-rate percent 100;
 priority low;
 }
}
}

```



**TIP:** To quickly configure the interfaces, issue the `load merge terminal` command and then copy the hierarchy and paste it into the switch terminal window.

## Verification

To verify that the PFC configuration for FCoE traffic components has been created and is operating properly, perform these tasks:

- [Verifying That Priority-Based Flow Control Has Been Enabled on page 248](#)
- [Verifying the Ingress Interface PFC Configuration on page 248](#)

### Verifying That Priority-Based Flow Control Has Been Enabled

**Purpose** Verify that PFC is enabled on the FCoE queue to enable lossless transport.

**Action** List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Type: Input, Name: fcoe-cnp, Index: 51697
```

```
Cable Length: 100 m
```

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2500 |
| 100      | Disabled |      |
| 101      | Disabled |      |
| 110      | Disabled |      |
| 111      | Disabled |      |

```
Type: Output
```

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 000      | 0                   |
| 001      |                     |
| 010      | 1                   |
| 011      | 2                   |
| 100      | 3                   |
| 101      | 4                   |
| 110      | 5                   |
| 111      | 6                   |
|          | 7                   |

**Meaning** The **show class-of-service congestion-notification** operational command lists all of the congestion notification profiles and which IEEE 802.1p code points have PFC enabled. The command output shows that PFC is enabled on code point 011 for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

### Verifying the Ingress Interface PFC Configuration

**Purpose** Verify that the classifier **fcoe-classifier** and the congestion notification profile **fcoe-cnp** are configured on ingress interfaces **xe-0/0/31**, **xe-0/0/32**, **xe-0/0/33**, and **xe-0/0/34**.

**Action** List the ingress interfaces using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/31
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}

user@switch> show configuration class-of-service interfaces xe-0/0/32
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}

user@switch> show configuration class-of-service interfaces xe-0/0/33
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}

user@switch> show configuration class-of-service interfaces xe-0/0/34
congestion-notification-profile fcoe-cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe-classifier;
 }
}
```

**Meaning** The **show configuration class-of-service interfaces** commands list the congestion notification profile that is mapped to the interface (**fcoe-cnp**) and the IEEE 802.1p classifier associated with the interface (**fcoe-classifier**).

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
  - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
  - [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 21](#)
  - [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)

## Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG

Multichassis link aggregation groups (MC-LAGs) provide redundancy and load balancing between two QFX Series switches, multihoming support for client devices such as servers, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).



**NOTE:** This example uses Junos OS without support for the Enhanced Layer 2 Software (ELS) configuration style. If your switch runs software that supports ELS, see *Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG*.

You can use an MC-LAG to provide a redundant aggregation layer for Fiber Channel over Ethernet (FCoE) traffic in an *inverted-U* topology. To support lossless transport of FCoE traffic across an MC-LAG, you must configure the appropriate class of service (CoS) on both of the QFX Series switches with MC-LAG port members. The CoS configuration must be the same on both of the MC-LAG switches because an MC-LAG does not carry forwarding class and IEEE 802.1p priority information.



**NOTE:** This example describes how to configure CoS to provide lossless transport for FCoE traffic across an MC-LAG that connects two QFX Series switches. It also describes how to configure CoS on the FCoE transit switches that connect FCoE hosts to the QFX Series switches that form the MC-LAG.

This example does *not* describe how to configure the MC-LAG itself. For a detailed example of MC-LAG configuration, see *Example: Configuring Multichassis Link Aggregation*. However, this example includes a subset of MC-LAG configuration that only shows how to configure interface membership in the MC-LAG.

Ports that are part of an FCoE-FC gateway configuration (a virtual FCoE-FC gateway fabric) do not support MC-LAGs. Ports that are members of an MC-LAG act as FCoE passthrough transit switch ports.

QFX Series switches support MC-LAGs. QFabric system Node devices do not support MC-LAGs, and QFX3500 and QFX3600 Virtual Chassis switches do not support FCoE.

This topic describes:

- [Requirements on page 250](#)
- [Overview on page 251](#)
- [Configuration on page 256](#)
- [Verification on page 264](#)

## Requirements

This example uses the following hardware and software components:

- Two Juniper Networks QFX3500 Switches that form an MC-LAG for FCoE traffic.
- Two Juniper Networks QFX3500 Switches that provide FCoE server access in transit switch mode and that connect to the MC-LAG switches. These switches can be standalone QFX3500 switches or they can be Node devices in a QFabric system.

- FCoE servers (or other FCoE hosts) connected to the transit switches.
- Junos OS Release 12.2 or later for the QFX Series.

## Overview

FCoE traffic requires lossless transport. This example shows you how to:

- Configure CoS for FCoE traffic on the two QFX3500 switches that form the MC-LAG, including priority-based flow control (PFC) and enhanced transmission selection (ETS; hierarchical scheduling of resources for the FCoE forwarding class priority and for the forwarding class set priority group).



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- Configure CoS for FCoE on the two FCoE transit switches that connect FCoE hosts to the MC-LAG switches and enable FIP snooping on the FCoE VLAN at the FCoE transit switch access ports.
- Disable IGMP snooping on the FCoE VLAN.



**NOTE:** This is only necessary if IGMP snooping is enabled on the VLAN. Before Junos OS Release 13.2, IGMP snooping was enabled by default on VLANs. Beginning with Junos OS Release 13.2, IGMP snooping is enabled by default only on the default VLAN.

- Configure the appropriate port mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.

## Topology

QFX3500 switches that act as transit switches support MC-LAGs for FCoE traffic in an inverted-U network topology, as shown in [Figure 13 on page 252](#).

Figure 13: Supported Topology for an MC-LAG on an FCoE Transit Switch




Table 87 on page 252 shows the configuration components for this example.

Table 87: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology

| Component                                                                  | Settings                                                                                                                 |
|----------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| Hardware                                                                   | Four QFX3500 switches (two to form the MC-LAG as passthrough transit switches and two transit switches for FCoE access). |
| Forwarding class (all switches)                                            | Default <b>fcoe</b> forwarding class.                                                                                    |
| Classifier (forwarding class mapping of incoming traffic to IEEE priority) | Default IEEE 802.1p trusted classifier on all FCoE interfaces.                                                           |



Table 87: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology (*continued*)

| Component                                                | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LAGs and MC-LAG                                          | <p>S1—Ports xe-0/0/10 and x-0/0/11 are members of LAG <b>ae0</b>, which connects Switch S1 to Switch S2. Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG <b>ae1</b>. All ports are configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.</p> <p>S2—Ports xe-0/0/10 and x-0/0/11 are members of LAG <b>ae0</b>, which connects Switch S2 to Switch S1. Ports xe-0/0/20 and xe-0/0/21 are members of MC-LAG <b>ae1</b>. All ports are configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>.</p> <p><b>NOTE:</b> Ports xe-0/0/20 and xe-0/0/21 on Switches S1 and S2 are the members of the MC-LAG.</p> <p>TS1—Ports xe-0/0/25 and x-0/0/26 are members of LAG <b>ae1</b>, configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>. Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in <b>tagged-access</b> port mode, with an MTU of <b>2180</b>.</p> <p>TS2—Ports xe-0/0/25 and x-0/0/26 are members of LAG <b>ae1</b>, configured in <b>trunk</b> port mode, as <b>fcoe-trusted</b>, and with an MTU of <b>2180</b>. Ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are configured in <b>tagged-access</b> port mode, with an MTU of <b>2180</b>.</p> |
| FCoE queue scheduler (all switches)                      | <p><b>fcoe-sched:</b><br/>           Minimum bandwidth <b>3g</b><br/>           Maximum bandwidth <b>100%</b><br/>           Priority <b>low</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Forwarding class-to-scheduler mapping (all switches)     | <p>Scheduler map <b>fcoe-map</b>:<br/>           Forwarding class <b>fcoe</b><br/>           Scheduler <b>fcoe-sched</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Forwarding class set (FCoE priority group, all switches) | <p><b>fcoe-pg:</b><br/>           Forwarding class <b>fcoe</b></p> <p>Egress interfaces:</p> <ul style="list-style-type: none"> <li>• S1—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• S2—LAG <b>ae0</b> and MC-LAG <b>ae1</b></li> <li>• TS1—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> <li>• TS2—LAG <b>ae1</b>, interfaces <b>xe-0/0/30</b>, <b>xe-0/0/31</b>, <b>xe-0/0/32</b>, and <b>xe-0/0/33</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| Traffic control profile (all switches)                   | <p><b>fcoe-tcp:</b><br/>           Scheduler map <b>fcoe-map</b><br/>           Minimum bandwidth <b>3g</b><br/>           Maximum bandwidth <b>100%</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

**Table 87: Components of the CoS for FCoE Traffic Across an MC-LAG Configuration Topology** (*continued*)

| Component                                          | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PFC congestion notification profile (all switches) | <p><b>fcoe-cnp:</b><br/>Code point 011</p> <p>Ingress interfaces:</p> <ul style="list-style-type: none"> <li>• S1—LAG ae0 and MC-LAG ae1</li> <li>• S2—LAG ae0 and MC-LAG ae1</li> <li>• TS1—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33</li> <li>• TS2—LAG ae1, interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33</li> </ul>                                                                                                                                                                   |
| FCoE VLAN name and tag ID                          | <p>Name—<b>fcoe_vlan</b><br/>ID—100</p> <p>Include the FCoE VLAN on the interfaces that carry FCoE traffic on all four switches.</p> <p>Disable IGMP snooping on the interfaces that belong to the FCoE VLAN on all four switches.</p>                                                                                                                                                                                                                                                                                         |
| FIP snooping                                       | <p>Enable FIP snooping on Transit Switches TS1 and TS2 on the FCoE VLAN. Configure the LAG interfaces that connect to the MC-LAG switches as FCoE trusted interfaces so that they do not perform FIP snooping.</p> <p>This example enables VN2VN_Port FIP snooping on the FCoE transit switch interfaces connected to the FCoE servers. The example is equally valid with VN2VF_Port FIP snooping enabled on the transit switch access ports. The method of FIP snooping you enable depends on your network configuration.</p> |



**NOTE:** This example uses the default IEEE 802.1p trusted BA classifier, which is automatically applied to trunk mode and tagged access mode ports if you do not apply an explicitly configured classifier.

To configure CoS for FCoE traffic across an MC-LAG:

- Use the default FCoE forwarding class and forwarding-class-to-queue mapping (do not explicitly configure the FCoE forwarding class or output queue). The default FCoE forwarding class is **fcoe**, and the default output queue is queue 3.



**NOTE:** In Junos OS Release 12.2, traffic mapped to explicitly configured forwarding classes, even lossless forwarding classes such as **fcoe**, is treated as lossy (**best-effort**) traffic and does *not* receive lossless treatment. To receive lossless treatment in Release 12.2, traffic must use one of the default lossless forwarding classes (**fcoe** or **no-loss**).

In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in the explicit forwarding class configuration to configure a lossless forwarding class.

- Use the default trusted BA classifier, which maps incoming packets to forwarding classes by the IEEE 802.1p code point (CoS priority) of the packet. The trusted classifier is the default classifier for interfaces in trunk and tagged-access port modes. The default trusted classifier maps incoming packets with the IEEE 802.1p code point 3 (011) to the FCoE forwarding class. If you choose to configure the BA classifier instead of using the default classifier, you must ensure that FCoE traffic is classified into forwarding classes in exactly the same way on both MC-LAG switches. Using the default classifier ensures consistent classifier configuration on the MC-LAG ports.
- Configure a congestion notification profile that enables PFC on the FCoE code point (code point 011 in this example). The congestion notification profile configuration must be the same on both MC-LAG switches.
- Apply the congestion notification profile to the interfaces.
- Configure enhanced transmission selection (ETS, also known as hierarchical scheduling) on the interfaces to provide the bandwidth required for lossless FCoE transport. Configuring ETS includes configuring bandwidth scheduling for the FCoE forwarding class, a forwarding class set (priority group) that includes the FCoE forwarding class, and a traffic control profile to assign bandwidth to the forwarding class set that includes FCoE traffic.
- Apply the ETS scheduling to the interfaces.
- Configure the port mode, MTU, and FCoE trusted or untrusted state for each interface to support lossless FCoE transport.

In addition, this example describes how to enable FIP snooping on the Transit Switch TS1 and TS2 ports that are connected to the FCoE servers and how to disable IGMP snooping on the FCoE VLAN. To provide secure access, FIP snooping must be enabled on the FCoE access ports.

This example focuses on the CoS configuration to support lossless FCoE transport across an MC-LAG. This example does not describe how to configure the properties of MC-LAGs and LAGs, although it does show you how to configure the port characteristics required

to support lossless transport and how to assign interfaces to the MC-LAG and to the LAGs.

Before you configure CoS, configure:

- The MC-LAGs that connect Switches S1 and S2 to Switches TS1 and TS2. (*Example: Configuring Multichassis Link Aggregation* describes how to configure MC-LAGs.)
- The LAGs that connect the Transit Switches TS1 and TS2 to MC-LAG Switches S1 and S2. (*Configuring Link Aggregation* describes how to configure LAGs.)
- The LAG that connects Switch S1 to Switch S2.

## Configuration

To configure CoS for lossless FCoE transport across an MC-LAG, perform these tasks:

- [Configuring MC-LAG Switches S1 and S2 on page 258](#)
- [Configuring FCoE Transit Switches TS1 and TS2 on page 259](#)
- [Results on page 261](#)

### CLI Quick Configuration

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for MC-LAG Switch S1 and MC-LAG Switch S2 at the **[edit]** hierarchy level. The configurations on Switches S1 and S2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

### Switch S1 and Switch S2

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate
3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae0 congestion-notification-profile fcoe-cnp
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set protocols igmp-snooping vlan fcoe_vlan disable
set interfaces xe-0/0/10 ether-options 802.3ad ae0
set interfaces xe-0/0/11 ether-options 802.3ad ae0
set interfaces xe-0/0/20 ether-options 802.3ad ae1
set interfaces xe-0/0/21 ether-options 802.3ad ae1
set interfaces ae0 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces ae0 mtu 2180
set interfaces ae1 mtu 2180
set ethernet-switching-options secure-access-port interface ae0 fcoe-trusted
```

```
set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
```

To quickly configure CoS for lossless FCoE transport across an MC-LAG, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI for Transit Switch TS1 and Transit Switch TS2 at the **[edit]** hierarchy level. The configurations on Switches TS1 and TS2 are identical because the CoS configuration must be identical, and because this example uses the same ports on both switches.

## Switch TS1 and Switch TS2

```
set class-of-service schedulers fcoe-sched priority low transmit-rate 3g
set class-of-service schedulers fcoe-sched shaping-rate percent 100
set class-of-service scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
set class-of-service forwarding-class-sets fcoe-pg class fcoe
set class-of-service traffic-control-profiles fcoe-tcp scheduler-map fcoe-map guaranteed-rate
3g
set class-of-service traffic-control-profiles fcoe-tcp shaping-rate percent 100
set class-of-service interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg
output-traffic-control-profile fcoe-tcp
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces ae1 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/30 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
set vlans fcoe_vlan vlan-id 100
set protocols igmp-snooping vlan fcoe_vlan disable
set interfaces xe-0/0/25 ether-options 802.3ad ae1
set interfaces xe-0/0/26 ether-options 802.3ad ae1
set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan members fcoe_vlan
set interfaces xe-0/0/30 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces xe-0/0/31 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces xe-0/0/32 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces xe-0/0/33 unit 0 family ethernet-switching port-mode tagged-access vlan members
fcoe_vlan
set interfaces ae1 mtu 2180
set interfaces xe-0/0/30 mtu 2180
set interfaces xe-0/0/31 mtu 2180
set interfaces xe-0/0/32 mtu 2180
set interfaces xe-0/0/33 mtu 2180
set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
set ethernet-switching-options secure-access-port vlan fcoe_vlan examine-fip examine-vn2v2
beacon-period 90000
```

## Configuring MC-LAG Switches S1 and S2

### Step-by-Step Procedure

To configure CoS resource scheduling (ETS), PFC, the FCoE VLAN, and the LAG and MC-LAG interface membership and characteristics to support lossless FCoE transport across an MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point **011**, so you do not configure them):

1. Configure output scheduling for the FCoE queue:  

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```
2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):  

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```
3. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:  

```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
4. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:  

```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
5. Apply the FCoE forwarding class set and traffic control profile to the LAG and MC-LAG interfaces:  

```
[edit class-of-service]
user@switch# set interfaces ae0 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile
fcoe-tcp
```
6. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point **011**:  

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point
011 pfc
```
7. Apply the PFC configuration to the LAG and MC-LAG interfaces:  

```
[edit class-of-service]
user@switch# set interfaces ae0 congestion-notification-profile fcoe-cnp
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
```
8. Configure the VLAN for FCoE traffic (**fcoe\_vlan**):  

```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```
9. Disable IGMP snooping on the FCoE VLAN:  

```
[edit protocols]
user@switch# set igmp-snooping vlan fcoe_vlan disable
```

10. Add the member interfaces to the LAG between the two MC-LAG switches:
 

```
[edit interfaces]
user@switch# set xe-0/0/10 ether-options 802.3ad ae0
user@switch# set xe-0/0/11 ether-options 802.3ad ae0
```
11. Add the member interfaces to the MC-LAG:
 

```
[edit interfaces]
user@switch# set xe-0/0/20 ether-options 802.3ad ae1
user@switch# set xe-0/0/21 ether-options 802.3ad ae1
```
12. Configure the port mode as **trunk** and membership in the FCoE VLAN (**fcoe\_vlan**) for the LAG (**ae0**) and for the MC-LAG (**ae1**):
 

```
[edit interfaces]
user@switch# set interfaces ae0 unit 0 family ethernet-switching port-mode trunk vlan
members fcoe_vlan
user@switch# set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan
members fcoe_vlan
```
13. Set the MTU to **2180** for the LAG and MC-LAG interfaces. 2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes:
 

```
[edit interfaces]
user@switch# set ae0 mtu 2180
user@switch# set ae1 mtu 2180
```
14. Set the LAG and MC-LAG interfaces as FCoE trusted ports. Ports that connect to other switches should be trusted and should not perform FIP snooping:
 

```
[edit]
user@switch# set ethernet-switching-options secure-access-port interface ae0 fcoe-trusted
user@switch# set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
```

### Configuring FCoE Transit Switches TS1 and TS2

#### Step-by-Step Procedure

The CoS configuration on FCoE Transit Switches TS1 and TS2 is similar to the CoS configuration on MC-LAG Switches S1 and S2. However, the port configurations differ, and you must enable FIP snooping on the Switch TS1 and Switch TS2 FCoE access ports.

To configure resource scheduling (ETS), PFC, the FCoE VLAN, and the LAG interface membership and characteristics to support lossless FCoE transport across the MC-LAG (this example uses the default **fcoe** forwarding class and the default classifier to map incoming FCoE traffic to the FCoE IEEE 802.1p code point **011**, so you do not configure them):

1. Configure output scheduling for the FCoE queue:
 

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low transmit-rate 3g
user@switch# set schedulers fcoe-sched shaping-rate percent 100
```
2. Map the FCoE forwarding class to the FCoE scheduler (**fcoe-sched**):
 

```
[edit class-of-service]
user@switch# set scheduler-maps fcoe-map forwarding-class fcoe scheduler fcoe-sched
```
3. Configure the forwarding class set (**fcoe-pg**) for the FCoE traffic:

- ```
[edit class-of-service]
user@switch# set forwarding-class-sets fcoe-pg class fcoe
```
4. Define the traffic control profile (**fcoe-tcp**) to use on the FCoE forwarding class set:


```
[edit class-of-service]
user@switch# set traffic-control-profiles fcoe-tcp scheduler-map fcoe-map
guaranteed-rate 3g
user@switch# set traffic-control-profiles fcoe-tcp shaping-rate percent 100
```
 5. Apply the FCoE forwarding class set and traffic control profile to the LAG interface and to the FCoE access interfaces:


```
[edit class-of-service]
user@switch# set interfaces ae1 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/30 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/31 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/32 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
user@switch# set class-of-service interfaces xe-0/0/33 forwarding-class-set fcoe-pg output-traffic-control-profile fcoe-tcp
```
 6. Enable PFC on the FCoE priority by creating a congestion notification profile (**fcoe-cnp**) that applies FCoE to the IEEE 802.1 code point 011:


```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
```
 7. Apply the PFC configuration to the LAG interface and to the FCoE access interfaces:


```
[edit class-of-service]
user@switch# set interfaces ae1 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/30 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe-cnp
user@switch# set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe-cnp
```
 8. Configure the VLAN for FCoE traffic (**fcoe_vlan**):


```
[edit vlans]
user@switch# set fcoe_vlan vlan-id 100
```
 9. Disable IGMP snooping on the FCoE VLAN:


```
[edit protocols]
user@switch# set igmp-snooping vlan fcoe_vlan disable
```
 10. Add the member interfaces to the LAG:


```
[edit interfaces]
user@switch# set xe-0/0/25 ether-options 802.3ad ae1
user@switch# set xe-0/0/26 ether-options 802.3ad ae1
```
 11. On the LAG (**ae1**), configure the port mode as **trunk** and membership in the FCoE VLAN (**fcoe_vlan**):


```
[edit interfaces]
user@switch# set interfaces ae1 unit 0 family ethernet-switching port-mode trunk vlan
members fcoe_vlan
```

12. On the FCoE access interfaces (xe-0/0/30, xe-0/0/31, xe-0/0/32, xe-0/0/33), configure the port mode as **tagged-access** and membership in the FCoE VLAN (**fcoe_vlan**):

```
[edit interfaces]
user@switch# set interfaces xe-0/0/30 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/31 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/32 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
user@switch# set interfaces xe-0/0/33 unit 0 family ethernet-switching port-mode
tagged-access vlan members fcoe_vlan
```

13. Set the MTU to **2180** for the LAG and FCoE access interfaces. 2180 bytes is the minimum size required to handle FCoE packets because of the payload and header sizes; you can configure the MTU to a higher number of bytes if desired, but not less than 2180 bytes:

```
[edit interfaces]
user@switch# set ae1 mtu 2180
user@switch# set xe-0/0/30 mtu 2180
user@switch# set xe-0/0/31 mtu 2180
user@switch# set xe-0/0/32 mtu 2180
user@switch# set xe-0/0/33 mtu 2180
```

14. Set the LAG interface as an FCoE trusted port. Ports that connect to other switches should be trusted and should not perform FIP snooping:

```
[edit]
user@switch# set ethernet-switching-options secure-access-port interface ae1 fcoe-trusted
```



NOTE: Access ports xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33 are not configured as FCoE trusted ports. The access ports remain in the default state as untrusted ports because they connect directly to FCoE devices and must perform FIP snooping to ensure network security.

15. Enable FIP snooping on the FCoE VLAN to prevent unauthorized FCoE network access (this example uses VN2VN_Port FIP snooping; the example is equally valid if you use VN2VF_Port FIP snooping):

```
[edit]
user@switch# set ethernet-switching-options secure-access-port vlan fcoe_vlan
examine-fip examine-vn2vn beacon-period 90000
```

Results

Display the results of the CoS configuration on MC-LAG Switch S1 and on MC-LAG Switch S2 (the results on both switches are the same):

```
user@switch> show configuration class-of-service
```

```
traffic-control-profiles {
  fcoe-tcp {
    scheduler-map fcoe-map;
    shaping-rate percent 100;
    guaranteed-rate 3000000000;
  }
}
forwarding-class-sets {
  fcoe-pg {
    class fcoe;
  }
}
congestion-notification-profile {
  fcoe-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
      }
    }
  }
}
interfaces {
  ae0 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  ae1 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
}
scheduler-maps {
  fcoe-map {
    forwarding-class fcoe scheduler fcoe-sched;
  }
}
schedulers {
  fcoe-sched {
    transmit-rate 3000000000;
    shaping-rate percent 100;
    priority low;
  }
}
```



NOTE: The forwarding class and classifier configurations are not shown because the show command does not display default portions of the configuration.

For MC-LAG verification commands, see *Example: Configuring Multichassis Link Aggregation*.

Display the results of the CoS configuration on FCoE Transit Switch TS1 and on FCoE Transit Switch TS2 (the results on both transit switches are the same):

```
user@switch> show configuration class-of-service
traffic-control-profiles {
  fcoe-tcp {
    scheduler-map fcoe-map;
    shaping-rate percent 100;
    guaranteed-rate 30000000000;
  }
}
forwarding-class-sets {
  fcoe-pg {
    class fcoe;
  }
}
congestion-notification-profile {
  fcoe-cnp {
    input {
      ieee-802.1 {
        code-point 011 {
          pfc;
        }
      }
    }
  }
}
}
interfaces {
  xe-0/0/30 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  xe-0/0/31 {
    forwarding-class-set {
      fcoe-pg {
        output-traffic-control-profile fcoe-tcp;
      }
    }
    congestion-notification-profile fcoe-cnp;
  }
  xe-0/0/32 {
```

```
forwarding-class-set {
  fcoe-pg {
    output-traffic-control-profile fcoe-tcp;
  }
}
congestion-notification-profile fcoe-cnp;
}
xe-0/0/33 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
ae1 {
  forwarding-class-set {
    fcoe-pg {
      output-traffic-control-profile fcoe-tcp;
    }
  }
  congestion-notification-profile fcoe-cnp;
}
}
scheduler-maps {
  fcoe-map {
    forwarding-class fcoe scheduler fcoe-sched;
  }
}
schedulers {
  fcoe-sched {
    transmit-rate 3000000000;
    shaping-rate percent 100;
    priority low;
  }
}
```



NOTE: The forwarding class and classifier configurations are not shown because the show command does not display default portions of the configuration.

Verification

To verify that the CoS components and FIP snooping have been configured and are operating properly, perform these tasks. Because this example uses the default **fcoe** forwarding class and the default IEEE 802.1p trusted classifier, the verification of those configurations is not shown:

- [Verifying That the Output Queue Schedulers Have Been Created on page 265](#)
- [Verifying That the Priority Group Output Scheduler \(Traffic Control Profile\) Has Been Created on page 266](#)

- [Verifying That the Forwarding Class Set \(Priority Group\) Has Been Created on page 266](#)
- [Verifying That Priority-Based Flow Control Has Been Enabled on page 267](#)
- [Verifying That the Interface Class of Service Configuration Has Been Created on page 267](#)
- [Verifying That the Interfaces Are Correctly Configured on page 269](#)
- [Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces on page 272](#)
- [Verifying That the FIP Snooping Mode Is Correct on FCoE Transit Switches TS1 and TS2 on page 272](#)
- [Verifying That IGMP Snooping Is Disabled on the FCoE VLAN on page 273](#)

Verifying That the Output Queue Schedulers Have Been Created

Purpose Verify that the output queue scheduler for FCoE traffic has the correct bandwidth parameters and priorities, and is mapped to the correct forwarding class (output queue). Queue scheduler verification is the same on each of the four switches.

Action List the scheduler map using the operational mode command **show class-of-service scheduler-map fcoe-map**:

```
user@switch> show class-of-service scheduler-map fcoe-map
Scheduler map: fcoe-map, Index: 9023
```

```
Scheduler: fcoe-sched, Forwarding class: fcoe, Index: 37289
Transmit rate: 3000000000 bps, Rate Limit: none, Buffer size: remainder,
Buffer Limit: none, Priority: low
Excess Priority: unspecified
Shaping rate: 100 percent,
drop-profile-map-set-type: mark
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low            any       1      <default-drop-profile>
  Medium high    any       1      <default-drop-profile>
  High           any       1      <default-drop-profile>
```

Meaning The **show class-of-service scheduler-map fcoe-map** command lists the properties of the scheduler map **fcoe-map**. The command output includes:

- The name of the scheduler map (**fcoe-map**)
- The name of the scheduler (**fcoe-sched**)
- The forwarding classes mapped to the scheduler (**fcoe**)
- The minimum guaranteed queue bandwidth (transmit rate **3000000000 bps**)
- The scheduling priority (**low**)
- The maximum bandwidth in the priority group the queue can consume (shaping rate **100 percent**)
- The drop profile loss priority for each drop profile name. This example does not include drop profiles because you do not apply drop profiles to FCoE traffic.

Verifying That the Priority Group Output Scheduler (Traffic Control Profile) Has Been Created

- Purpose** Verify that the traffic control profile **fcoe-tcp** has been created with the correct bandwidth parameters and scheduler mapping. Priority group scheduler verification is the same on each of the four switches.
- Action** List the FCoE traffic control profile properties using the operational mode command **show class-of-service traffic-control-profile fcoe-tcp**:
- ```
user@switch> show class-of-service traffic-control-profile fcoe-tcp
Traffic control profile: fcoe-tcp, Index: 18303
 Shaping rate: 100 percent
 Scheduler map: fcoe-map
 Guaranteed rate: 3000000000
```
- Meaning** The **show class-of-service traffic-control-profile fcoe-tcp** command lists all of the configured traffic control profiles. For each traffic control profile, the command output includes:
- The name of the traffic control profile (**fcoe-tcp**)
  - The maximum port bandwidth the priority group can consume (shaping rate **100 percent**)
  - The scheduler map associated with the traffic control profile (**fcoe-map**)
  - The minimum guaranteed priority group port bandwidth (guaranteed rate **3000000000** in bps)

### Verifying That the Forwarding Class Set (Priority Group) Has Been Created

---

- Purpose** Verify that the FCoE priority group has been created and that the **fcoe** priority (forwarding class) belongs to the FCoE priority group. Forwarding class set verification is the same on each of the four switches.
- Action** List the forwarding class sets using the operational mode command **show class-of-service forwarding-class-set fcoe-pg**:
- ```
user@switch> show class-of-service forwarding-class-set fcoe-pg
Forwarding class set: fcoe-pg, Type: normal-type, Forwarding class set index:
31420
  Forwarding class                                Index
  fcoe                                              1
```
- Meaning** The **show class-of-service forwarding-class-set fcoe-pg** command lists all of the forwarding classes (priorities) that belong to the **fcoe-pg** priority group, and the internal index number of the priority group. The command output shows that the forwarding class set **fcoe-pg** includes the forwarding class **fcoe**.

Verifying That Priority-Based Flow Control Has Been Enabled

Purpose Verify that PFC is enabled on the FCoE code point. PFC verification is the same on each of the four switches.

Action List the FCoE congestion notification profile using the operational mode command **show class-of-service congestion-notification fcoe-cnp**:

```
user@switch> show class-of-service congestion-notification fcoe-cnp
Type: Input, Name: fcoe-cnp, Index: 6879
Cable Length: 100 m
  Priority  PFC          MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Enabled   2500
  100      Disabled
  101      Disabled
  110      Disabled
  111      Disabled
Type: Output
  Priority  Flow-Control-Queues
  000      0
  001      0
  010      1
  011      2
  100      3
  101      4
  110      5
  111      6
  111      7
```

Meaning The **show class-of-service congestion-notification fcoe-cnp** command lists all of the IEEE 802.1p code points in the congestion notification profile that have PFC enabled. The command output shows that PFC is enabled on code point **011 (fcoe queue)** for the **fcoe-cnp** congestion notification profile.

The command also shows the default cable length (100 meters), the default maximum receive unit (2500 bytes), and the default mapping of priorities to output queues because this example does not include configuring these options.

Verifying That the Interface Class of Service Configuration Has Been Created

Purpose Verify that the CoS properties of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches TS1 and TS2.

Action List the interface CoS configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
ae0 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}

ae1 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
```

List the interface CoS configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration class-of-service interfaces**:

```
user@switch> show configuration class-of-service interfaces
xe-0/0/30 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
xe-0/0/31 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
xe-0/0/32 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
xe-0/0/33 {
    forwarding-class-set {
        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}
ae1 {
    forwarding-class-set {
```



```

        fcoe-pg {
            output-traffic-control-profile fcoe-tcp;
        }
    }
    congestion-notification-profile fcoe-cnp;
}

```

Meaning The **show configuration class-of-service interfaces** command lists the class of service configuration for all interfaces. For each interface, the command output includes:

- The name of the interface (for example, **ae0** or **xe-0/0/30**)
- The name of the forwarding class set associated with the interface (**fcoe-pg**)
- The name of the traffic control profile associated with the interface (output traffic control profile, **fcoe-tcp**)
- The name of the congestion notification profile associated with the interface (**fcoe-cnp**)



NOTE: Interfaces that are members of a LAG are not shown individually. The LAG or MC-LAG CoS configuration is applied to all interfaces that are members of the LAG or MC-LAG. For example, the interface CoS configuration output on MC-LAG Switches S1 and S2 shows the LAG CoS configuration but does not show the CoS configuration of the member interfaces separately. The interface CoS configuration output on FCoE Transit Switches TS1 and TS2 shows the LAG CoS configuration but also shows the configuration for interfaces xe-0/0/30, xe-0/0/31, xe-0/0/32, and xe-0/0/33, which are not members of a LAG.

Verifying That the Interfaces Are Correctly Configured

Purpose Verify that the LAG membership, MTU, VLAN membership, and port mode of the interfaces are correct. The verification output on MC-LAG Switches S1 and S2 differs from the output on FCoE Transit Switches T1 and T2.

Action List the interface configuration on MC-LAG Switches S1 and S2 using the operational mode command **show configuration interfaces**:

```

user@switch> show configuration interfaces
xe-0/0/10 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/11 {
    ether-options {
        802.3ad ae0;
    }
}
xe-0/0/20 {
    ether-options {

```

```

        802.3ad ae1;
    }
}
xe-0/0/21 {
    ether-options {
        802.3ad ae1;
    }
}
ae0 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
ae1 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
}

```

List the interface configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration interfaces**:

```

user@switch> show configuration interfaces
xe-0/0/25 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/26 {
    ether-options {
        802.3ad ae1;
    }
}
xe-0/0/30 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
xe-0/0/31 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {

```

```

        port-mode tagged-access;
        vlan {
            members fcoe_vlan;
        }
    }
}
xe-0/0/32 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
xe-0/0/33 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode tagged-access;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}
ae1 {
    mtu 2180;
    unit 0 {
        family ethernet-switching {
            port-mode trunk;
            vlan {
                members fcoe_vlan;
            }
        }
    }
}

```

Meaning The **show configuration interfaces** command lists the configuration of each interface by interface name.

For each interface that is a member of a LAG, the command lists only the name of the LAG to which the interface belongs.

For each LAG interface and for each interface that is not a member of a LAG, the command output includes:

- The MTU (**2180**)
- The unit number of the interface (**0**)
- The port mode (**trunk** mode for interfaces that connect two switches, **tagged-access** mode for interfaces that connect to FCoE hosts)
- The name of the VLAN in which the interface is a member (**fcoe_vlan**)

Verifying That FIP Snooping Is Enabled on the FCoE VLAN on FCoE Transit Switches TS1 and TS2 Access Interfaces

Purpose Verify that FIP snooping is enabled on the FCoE VLAN access interfaces. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Transit Switch TS1 and TS2 FCoE access ports.

Action List the port security configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show configuration ethernet-switching-options secure-access-port**:

```
user@switch> show configuration ethernet-switching-options secure-access-port
interface ae1.0 {
    fcoe-trusted;
}
vlan fcoe_vlan {
    examine-fip {
        examine-vn2vn {
            beacon-period 90000;
        }
    }
}
```

Meaning The **show configuration ethernet-switching-options secure-access-port** command lists port security information, including whether a port is trusted. The command output shows that:

- LAG port **ae1.0**, which connects the FCoE transit switch to the MC-LAG switches, is configured as an FCoE trusted interface. FIP snooping is not performed on the member interfaces of the LAG (**xe-0/0/25** and **xe-0/0/26**).
- FIP snooping is enabled (**examine-fip**) on the FCoE VLAN (**fcoe_vlan**), the type of FIP snooping is VN2VN_Port FIP snooping (**examine-vn2vn**) and the beacon period is set to 90000 milliseconds. On Transit Switches TS1 and TS2, all interface members of the FCoE VLAN perform FIP snooping unless the interface is configured as FCoE trusted. On Transit Switches TS1 and TS2, interfaces **xe-0/0/30**, **xe-0/0/31**, **xe-0/0/32**, and **xe-0/0/33** perform FIP snooping because they are not configured as FCoE trusted. The interface members of LAG **ae1** (**xe-0/0/25** and **xe-0/0/26**) do not perform FIP snooping because the LAG is configured as FCoE trusted.

Verifying That the FIP Snooping Mode Is Correct on FCoE Transit Switches TS1 and TS2

Purpose Verify that the FIP snooping mode is correct on the FCoE VLAN. FIP snooping is enabled only on the FCoE access interfaces, so it is enabled only on FCoE Transit Switches TS1 and TS2. FIP snooping is not enabled on MC-LAG Switches S1 and S2 because FIP snooping is done at the Transit Switch TS1 and TS2 FCoE access ports.

Action List the FIP snooping configuration on FCoE Transit Switches TS1 and TS2 using the operational mode command **show fip snooping brief**:

```
user@switch> show fip snooping brief
VLAN: fcoe_vlan,      Mode: VN2VN Snooping
FC-MAP: 0e:fd:00
...
```



NOTE: The output has been truncated to show only the relevant information.

Meaning The **show fip snooping brief** command lists FIP snooping information, including the FIP snooping VLAN and the FIP snooping mode. The command output shows that:

- The VLAN on which FIP snooping is enabled is **fcoe_vlan**
- The FIP snooping mode is VN2VN_Port FIP snooping (**VN2VN Snooping**)

Verifying That IGMP Snooping Is Disabled on the FCoE VLAN

Purpose Verify that IGMP snooping is disabled on the FCoE VLAN on all four switches.

Action List the IGMP snooping protocol information on each of the four switches using the **show configuration protocols igmp-snooping** command:

```
user@switch> show configuration protocols igmp-snooping
vlan fcoe_vlan {
    disable;
}
```

Meaning The **show configuration protocols igmp-snooping** command lists the IGMP snooping configuration for the VLANs configured on the switch. The command output shows that IGMP snooping is disabled on the FCoE VLAN (**fcoe_vlan**).

Related Documentation

- *Example: Configuring Multichassis Link Aggregation*
- *Configuring Link Aggregation*
- [Example: Configuring CoS PFC for FCoE Traffic on page 241](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- *Understanding Multichassis Link Aggregation*
- *Understanding MC-LAGs on an FCoE Transit Switch*

Example: Configuring Unicast Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. Classifiers associate packets with a forwarding class and loss priority and assign packets

to output queues based on the associated forwarding class. You apply classifiers to ingress interfaces.

- [Requirements on page 274](#)
- [Overview on page 274](#)
- [Configuring Unicast Classifiers on page 275](#)
- [Verification on page 275](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

Junos OS supports three general types of classifiers:

- Behavior aggregate or CoS value traffic classifiers—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP or DSCP IPv6) value, IEEE 802.1p value, or MPLS EXP value. (EXP classifiers can be applied only to **family mpls** interfaces.)
- Fixed classifiers. Fixed classifiers classify all ingress traffic on a physical interface into one forwarding class, regardless of the CoS bits in the packet header.
- Multifield traffic classifiers—Examine multiple fields in the packet, such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.



.....

NOTE: You must assign unicast traffic and multidestination (multicast, broadcast, and destination lookup fail) traffic to different classifiers. One classifier cannot include both unicast and multidestination forwarding classes. A unicast classifier can include only forwarding classes for unicast traffic.

.....

This example describes how to configure a BA classifier called **ba-ucast-classifier** as the default IEEE 802.1 map and apply it to ingress interface **xe-0/0/10**. The BA classifier assigns loss priorities, as shown in [Table 88 on page 275](#), to incoming packets in the four forwarding classes.

You can use the same procedure to set multifield classifiers (except that you use firewall filter rules).

Table 88: ba-ucast-classifier Loss Priority Assignments

Unicast Forwarding Class	For CoS Traffic Type	ba-ucast-classifier Loss Priority to IEEE 802.1p Code Point Mapping	Packet Drop Attribute
be	Best-effort traffic	Low loss priority code point: 000	drop
fcoe	Guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic	Low loss priority code point: 011	no-loss
no-loss	Guaranteed delivery for TCP traffic	Low loss priority code point: 100	no-loss
nc	Network-control traffic	Low loss priority code point: 110	drop

Configuring Unicast Classifiers

To configure a unicast IEEE 802.1 BA classifier named **ba-ucast-classifier** as the default IEEE 802.1 map:

- Associate code point **000** with forwarding class **be** and loss priority **low**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-ucast-classifier import default forwarding-class be
loss-priority low code-points 000
```
- Associate code point **011** with forwarding class **fcoe** and loss priority **low**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-ucast-classifier forwarding-class fcoe loss-priority low
code-points 011
```
- Associate code point **100** with forwarding class **no-loss** and loss priority **low**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-ucast-classifier forwarding-class no-loss loss-priority low
code-points 100
```
- Associate code point **110** with forwarding class **nc** and loss priority **low**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-ucast-classifier forwarding-class nc loss-priority low
code-points 110
```
- Apply the unicast classifier to ingress interface **xe-0/0/10**:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/10 unit 0 classifiers ieee-802.1 ba-ucast-classifier
```

Verification

To verify the unicast classifier configuration, perform these tasks:

- [Verifying the Unicast Classifier Configuration on page 276](#)
- [Verifying the Ingress Interface Configuration on page 276](#)

Verifying the Unicast Classifier Configuration

Purpose Verify that you configured the unicast classifier with the correct forwarding classes, loss priorities, and code points.

Action List the classifier configuration using the operational mode command **show configuration class-of-service classifiers ieee-802.1 ba-ucast-classifier**:

```
user@switch> show configuration class-of-service classifiers ieee-802.1 ba-ucast-classifier
  forwarding-class be {
    loss-priority low code-points 000;
  }
  forwarding-class fcoe {
    loss-priority low code-points 011;
  }
  forwarding-class no-loss {
    loss-priority low code-points 100;
  }
  forwarding-class nc
    loss-priority low code-points 110;
  }
```

Verifying the Ingress Interface Configuration

Purpose Verify that the unicast classifier **ba-ucast-classifier** is attached to ingress interface **xe-0/0/10**.

Action List the ingress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/10**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/10
congestion-notification-profile fcoe-cnp;
unit 0 {
  classifiers {
    ieee-802.1 ba-ucast-classifier;
  }
}
```

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
 - [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 277](#)
 - [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 404](#)
 - [Configuring a Global MPLS EXP Classifier](#)
 - [Configuring Rewrite Rules for MPLS EXP Classifiers](#)
 - [Monitoring CoS Classifiers on page 545](#)
 - [Understanding CoS Classifiers on page 51](#)
 - [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57](#)

Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. Classifiers associate packets with a forwarding class and loss priority and assign packets to output queues based on the associated forwarding class.

- [Requirements on page 277](#)
- [Overview on page 277](#)
- [Configuring Multidestination Classifiers on page 278](#)
- [Verification on page 278](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

Junos OS supports three general types of classifiers:

- Behavior aggregate or CoS value traffic classifiers—Examine the CoS value in the packet header. The value in this single field determines the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the Differentiated Services code point (DSCP) value or IEEE 802.1p value.



NOTE: DSCP IPv6 multidestination classifiers are not supported. IPv6 multidestination traffic uses the DSCP multidestination classifier.

- Fixed classifiers. Fixed classifiers classify all ingress traffic on a physical interface into one forwarding class, regardless of the CoS bits in the packet header.
- Multifield traffic classifiers—Examine multiple fields in the packet such as source and destination addresses and source and destination port numbers of the packet. With multifield classifiers, you set the forwarding class and loss priority of a packet based on firewall filter rules.

Multidestination classifiers apply to all of the switch interfaces and handle multicast, broadcast, and destination lookup fail (DLF) traffic. You cannot apply a multidestination classifier to a single interface or to a range of interfaces.



NOTE: You must assign unicast traffic and multicast traffic to different classifiers. One classifier cannot include both unicast and multicast forwarding classes. A multidestination classifier can include only forwarding classes for multicast traffic.

The following example describes how to configure a BA classifier called **ba-mcast-classifier**, which is applied to all of the switch interfaces. The BA classifier assigns loss priorities, as shown in [Table 89 on page 278](#), to incoming packets in the multdestination forwarding class.

You can use the same procedure to set multifield classifiers (except that you use firewall filter rules).

Table 89: BA-mcast-classifier Loss Priority Assignments

Multicast Forwarding Class	For CoS Traffic Type	ba-mcast-classifier Assignment
mcast	Best-effort multicast traffic	Low loss priority code point: 000

Configuring Multidestination Classifiers

To configure a multicast IEEE 802.1 BA classifier named **ba-mcast-classifier**:

1. Associate code point 000 with forwarding class **mcast** and loss priority **low**:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 ba-mcast-classifier forwarding-class mcast loss-priority low
code-points 000
```

2. Configure the classifier as a multdestination classifier:

```
[edit class-of-service]
user@switch# set multi-destination classifiers ieee-802.1 ba-mcast-classifier
```

Verification

To verify the multdestination classifier configuration, perform these tasks:

- [Verifying the IEEE 802.1 Multidestination Classifier on page 278](#)
- [Verifying the Multidestination Classifier Configuration on page 278](#)

Verifying the IEEE 802.1 Multidestination Classifier

Purpose Verify that the classifier **ba-mcast-classifier** is configured as the IEEE 802.1 multdestination classifier:

Action Verify the results of the classifier configuration using the operational mode command **show configuration class-of-service multi-destination classifiers ieee-802.1**:

```
user@switch> show configuration class-of-service multi-destination classifiers ieee-802.1
ba-mcast-classifier;
```

Verifying the Multidestination Classifier Configuration

Purpose Verify that you configured the multdestination classifier with the correct forwarding classes, loss priorities, and code points.

Action List the classifier configuration using the operational mode command **show configuration class-of-service classifiers ieee-802.1 ba-mcast-classifier**:

```
user@switch> show configuration class-of-service classifiers ieee-802.1 ba-mcast-classifier
  forwarding-class mcast {
    loss-priority low code-points 000;
  }
```

- Related Documentation**
- [Example: Configuring Unicast Classifiers on page 273](#)
 - [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 406](#)
 - [Monitoring CoS Classifiers on page 545](#)
 - [Understanding CoS Classifiers on page 51](#)
 - [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57](#)

Example: Configuring WRED Drop Profiles

You can configure an interpolated weighted random early detection (WRED) profile to control packet drop characteristics for different traffic loss priorities.



NOTE: You cannot enable WRED on multidestination (multicast) queues. You can enable WRED only on unicast queues.

Also, do not enable WRED on lossless traffic flows. Use priority-based flow control (PFC) to prevent packet loss on lossless forwarding classes.

- [Requirements on page 279](#)
- [Overview on page 279](#)
- [Configuring a Drop Profile on page 281](#)
- [Verification on page 281](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

You associate a WRED profile with a loss priority in a scheduler. When you attach the scheduler to a forwarding class (queue), you apply the interpolated drop profile to traffic of the specified loss priority in that queue. *Interpolated* means that the switch creates a smooth drop curve from a drop start point to a drop end point, with a maximum drop rate that is reached at the drop end point:

- Drop start point—Percentage of average queue fill level when the WRED algorithm starts to drop packets. Before the drop start point, no packets are scheduled to drop.
- Drop end point—Average queue fill level at which all subsequently arriving packets are dropped. When the queue fill levels falls below the drop end point, packets begin to be forwarded again. (At the drop end point, the packet drop probability becomes 100 percent.)
- Maximum drop rate—Drop probability when the average queue fill level reaches the drop end point.


You set the drop start point and the drop end point by specifying two queue fill level percentage values. The first value is the drop start point and the second value is the drop end point.

You set the maximum drop rate by specifying two drop probability percentage values. The first value is always zero (0), which is the minimum drop rate, the probability of dropping a packet at the drop start point. The second value is the maximum drop rate at the drop end point.

The drop rate is zero until the queue fill level reaches the drop start point. As the queue continues to fill, packets drop in smooth linear curve until the queue reaches the drop end point, when packets drop at the maximum drop rate. If the queue fills beyond the drop end point, all packets that match the drop profile are dropped.

Figure 14 on page 280 shows the graph for a drop profile with a drop start point of 30 percent, a drop end point of 50 percent, and a maximum drop rate of 80 percent.

Figure 14: WRED Drop Profile Packet Drop Example



The graph shows that when the queue fill level is less than 30 percent, the packet drop rate is zero. When the queue fill level reaches 30 percent, packets begin to drop. As the queue fills, a higher percentage of packets drop. When the queue fill level reaches 50 percent, the packet drop rate has climbed to 80 percent. When the queue fill level exceeds 50 percent, all packets drop.

This example describes how to configure the drop profile shown in Figure 14 on page 280. The drop profile will have:

- The name **be-dp1**
- 30 percent for the drop start point (first **fill-level** setting)
- 50 percent for the drop end point (second **fill-level** setting)
- 0 percent for the minimum drop rate (first **drop-probability** setting)
- 80 percent for the maximum drop rate (second **drop-probability** setting)

You apply a drop profile by configuring a drop profile map that maps the drop profile to a packet loss priority and associates the drop profile and packet loss priority with a scheduler. When you associate the scheduler with a forwarding class (queue), the switch applies the drop profile to the packets in the forwarding class that have a matching packet loss priority.

Configuring a Drop Profile

1. Set the drop start point at **30** percent, the drop end point at **50** percent, the minimum drop rate at **0** percent, and the maximum drop rate at **80** percent for the drop profile **be-dp1**:

```
[edit class-of-service]
user@switch# set drop-profile be-dp1 interpolate fill-level 30 fill-level 50 drop-probability
0 drop-probability 80
```

Verification

Verifying the Drop Profile Configuration

Purpose	Verify that you configured the drop profile be-dp1 with the correct drop start and end points and with the correct drop rates.
Action	<p>Verify the results of the drop profile configuration using the operational mode command show configuration class-of-service drop-profiles be-dp1:</p> <pre>user@switch> show configuration class-of-service drop-profiles be-dp1 interpolate { fill-level [30 50]; drop-probability [0 80]; }</pre>
Related Documentation	<ul style="list-style-type: none"> • Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220 • Example: Configuring Queue Schedulers on page 289 • Example: Configuring Drop Profile Maps on page 281 • Configuring CoS WRED Drop Profiles on page 407 • Understanding CoS WRED Drop Profiles on page 149

Example: Configuring Drop Profile Maps

A drop-profile map associates a WRED profile for traffic of a specified loss priority with a scheduler. When you use a scheduler map to map a scheduler to a forwarding class,

the drop profile map associated with the scheduler applies the specified WRED profile to traffic in the forwarding class that matches the specified loss priority.

- [Requirements on page 282](#)
- [Overview on page 282](#)
- [Configuring a Drop Profile Map on page 282](#)
- [Verification on page 283](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

Drop profile maps enable you to configure different drop profiles for traffic of different loss priorities within the same scheduler. You can associate different drop profiles with low-priority, medium-high priority, and high-priority traffic within a single scheduler, and then map that scheduler to a forwarding class. This applies the appropriate drop profile to traffic of each loss priority in a forwarding class. Drop profile maps apply to all traffic protocols.

The following example describes how to configure a drop profile map for a scheduler named **mylan** that includes:

- A drop profile called **lp-profile** for low-priority traffic
- A drop profile called **mh-profile** for medium-high priority traffic
- A drop profile called **h-profile** for high-priority traffic

You apply the drop profiles in the drop profile map to a forwarding class by associating the scheduler **mylan** with a forwarding class in a scheduler map.

Configuring a Drop Profile Map

1. Configure the drop profile for low-priority traffic:

```
[edit class-of-service]
user@switch# set schedulers mylan drop-profile-map loss-priority low protocol any
drop-profile lp-profile
```

2. Configure the drop profile for medium-high priority traffic:

```
[edit class-of-service]
user@switch# set schedulers mylan drop-profile-map loss-priority medium-high protocol
any drop-profile mh-profile
```

3. Configure the drop profile for high-priority traffic:

```
[edit class-of-service]
user@switch# set schedulers mylan drop-profile-map loss-priority high protocol any
drop-profile h-profile
```

Verification

Verifying the Drop Profile Map Configuration

Purpose Verify that you configured the drop profile map for the scheduler **mylan** with the correct loss priorities and drop profiles.

Action Verify the results of the drop profile map configuration using the operational mode command **show configuration class-of-service schedulers mylan**:

```
user@switch> show configuration class-of-service schedulers mylan
transmit-rate 3g;
shaping-rate percent 100;
priority low;
drop-profile-map loss-priority low protocol any drop-profile lp-profile;
drop-profile-map loss-priority medium-high protocol any drop-profile mh-profile;
drop-profile-map loss-priority high protocol any drop-profile h-profile;
```



NOTE: This example does not include configuring scheduler bandwidth and priority. This information (transmit rate, shaping rate, and priority) is shown for completeness.

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring WRED Drop Profiles on page 279](#)
- [Configuring CoS Drop Profile Maps on page 408](#)
- [Understanding CoS WRED Drop Profiles on page 149](#)

Example: Configuring Forwarding Classes

Forwarding classes allow you to group packets for transmission. You assign packets to unicast or multidestination (multicast, broadcast, and destination lookup fail) output queues based on forwarding classes.

- [Requirements on page 283](#)
- [Overview on page 284](#)
- [Configuring Forwarding Classes on page 285](#)
- [Verification on page 286](#)

Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch

- Junos OS Release 11.1 or later for the QFX Series

Overview

The switch supports a total of 12 forwarding classes. In order to forward traffic, you must map (assign) the forwarding classes to unicast or multidestination output queues. The switch has 12 queues. Queues 0 through 7 are for unicast traffic, and queues 8 through 11 are for multidestination traffic. The switch supports up to two lossless forwarding classes.

By default, four categories of unicast forwarding classes and one multidestination forwarding class are defined. You can define the remaining seven forwarding classes and configure them as unicast or multidestination by mapping them to unicast or multidestination queues. The type of queue, unicast or multidestination, determines the type of forwarding class.

The four default unicast forwarding classes are:

- **be**—Best-effort traffic
- **fcoe**—Guaranteed delivery for Fibre Channel over Ethernet traffic
- **no-loss**—Guaranteed delivery for TCP no-loss traffic
- **nc**—Network control traffic

The default multidestination forwarding class is:

- **mcast**—Multidestination traffic

Map forwarding classes to queues using the **class** statement, which enables you to configure up to 12 forwarding classes. You can map more than one forwarding class to a single queue, but all forwarding classes mapped to a particular queue must be of the same type, either unicast or multicast. You cannot mix unicast and multicast forwarding classes on the same queue. The statement format is:

```
[edit class-of-service forwarding-classes]
user@switch# class class-name queue-num queue-number;
```



NOTE: If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless **fcoe** and **no-loss** forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (**best-effort**) traffic and does *not* receive lossless treatment.

In Junos OS Release 12.3 and later, you can include the *no-loss* packet drop attribute in explicit forwarding class configurations to configure a lossless forwarding class.



NOTE: Junos OS Release 11.3R1 and earlier supported an alternate method of mapping forwarding classes to queues that allowed you to map only one forwarding class to a queue using the statement:

```
[edit class-of-service forwarding-classes]
user@switch# queue queue-number class-name
```

The `queue` statement has been deprecated and is no longer valid in Junos OS Release 11.3R2 and later. If you have a configuration that uses the `queue` statement to map forwarding classes to queues, edit the configuration to replace the `queue` statement with the `class` statement.



NOTE: The QFX Series uses hierarchical scheduling to control output queue forwarding. When you define a forwarding class that will carry traffic on the switch (the behavior aggregate classifier has a forwarding class and you expect traffic for the forwarding class), you must also define a scheduling policy for the forwarding class. Defining a scheduling policy means:

- Mapping a scheduler to the forwarding class in a scheduler map
- Including the forwarding class in a forwarding class set
- Associating the scheduler map with a traffic control profile
- Attaching the traffic control profile to a forwarding class set and an interface

Table 90 on page 285 shows the configuration forwarding-class-to-queue mapping for this example:

Table 90: Forwarding-Class-to-Queue Example Configuration

Forwarding Class	Queue
best-effort	0
nc	7
mcast	8

Configuring Forwarding Classes

To configure CoS forwarding classes, map the forwarding classes to queues:

1. Map the **best-effort** forwarding class to queue 0:

```
[edit class-of-service forwarding-classes]
user@switch# set class best-effort queue-num 0
```

2. Map the **nc** forwarding class to queue 7:

```
[edit class-of-service forwarding-classes]
```

```
user@switch# set class nc queue-num 7
```

3. Map the **mcast-be** forwarding class to queue 8:

```
[edit class-of-service forwarding-classes]  
user@switch# set class mcast-be queue-num 8
```

Verification

Verifying the Forwarding-Class-to-Queue Mapping

Purpose	Verify the forwarding-class-to-queue mapping. (The system shows only the explicitly configured forwarding classes; it does not show default forwarding classes such as fcoe and no-loss .)
Action	Verify the results of the forwarding class configuration using the operational mode command show configuration class-of-service forwarding-classes : user@switch> show configuration class-of-service forwarding-classes class best-effort queue-num 0; class network-control queue-num 7; class mcast queue-num 8;
Related Documentation	<ul style="list-style-type: none">• Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220• Defining CoS Forwarding Classes on page 408• Monitoring CoS Forwarding Classes on page 546• Overview of CoS Changes Introduced in Junos OS Release 11.3 on page 13• Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 21• Understanding CoS Forwarding Classes on page 68

Example: Configuring Forwarding Class Sets

A forwarding class set (fc-set) is a priority group for enhanced transmission selection (ETS) traffic control. Each fc-set consists of one or more forwarding classes (output queues).

ETS enables you to configure link resources (bandwidth and bandwidth sharing characteristics) for a priority group, and then allocate the priority group's resources among the forwarding classes that belong to the priority group. This is called two-tier, or hierarchical, scheduling. Traffic control profiles control the scheduling for the priority group, and schedulers control the scheduling for individual forwarding classes.

- [Requirements on page 287](#)
- [Overview on page 287](#)
- [Configuring Forwarding Class Sets on page 288](#)
- [Verification on page 288](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

You can configure up to three unicast fc-sets and one multicast fc-set. A common way to configure unicast priority groups is to configure separate fc-sets for local area network (LAN) traffic, storage area network (SAN) traffic, and high-performance computing (HPC) traffic, and then assign the appropriate forwarding classes to each fc-set.



NOTE: If you configure strict-high priority queues, you must create an fc-set that is dedicated only to strict-high priority traffic. Only one fc-set can contain strict-high priority queues. Queues that are not strict-high priority cannot belong to the same fc-set as strict-high priority queues. The multidestination fc-set cannot contain strict-high priority queues.

To apply ETS, you map one or more fc-sets to a physical egress port. You can map up to three forwarding class sets to each port. When you map an fc-set to a port, the port uses hierarchical scheduling to allocate port resources to the priority group (fc-set) and to allocate the priority group resources to the queues (forwarding classes) that belong to the priority group.

This example describes how to:

- Configure three fc-sets called **lan-pg**, **san-pg**, and **hpc-pg**.
- Assign forwarding classes to each of the fc-sets.
- Apply the fc-sets and their output traffic control profiles to an egress interface.

This example does not describe how to configure the forwarding classes assigned to the fc-sets or how to configure traffic control profiles. [Table 91 on page 287](#) shows the configuration components for this example:

Table 91: Components of the Forwarding Class Sets Configuration Example

Component	Settings
Hardware	QFX3500 switch
LAN traffic priority group	Forwarding class set: lan-pg Forwarding classes: best-effort-1 , best-effort-2
SAN traffic priority group	Forwarding class set: san-pg Forwarding classes: fcoe , fcoe-2

Table 91: Components of the Forwarding Class Sets Configuration Example (*continued*)

Component	Settings
HPC traffic priority group	Forwarding class set: hpc-pg Forwarding classes: nc, high-perf
Egress interface	xe-0/0/7

Configuring Forwarding Class Sets

1. Define the **lan-pg** priority group (fc-set) and assign to it the forwarding classes **best-effort-1** and **best-effort-2**:


```
[edit class-of-service]
user@switch# set forwarding-class-sets lan-pg class best-effort-1
user@switch# set forwarding-class-sets lan-pg class best-effort-2
```
2. Define the **san-pg** priority group and assign to it the forwarding classes **fcoe** and **fcoe-2**:


```
[edit class-of-service]
user@switch# set forwarding-class-sets san-pg class fcoe
user@switch# set forwarding-class-sets san-pg class fcoe-2
```
3. Define the **hpc-pg** priority group and assign to it the forwarding classes **nc** and **high-perf**:


```
[edit class-of-service]
user@switch# set forwarding-class-sets hpc-pg class nc
user@switch# set forwarding-class-sets hpc-pg class high-perf
```
4. Map the three forwarding class sets to an interface (the output traffic control profiles associated with the forwarding class sets determine the class of service scheduling for the priority groups):


```
[edit class-of-service]
user@switch# set interfaces xe-0/0/7 forwarding-class-set lan-pg
output-traffic-control-profile lan-tcp
user@switch# set interfaces xe-0/0/7 forwarding-class-set san-pg
output-traffic-control-profile san-tcp
user@switch# set interfaces xe-0/0/7 forwarding-class-set hpc-pg
output-traffic-control-profile hpc-tcp
```

Verification

To verify the priority group configuration, perform these tasks:

- [Verifying Forwarding Class Set Membership on page 288](#)
- [Verifying the Egress Interface Configuration on page 289](#)

Verifying Forwarding Class Set Membership

Purpose Verify that you configured the **lan-pg**, **san-pg**, and **hpc-pg** priority groups with the correct forwarding classes.

Action List the forwarding class set member configuration using the operational mode command **show configuration class-of-service forwarding-class-sets**:

```
user@switch> show configuration class-of-service forwarding-class-sets
lan-pg {
    class best-effort-1;
    class best-effort-2;
}
san-pg {
    class fcoe;
    class fcoe-2;
}
hpc-pg {
    class high-perf;
    class nc;
}
```

Verifying the Egress Interface Configuration

Purpose Verify that egress interface **xe-0/0/7** is associated with the **lan-pg**, **san-pg**, and **hpc-pg** priority groups and with the correct output traffic control profiles.

Action Display the egress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
    lan-pg {
        output-traffic-control-profile lan-tcp;
    }
    san-pg {
        output-traffic-control-profile san-tcp;
    }
    hpc-pg {
        output-traffic-control-profile hpc-tcp;
    }
}
```

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
 - [Example: Configuring Queue Schedulers on page 289](#)
 - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
 - [Defining CoS Forwarding Class Sets on page 410](#)
 - [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 73](#)

Example: Configuring Queue Schedulers

Schedulers define the CoS properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the priority of the queue, whether

explicit congestion notification (ECN) is enabled on the queue, and the WRED packet drop profiles associated with the queue.

- [Requirements on page 290](#)
- [Overview on page 290](#)
- [Configuring a CoS Scheduler on page 293](#)
- [Verification on page 293](#)

Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

Scheduler parameters define the following characteristics for the queues mapped to the scheduler:

- **transmit-rate**—Minimum bandwidth, also known as the committed information rate (CIR). Each queue mapped to the scheduler receives a minimum of either the configured amount of absolute bandwidth or the configured percentage of bandwidth. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue. You cannot configure a transmit rate for strict-high priority queues. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.



NOTE: The **transmit-rate** setting works only if you also configure the **guaranteed-rate** in the traffic control profile that is attached to the forwarding class set to which the queue belongs. If you do not configure the **guaranteed-rate**, the **transmit-rate** does not work. The sum of all queue transmit rates in a forwarding class set should not exceed the traffic control profile guaranteed rate. If you configure transmit rates whose sum exceeds the forwarding class set guaranteed rate, the commit check fails, and the system rejects the configuration.



NOTE: Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **shaping-rate**—Maximum bandwidth, also known as the peak information rate (PIR). Each queue receives a maximum of the configured amount of absolute bandwidth or the configured percentage of bandwidth, even if more bandwidth is available.



NOTE: Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **priority**—One of two bandwidth priorities that queues associated with a scheduler can receive:

- **low**—The scheduler has low priority.
- **strict-high**—The scheduler has strict-high priority. You can configure only one queue as a strict-high priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.

We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- **drop-profile-map**—Mapping of a drop profile to a loss priority and protocol to apply WRED to the scheduler.
- **buffer-size**—Size of the queue buffer as a percentage of the dedicated buffer space on the port, or as a proportional share of the dedicated buffer space on the port that remains after the explicitly configured queues are served.
- **explicit-congestion-notification**—Enables ECN on a best-effort queue. ECN enables end-to-end congestion notification between two ECN-enabled endpoints on TCP/IP based networks. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. ECN is disabled by default.



NOTE: Ingress port congestion can occur during periods of egress port congestion if an ingress port forwards traffic to more than one egress port, and at least one of those egress ports experiences congestion. If this occurs, the congested egress port can cause the ingress port to exceed its fair allocation of ingress buffer resources. When the ingress port exceeds its buffer resource allocation, frames are dropped at the ingress. Ingress port frame drop affects not only the congested egress ports, but also all of the egress ports to which the congested ingress port forwards traffic.

If a congested ingress port drops traffic that is destined for one or more uncongested egress ports, configure a weighted random early detection (WRED) drop profile and apply it to the egress queue that is causing the congestion. The drop profile prevents the congested egress queue from affecting egress queues on other ports by dropping frames at the egress instead of causing congestion at the ingress port.



NOTE: Do not configure drop profiles for the fcoe and no-loss forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

Scheduler maps associate schedulers with forwarding classes (queues). After defining schedulers and mapping them to queues in a scheduler map, to configure hardware queue scheduling (port scheduling) you:

1. Associate a scheduler map with a traffic control profile (a traffic control profile schedules resources for a group of forwarding classes, called a *forwarding class set* or *priority group*).
2. Attach a forwarding class and a traffic control profile to an interface.

You can associate up to four user-defined scheduler maps with forwarding class sets.

This process configures the hardware queues, packet schedulers, and WRED characteristics that operate according to the scheduler mapping. The traffic control profile uses the scheduler CoS properties to determine the resources that should be allocated to the individual output queues from the total resources available to the priority group.

Table 92 on page 292 shows the configuration components for this example.

Table 92: Components of the Queue Scheduler Configuration Example

Component	Settings
Hardware	QFX3500 switch
Scheduler	Name: be-sched Transmit rate: 20% Shaping rate: 40% Buffer size: 20% Priority: low Drop profile: be-dp ECN: disable (default)
Scheduler map	Name: be-map Forwarding class to associate with the be-sched scheduler: best-effort
Traffic control profile	Name: be-tcp NOTE: This topic does not describe how to define a traffic control profile.
Forwarding class set	Name: lan-pg

Configuring a CoS Scheduler

To configure a CoS scheduler using the CLI:

1. Create a scheduler (**be-sched**) with a minimum guaranteed bandwidth of 2 Gbps, a maximum bandwidth of 4 Gbps, low priority, and map it to the drop profile **be-dp**:

```
[edit class-of-service schedulers]
user@switch# set be-sched transmit-rate percent 20
user@switch# set be-sched shaping-rate percent 40
user@switch# set be-sched buffer-size percent 20
user@switch# set be-sched priority low
user@switch# set be-sched drop-profile-map loss-priority low protocol any drop-profile
be-dp
```



NOTE: Because ECN is disabled by default, no ECN configuration is shown.

2. Configure a scheduler map (**be-map**) that associates the scheduler (**be-sched**) with the forwarding class (**best-effort**):

```
[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched
```

3. Associate the scheduler map **be-map** with a traffic control profile (**be-tcp**):

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp scheduler-map be-map
```

4. Associate the traffic control profile **be-tcp** with a forwarding class set (**lan-pg**) and a 10-Gigabit Ethernet interface (**xe-0/0/7**):

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/7 forwarding-class-set lan-pg
output-traffic-control-profile be-tcp
```

5. Alternatively, you can assign the scheduler map (**be-map**) to all the 10-Gigabit Ethernet interfaces using wildcards (**xe-***):

```
[edit class-of-service interfaces]
user@switch# set xe-* forwarding-class-set lan-pg output-traffic-control-profile be-tcp
```

Verification

To verify that the queue scheduler has been created and is mapped to the correct interfaces, perform these tasks:

- [Verifying the Scheduler Configuration on page 293](#)
- [Verifying the Scheduler Map Configuration on page 294](#)
- [Verifying That the Scheduler Is Associated with the Interface on page 294](#)

Verifying the Scheduler Configuration

Purpose Verify that the queue scheduler **be-sched** has been created with a minimum guaranteed bandwidth of 2 Gbps, a maximum bandwidth of 4 Gbps, the priority set to **low**, and the drop profile **be-dp**.

Action Display the scheduler using the operational mode command **show configuration class-of-service schedulers be-sched**:

```
user@switch> show configuration class-of-service schedulers be-sched
transmit-rate percent 20;
shaping-rate percent 40;
buffer-size percent 20
priority low;
drop-profile-map loss-priority low protocol any drop-profile be-dp;
```

Verifying the Scheduler Map Configuration

Purpose Verify that the scheduler map **be-map** has been created and associates the forwarding class **best-effort** with the scheduler **be-sched**, and also that the scheduler map is attached to the traffic control profile **be-tcp**.

Action Display the scheduler map using the operational mode command **show configuration class-of-service scheduler-maps be-map**:

```
user@switch> show configuration class-of-service scheduler-maps be-map
forwarding-class best-effort scheduler be-sched;
```

Display the traffic control profile to verify that the scheduler map **be-map** is attached using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp scheduler-map**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp scheduler-map
scheduler-map be-map;
```



NOTE: This topic does not describe how to configure a traffic control profile or its allocation of port bandwidth. Using a traffic control profile to configure the port resource allocation to the priority group is necessary to implement hierarchical scheduling.

Verifying That the Scheduler Is Associated with the Interface

Purpose Verify that the forwarding class set (**lan-pg**) and the traffic control profile (**be-tcp**) that are associated with the queue scheduler are attached to the interface **xe-0/0/7**.

Action List the interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
    lan-pg {
        output-traffic-control-profile be-tcp;
    }
}
```

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)

- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 300](#)
- [Example: Configuring Maximum Output Bandwidth on page 305](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
- [Example: Configuring WRED Drop Profiles on page 279](#)
- [Example: Configuring ECN](#)
- [Defining CoS Queue Schedulers on page 410](#)
- [Monitoring CoS Scheduler Maps on page 549](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)
- [Understanding CoS Scheduling on QFabric System Node Device Fabric \(fte\) Ports on page 199](#)
- [Understanding Default CoS Scheduling on QFabric System Interconnect Devices \(Junos OS Release 13.1 and Later Releases\) on page 203](#)
- [Understanding CoS Buffer Configuration on page 131](#)

Example: Configuring Queue Scheduling Priority

You can configure the bandwidth scheduling priority of individual queues by specifying the priority in a scheduler, and then using a scheduler map to associate the scheduler with a queue.

- [Requirements on page 295](#)
- [Overview on page 295](#)
- [Configuring Queue Scheduling Priority on page 296](#)
- [Verification on page 297](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

Queues can have one of two bandwidth priorities:

- **strict-high**—You can configure only one queue as a strict-high or high-priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.



NOTE: If you configure strict-high priority queues, you must create an fc-set that is dedicated only to strict-high priority traffic. Only one fc-set can contain strict-high priority queues. Queues that are not strict-high priority cannot belong to the same fc-set as strict-high priority queues. The multidestination fc-set cannot contain strict-high priority queues.

We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- **low**—Low priority. Traffic with this priority is serviced after any queue that has a **strict-high** priority.

Table 93 on page 296 shows the configuration components for this example.

This example describes how to set the queue priority for two forwarding classes (queues) named **fcoe** and **no-loss**. Both queues have a priority of **low**. The scheduler for the **fcoe** queue is named **fcoe-sched** and the scheduler for the **no-loss** queue is named **nl-sched**. One scheduler map, **schedmap1**, associates the schedulers to the queues.

Table 93: Components of the Queue Scheduler Priority Configuration Example

Component	Settings
Hardware	QFX3500 switch
Schedulers	fcoe-sched for FCoE traffic nl-sched for no-loss traffic
Priority	low for FCoE traffic low for no-loss traffic
Scheduler map	schedmap1 : FCoE mapping: scheduler fcoe-sched to forwarding class fcoe No-loss mapping: scheduler nl-sched to forwarding class no-loss

Configuring Queue Scheduling Priority

To configure queue priority using the CLI:

1. Create the FCoE scheduler with **low** priority:

```
[edit class-of-service]
user@switch# set schedulers fcoe-sched priority low
```
2. Create the no-loss scheduler with **low** priority:

```
[edit class-of-service]
```

- ```
user@switch# set schedulers nl-sched priority low
```
3. Associate the schedulers with the desired queues in the scheduler map:
- ```
[edit class-of-service]
user@switch# set scheduler-maps schedmap1 forwarding-class fcoe scheduler fcoe-sched
user@switch# set scheduler-maps schedmap1 forwarding-class no-loss scheduler nl-sched
```

Verification

To verify that you configured the queue scheduling priority for bandwidth and mapped the schedulers to the correct forwarding classes, perform these tasks:

- [Verifying the Queue Scheduling Priority on page 297](#)
- [Verifying the Scheduler-to-Forwarding-Class Mapping on page 297](#)

Verifying the Queue Scheduling Priority

Purpose Verify that you configured the queue schedulers **fcoe-sched** and **nl-sched** with **low** queue scheduling priority.

Action Display the **fcoe-sched** scheduler priority configuration using the operational mode command **show configuration class-of-service schedulers fcoe-sched priority**:

```
user@switch> show configuration class-of-service schedulers fcoe-sched priority
priority low;
```

Display the **nl-sched** scheduler priority configuration using the operational mode command **show configuration class-of-service schedulers nl-sched priority**:

```
user@switch> show configuration class-of-service schedulers nl-sched priority
priority low;
```

Verifying the Scheduler-to-Forwarding-Class Mapping

Purpose Verify that you configured the scheduler map **schedmap1** to map scheduler **fcoe-sched** to forwarding class **fcoe** and schedule **nl-sched** to forwarding class **no-loss**.

Action Display the scheduler map **schedmap1** using the operational mode command **show configuration class-of-service scheduler-maps schedmap1**:

```
user@switch> show configuration class-of-service scheduler-maps schedmap1
forwarding-class fcoe scheduler fcoe-sched;
forwarding-class no-loss scheduler nl-sched;
```

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
 - [Example: Configuring Queue Schedulers on page 289](#)
 - [Defining CoS Queue Scheduling Priority on page 414](#)
 - [Monitoring CoS Scheduler Maps on page 549](#)
 - [Understanding CoS Output Queue Schedulers on page 86](#)

Example: Configuring Traffic Control Profiles (Priority Group Scheduling)

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (queues) mapped to a forwarding class set share the bandwidth resources that you configure in the traffic control profile. A scheduler map associates forwarding classes with schedulers to define how the individual queues in a forwarding class set share the bandwidth allocated to that forwarding class set.

- [Requirements on page 298](#)
- [Overview on page 298](#)
- [Configuring a Traffic Control Profile on page 299](#)
- [Verification on page 300](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

The parameters you configure in a traffic control profile define the following characteristics for the priority group:

- **guaranteed-rate**—Minimum bandwidth, also known as the committed information rate (CIR). Each priority group receives a minimum of either the configured amount of absolute bandwidth or the configured percentage of bandwidth. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.



NOTE: In order for the **transmit-rate** option (minimum bandwidth for a queue that you set using scheduler configuration) to work properly, you must configure the **guaranteed-rate** for the priority group. If a priority group does not have a guaranteed minimum bandwidth, the queues (forwarding classes) that belong to the priority group cannot have a guaranteed minimum bandwidth.



NOTE: Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **shaping-rate**—Maximum bandwidth, also known as the peak information rate (PIR). Each priority group receives a maximum of the configured amount of absolute

bandwidth or the configured percentage of bandwidth, even if more bandwidth is available.



NOTE: Include the preamble bytes and interframe gap bytes as well as the data bytes in your bandwidth calculations.

- **scheduler-map**—Bandwidth and scheduling characteristics for the queues, defined by mapping forwarding classes to schedulers. (The queue scheduling characteristics represent amounts or percentages of the priority group bandwidth, not the amounts or percentages of total link bandwidth.)



NOTE: Because a port can have more than one priority group, when you assign resources to a priority group, keep in mind that the total port bandwidth must serve all of the queues associated with that port.

For example, if you map three priority groups to a 10-Gigabit Ethernet port, the queues associated with all three of the priority groups share the 10-Gbps bandwidth as defined by the traffic control profiles. Therefore, the total combined guaranteed-rate value of the three priority groups should not exceed 10 Gbps. If you configure guaranteed rates whose sum exceeds the port bandwidth, the system sends a syslog message to notify you that the configuration is not valid. However, the system does not perform a commit check. If you commit a configuration in which the sum of the guaranteed rates exceeds the port bandwidth, the hierarchical scheduler behaves unpredictably.

The sum of the queue (forwarding class) transmit rates cannot exceed the total guaranteed-rate of the priority group to which the queues belong. If you configure transmit rates whose sum exceeds the priority group guaranteed rate, the commit check fails and the system rejects the configuration.

If you configure the guaranteed-rate of a priority group as a percentage, configure all of the transmit rates associated with that priority group as percentages. In this case, if any of the transmit rates are configured as absolute values instead of percentages, the configuration is not valid and the system sends a syslog message.

Configuring a Traffic Control Profile

This example describes how to configure a traffic control profile named **san-tcp** with a scheduler map named **san-map1** and allocate to it a minimum bandwidth of 4 Gbps and a maximum bandwidth of 8 Gbps:

1. Create the traffic control profile and set the **guaranteed-rate** (minimum guaranteed bandwidth) to **4g**:

```
[edit class-of-service]
```

```
user@switch# set traffic-control-profiles san-tcp guaranteed-rate 4g
```

2. Set the **shaping-rate** (maximum guaranteed bandwidth) to **8g**:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp shaping-rate 8g
```

3. Associate the scheduler map **san-map1** with the traffic control profile:

```
[edit class-of-service]
user@switch# set traffic-control-profiles san-tcp scheduler-map san-map1
```

Verification

Verifying the Traffic Control Profile Configuration

Purpose Verify that the traffic control profile **san-tcp** has been created with a minimum guaranteed bandwidth of 4 Gbps, a maximum bandwidth of 8 Gbps, and the scheduler map **san-map1**.

Action List the traffic control profile using the operational mode command **show configuration class-of-service traffic-control-profiles san-tcp**:

```
user@switch> show configuration class-of-service traffic-control-profiles san-tcp
scheduler-map san-map1;
shaping-rate percent 8g;
guaranteed-rate 4g;
```

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 300](#)
- [Example: Configuring Maximum Output Bandwidth on page 305](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 416](#)
- [Understanding CoS Traffic Control Profiles on page 98](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)

Example: Configuring Minimum Guaranteed Output Bandwidth

Scheduling the minimum guaranteed output bandwidth for a queue (forwarding class) requires configuring both tiers of the two-tier hierarchical scheduler. One tier is scheduling the resources for the individual queue. The other tier is scheduling the resources for the priority group (forwarding class set) to which the queue belongs.

- [Requirements on page 300](#)
- [Overview on page 301](#)
- [Configuring Guaranteed Minimum Bandwidth on page 302](#)
- [Verification on page 303](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch

- Junos OS Release 11.1 or later for the QFX Series

Overview

The priority group minimum guaranteed bandwidth defines the minimum total amount of bandwidth available for all of the queues in the priority group to meet their minimum bandwidth requirements.

The **transmit-rate** setting in the scheduler configuration determines the minimum guaranteed bandwidth for an individual queue. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue.

The **guaranteed-rate** setting in the traffic control profile configuration determines the minimum guaranteed bandwidth for a priority group. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.



NOTE: You must configure both the **transmit-rate** value for the queue and the **guaranteed-rate** value for the priority group in order to set a valid minimum bandwidth guarantee for a queue. (If the priority group does not have a guaranteed minimum bandwidth, there is no guaranteed bandwidth pool from which the queue can take its guaranteed minimum bandwidth.)

The sum of the queue transmit rates in a priority group should not exceed the guaranteed rate for the priority group. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)



NOTE: When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.



NOTE: You cannot configure minimum guaranteed bandwidth on strict-high priority queues or on a priority group that contains strict-high priority queues.

This example describes how to:

- Configure a transmit rate (minimum guaranteed queue bandwidth) of 2 Gbps for queues in a scheduler named **be-sched**.
- Configure a guaranteed rate (minimum guaranteed priority group bandwidth) of 4 Gbps for a priority group in a traffic control profile named **be-tcp**.
- Assign the scheduler to a queue named **best-effort** by using a scheduler map named **be-map**.
- Associate the scheduler map **be-map** with the traffic control profile **be-tcp**.
- Assign the queue **best-effort** to a priority group named **be-pg**.
- Assign the priority group and the minimum guaranteed bandwidth scheduling to the egress interface **xe-0/0/7**.

Table 94 on page 302 shows the configuration components for this example:

Table 94: Components of the Minimum Guaranteed Output Bandwidth Configuration Example

Component	Settings
Hardware	QFX3500 switch
Minimum guaranteed queue bandwidth	Transmit rate: 2g
Minimum guaranteed priority group bandwidth	Guaranteed rate: 4g
Scheduler	be-sched
Scheduler map	be-map
Traffic control profile	be-tcp
Forwarding class set (priority group)	be-pg
Queue (forwarding class)	best-effort
Egress interface	xe-0/0/7

Configuring Guaranteed Minimum Bandwidth

To configure the minimum guaranteed bandwidth hierarchical scheduling for a queue and a priority group:

1. Configure the minimum guaranteed queue bandwidth of 2 Gbps for scheduler **be-sched**:

```
[edit class-of-service schedulers]
user@switch# set be-sched transmit-rate 2g
```
2. Configure the minimum guaranteed priority group bandwidth of 4 Gbps for traffic control profile **be-tcp**:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp guaranteed-rate 4g
```

3. Associate the scheduler **be-sched** with the **best-effort** queue in the scheduler map **be-map**:

```
[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched
```

4. Associate the scheduler map with the traffic control profile:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp scheduler-map be-map
```

5. Assign the **best-effort** queue to the priority group **be-pg**:

```
[edit class-of-service forwarding-class-sets]
user@switch# set be-pg class best-effort
```

6. Apply the configuration to interface **xe-0/0/7**:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

Verification

To verify the minimum guaranteed output bandwidth configuration, perform these tasks:

- [Verifying the Minimum Guaranteed Queue Bandwidth on page 303](#)
- [Verifying the Priority Group Minimum Guaranteed Bandwidth and Scheduler Map Association on page 303](#)
- [Verifying the Scheduler Map Configuration on page 304](#)
- [Verifying Queue \(Forwarding Class\) Membership in the Priority Group on page 304](#)
- [Verifying the Egress Interface Configuration on page 304](#)

Verifying the Minimum Guaranteed Queue Bandwidth

- Purpose** Verify that you configured the minimum guaranteed queue bandwidth as **2g** in the scheduler **be-sched**.
- Action** Display the minimum guaranteed bandwidth in the **be-sched** scheduler configuration using the operational mode command **show configuration class-of-service schedulers be-sched transmit-rate**:

```
user@switch> show configuration class-of-service schedulers be-sched transmit-rate
2g;
```

Verifying the Priority Group Minimum Guaranteed Bandwidth and Scheduler Map Association

- Purpose** Verify that the minimum guaranteed priority group bandwidth is **4g** and the attached scheduler map is **be-map** in the traffic control profile **be-tcp**.

Action Display the minimum guaranteed bandwidth in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp guaranteed-rate**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp guaranteed-rate 4g;
```

Display the scheduler map in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp scheduler-map**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp scheduler-map scheduler-map be-map;
```

Verifying the Scheduler Map Configuration

Purpose Verify that the scheduler map **be-map** maps the forwarding class **best-effort** to the scheduler **be-sched**.

Action Display the **be-map** scheduler map configuration using the operational mode command **show configuration class-of-service schedulers maps be-map**:

```
user@switch> show configuration class-of-service scheduler-maps be-map forwarding-class best-effort scheduler be-sched;
```

Verifying Queue (Forwarding Class) Membership in the Priority Group

Purpose Verify that the forwarding class set **be-pg** includes the forwarding class **best-effort**.

Action Display the **be-pg** forwarding class set configuration using the operational mode command **show configuration class-of-service forwarding-class-sets be-pg**:

```
user@switch> show configuration class-of-service forwarding-class-sets be-pg class best-effort;
```

Verifying the Egress Interface Configuration

Purpose Verify that the forwarding class set **be-pg** and the traffic control profile **be-tcp** are attached to egress interface **xe-0/0/7**.

Action Display the egress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7 forwarding-class-set {
    be-pg {
        output-traffic-control-profile be-tcp;
    }
}
```

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)

- [Example: Configuring Queue Scheduling Priority on page 295](#)
- [Example: Configuring Forwarding Class Sets on page 286](#)
- [Understanding CoS Traffic Control Profiles on page 98](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)

Example: Configuring Maximum Output Bandwidth

Scheduling the maximum output bandwidth for a queue (forwarding class) requires configuring both tiers of the hierarchical scheduler. One tier is scheduling the resources for the individual queue. The other tier is scheduling the resources for the priority group (forwarding class set) to which the queue belongs.

- [Requirements on page 305](#)
- [Overview on page 305](#)
- [Configuring Maximum Bandwidth on page 306](#)
- [Verification on page 307](#)

Requirements

This example uses the following hardware and software components:

- A Juniper Networks QFX3500 Switch
- Junos OS Release 11.1 or later for the QFX Series

Overview

The priority group maximum bandwidth defines the maximum total amount of bandwidth available for all of the queues in the priority group.

The **shaping-rate** setting in the scheduler configuration determines the maximum bandwidth for an individual queue.

The **shaping-rate** setting in the traffic control profile configuration determines the maximum bandwidth for a priority group.



NOTE: When you configure bandwidth for a queue or a priority group, the switch considers only the data as the configured bandwidth. The switch does not account for the bandwidth consumed by the preamble and the interframe gap (IFG). Therefore, when you calculate and configure the bandwidth requirements for a queue or for a priority group, consider the preamble and the IFG as well as the data in the calculations.



NOTE: When you set the maximum bandwidth (**shaping-rate**) for a queue or for a priority group at 100 Kbps or less, the traffic shaping behavior is accurate only within ± 20 percent of the configured **shaping-rate** value.

This example describes how to:

- Configure a maximum rate of 4 Gbps for queues in a scheduler named **be-sched**.
- Configure a maximum rate of 6 Gbps for a priority group in a traffic control profile named **be-tcp**.
- Assign the scheduler to a queue named **best-effort** by using a scheduler map named **be-map**.
- Associate the scheduler map **be-map** with the traffic control profile **be-tcp**.
- Assign the queue **best-effort** to a priority group named **be-pg**.
- Assign the priority group and the bandwidth scheduling to the interface **xe-0/0/7**.

Table 95 on page 306 shows the configuration components for this example:

Table 95: Components of the Maximum Output Bandwidth Configuration Example

Component	Settings
Hardware	QFX3500 switch
Maximum queue bandwidth	Shaping rate: 4g
Maximum priority group bandwidth	Shaping rate: 6g
Scheduler	be-sched
Scheduler map	be-map
Traffic control profile	be-tcp
Forwarding class set (priority group)	be-pg
Queue (forwarding class)	best-effort
Egress interface	xe-0/0/7

Configuring Maximum Bandwidth

To configure the maximum bandwidth hierarchical scheduling for a queue and a priority group:

1. Configure the maximum queue bandwidth of 4 Gbps for scheduler **be-sched**:

```
[edit class-of-service schedulers]
user@switch# set be-sched shaping-rate 4g
```

2. Configure the maximum priority group bandwidth of 6 Gbps for traffic control profile **be-tcp**:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp shaping-rate 6g
```

3. Associate the scheduler **be-sched** with the **best-effort** queue in the scheduler map **be-map**:

```
[edit class-of-service scheduler-maps]
user@switch# set be-map forwarding-class best-effort scheduler be-sched
```

4. Associate the scheduler map with the traffic control profile:

```
[edit class-of-service traffic-control-profiles]
user@switch# set be-tcp scheduler-map be-map
```

5. Assign the **best-effort** queue to the priority group **be-pg**:

```
[edit class-of-service forwarding-class-sets]
user@switch# set be-pg class best-effort
```

6. Apply the configuration to interface **xe-0/0/7**:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set be-pg output-traffic-control-profile be-tcp
```

Verification

To verify the maximum output bandwidth configuration, perform these tasks:

- [Verifying the Maximum Queue Bandwidth on page 307](#)
- [Verifying the Priority Group Maximum Bandwidth and Scheduler Map Association on page 307](#)
- [Verifying the Scheduler Map Configuration on page 308](#)
- [Verifying Queue \(Forwarding Class\) Membership in the Priority Group on page 308](#)
- [Verifying the Egress Interface Configuration on page 308](#)

Verifying the Maximum Queue Bandwidth

Purpose Verify that you configured the maximum queue bandwidth as **4g** in the scheduler **be-sched**.

Action List the maximum bandwidth in the **be-sched** scheduler configuration using the operational mode command **show configuration class-of-service schedulers be-sched shaping-rate**:

```
user@switch> show configuration class-of-service schedulers be-sched shaping-rate
4g;
```

Verifying the Priority Group Maximum Bandwidth and Scheduler Map Association

Purpose Verify that the maximum priority group bandwidth is **6g** and the attached scheduler map is **be-map** in the traffic control profile **be-tcp**.

Action List the maximum bandwidth in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp shaping-rate**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp shaping-rate
4g;
```

List the scheduler map in the **be-tcp** traffic control profile configuration using the operational mode command **show configuration class-of-service traffic-control-profiles be-tcp scheduler-map**:

```
user@switch> show configuration class-of-service traffic-control-profiles be-tcp scheduler-map
scheduler-map be-map;
```

Verifying the Scheduler Map Configuration

Purpose Verify that the scheduler map **be-map** maps the forwarding class **best-effort** to the scheduler **be-sched**.

Action List the **be-map** scheduler map configuration using the operational mode command **show configuration class-of-service schedulers maps be-map**:

```
user@switch> show configuration class-of-service scheduler-maps be-map
forwarding-class best-effort scheduler be-sched;
```

Verifying Queue (Forwarding Class) Membership in the Priority Group

Purpose Verify that the forwarding class set **be-pg** includes the forwarding class **best-effort**.

Action List the **be-pg** forwarding class set configuration using the operational mode command **show configuration class-of-service forwarding-class-sets be-pg**:

```
user@switch> show configuration class-of-service forwarding-class-sets be-pg
class best-effort;
```

Verifying the Egress Interface Configuration

Purpose Verify that the forwarding class set **be-pg** and the traffic control profile **be-tcp** are attached to egress interface **xe-0/0/7**.

Action List the egress interface using the operational mode command **show configuration class-of-service interfaces xe-0/0/7**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/7
forwarding-class-set {
    be-pg {
        output-traffic-control-profile be-tcp;
    }
}
```

Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)

- [Example: Configuring Forwarding Class Sets on page 286](#)
- [Understanding CoS Traffic Control Profiles on page 98](#)
- [Understanding CoS Hierarchical Port Scheduling \(ETS\) on page 81](#)

Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch)

The default system configuration supports FCoE traffic on priority 3 (IEEE 802.1p code point 011). If the FCoE traffic on your converged Ethernet network uses priority 3, the only user configuration required for lossless transport is to enable PFC on code point 011 on the FCoE ingress interfaces.

However, if your network uses a different priority than 3 for FCoE traffic, you need to configure lossless FCoE transport on that priority. This example shows you how to configure lossless FCoE transport on a converged Ethernet network that uses priority 5 (IEEE 802.1p code point 101) for FCoE traffic instead of using priority 3.

- [Requirements on page 309](#)
- [Overview on page 309](#)
- [Configuration on page 311](#)
- [Verification on page 313](#)

Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch in transit switch (FIP snooping) mode
- Junos OS Release 12.3 or later for the QFX Series

Overview

Although FCoE traffic typically uses IEEE 802.1p priority 3 on converged Ethernet networks, some networks use a different priority for FCoE traffic. Regardless of the priority used, FCoE traffic must receive lossless treatment. Supporting lossless behavior for FCoE traffic when your network does not use priority 3 requires configuring:

- A lossless forwarding class for FCoE traffic.
- A behavior aggregate (BA) classifier to map the FCoE forwarding class to the appropriate IEEE 802.1p priority.
- A congestion notification profile (CNP) to enable PFC on the FCoE code point at the interface ingress and to configure flow control on the interface egress. Flow control on the interface egress enables the interface to respond to PFC messages received from the connected peer and pause the correct IEEE 802.1p priority on the correct output queue.



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- A DCBX application and an application map to support DCBX application TLV exchange for the lossless FCoE traffic on the configured FCoE priority. By default, DCBX is enabled on all Ethernet interfaces, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifiers, CNP, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

Topology

This example shows how to configure one lossless FCoE traffic class, map it to a priority other than priority 3, and configure flow control to ensure lossless behavior on the interfaces. This example uses two Ethernet interfaces, xe-0/0/25 and xe-0/0/26. The interfaces connect to a converged Ethernet network that uses IEEE 802.1p priority 5 (code point 101) for FCoE traffic.

The configuration on the two interfaces is the same. Both interfaces use the same explicitly configured lossless FCoE forwarding class and the same ingress classifier. Both interfaces enable PFC on priority 5 and enable flow control on the same output queue (which is mapped to the lossless FCoE forwarding class).

Table 96 on page 310 shows the configuration components for this example.

Table 96: Components of the Configuration Topology for FCoE Traffic That Does Not Use Priority 3

Component	Settings
Hardware	QFX3500 switch
Forwarding class	Name— fcoe1 Queue mapping—queue 5 Packet drop attribute— no-loss NOTE: A lossless forwarding class can be mapped to any output queue. However, because the fcoe1 forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.

Table 96: Components of the Configuration Topology for FCoE Traffic That Does Not Use Priority 3 (*continued*)

Component	Settings
BA classifier	<p>Name—fcoe_p5</p> <p>FCoE priority mapping—Forwarding class fcoe1 mapped to code point 101 (IEEE 802.1p priority 5) and a packet loss priority of low.</p>
PFC configuration (CNPs)	<p>CNP name—fcoe_p5_cnp</p> <p>Input CNP code point—101</p> <p>MRU—2240 bytes</p> <p>Cable length—100 meters</p> <p>Output CNP code point—101</p> <p>Output CNP flow control queue—5</p> <p>NOTE: When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for pause in the default configuration (queues 3 and 4) are not enabled for pause unless they are included in the explicitly configured output CNP.</p>
DCBX application mapping	<p>Application name—fcoe_p5_app</p> <p>Application EtherType—0x8906</p> <p>Application map name—fcoe_p5_app_map</p> <p>Application map code points—101</p> <p>NOTE: LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.</p>



NOTE: This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

Configuration

CLI Quick Configuration

To quickly configure a lossless FCoE forwarding class that uses a different priority than IEEE 802.1p priority 3 for FCoE traffic on an FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service interfaces xe-0/0/25 unit 0 classifiers ieee-802.1 fcoe_p5
```

```

set class-of-service interfaces xe-0/0/26 unit 0 classifiers ieee-802.1 fcoe_p5
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 101
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p5_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service interfaces xe-0/0/25 congestion-notification-profile fcoe_p5_cnp
set class-of-service interfaces xe-0/0/26 congestion-notification-profile fcoe_p5_cnp
set applications application fcoe_p5_app ether-type 0x8906
set policy-options application-maps fcoe_p5_app_map application fcoe_p5_app code-points 101
set protocols dcbx interface xe-0/0/25 application-map fcoe_p5_app_map
set protocols dcbx interface xe-0/0/26 application-map fcoe_p5_app_map

```

Configuring A Lossless FCoE Forwarding Class On IEEE 802.1p Priority 5

Step-by-Step Procedure

To configure a lossless forwarding class for FCoE traffic on IEEE 802.1p priority 5 (code point 101), classify FCoE traffic into the lossless forwarding class, configure a congestion notification profile to enable PFC on the FCoE priority and output queue, and configure DCBX application protocol TLV exchange for traffic on the FCoE priority:

1. Configure the lossless forwarding class (named **fcoe1** and mapped to output queue 5) for FCoE traffic on IEEE 802.1p priority 5:

```

[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss

```
2. Configure the ingress classifier (**fcoe_p5**). The classifier maps the FCoE priority (code point 101) to the lossless FCoE forwarding class **fcoe1**:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low code-points
101

```
3. Apply the classifier to interfaces **xe-0/0/25** and **xe-0/0/26**:

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/25 unit 0 classifiers ieee-802.1 fcoe_p5
user@switch# set interfaces xe-0/0/26 unit 0 classifiers ieee-802.1 fcoe_p5

```
4. Configure the CNP. The input stanza enables PFC on the FCoE priority (IEEE 802.1p code point 101), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queue 5 on the FCoE priority:

```

[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p5_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5

```
5. Apply the CNP to the interfaces:

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/25 congestion-notification-profile fcoe_p5_cnp
user@switch# set interfaces xe-0/0/26 congestion-notification-profile fcoe_p5_cnp

```
6. Configure the DCBX application for FCoE to map to the Ethernet interfaces, so that DCBX can exchange application protocol TLVs on the IEEE 802.1p priority 5 instead of on the default priority 3:

```
[edit]
user@switch# set applications application fcoe_p5_app ether-type 0x8906
```

7. Configure a DCBX application map to map the FCoE application to the correct IEEE 802.1p FCoE priority:

```
[edit]
user@switch# set policy-options application-maps fcoe_p5_app_map application
fcoe_p5_app code-points 101
```

8. Apply the application map to the Ethernet interfaces so that DCBX exchanges FCoE application TLVs on the correct code point:

```
[edit]
user@switch# set protocols dcbx interface xe-0/0/25 application-map fcoe_p5_app_map
user@switch# set protocols dcbx interface xe-0/0/26 application-map fcoe_p5_app_map
```

Verification

To verify the configuration and proper operation of the lossless forwarding class and IEEE 802.1p priority, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 313](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 314](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 314](#)
- [Verifying the Interface Configuration on page 315](#)
- [Verifying the DCBX Application Configuration on page 316](#)
- [Verifying the DCBX Application Map Configuration on page 316](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 316](#)

Verifying the Forwarding Class Configuration

Purpose Verify that the lossless forwarding class **fcoe1** has been created.

Action Show the forwarding class configuration by using the operational command **show class-of-service forwarding-class**:

```
user@switch# show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
fcoe	1	3	normal	Enabled
no-loss	2	4	normal	Enabled
network-control	3	7	normal	Disabled
fcoe1	4	5	normal	Enabled
mcast	8	8	normal	Disabled

Meaning The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue 5 with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

Verifying the Behavior Aggregate Classifier Configuration

Purpose Verify that the classifier maps the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

Action List the classifier configured to support lossless FCoE transport using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
```

Classifier: fcoe_p5, Code point type: ieee-802.1, Index: 63065	Code point	Forwarding class	Loss priority
101	fcoe1	low	

Meaning The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier.

Classifier **fcoe_p5** maps code point 101 (priority 5) to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**, and all other priorities to the **best-effort** forwarding class with a packet loss priority of **high**.

Verifying the PFC Flow Control Configuration (CNP)

Purpose Verify that PFC is enabled on the correct input priority and that flow control is configured on the correct output queue in the CNP.

Action Display the congestion notification profile using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
Name: fcoe_p5_cnp, Index: 12137
Type: Input
Cable Length: 100 m
  Priority    PFC          MRU
  000        Disabled
  001        Disabled
  010        Disabled
  011        Disabled
  100        Disabled
  101        Enabled    2240
  110        Disabled
  111        Disabled
Type: Output
  Priority    Flow-Control-Queues
  101
  5
```

Meaning The **show class-of-service congestion-notification** command shows the input and output stanzas of the configured CNPs.

The **fcoe_p5_cnp** CNP input stanza shows that PFC is enabled on code point **101** (priority 5), the MRU is **2240** bytes, and the cable length is **100** meters. The CNP output stanza shows that output flow control is configured on queue **5** for code point **101** (priority 5).

Verifying the Interface Configuration

Purpose Verify that the correct classifier and congestion notification profile are configured on the interfaces.

Action List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/25** and **show configuration class-of-service interfaces xe-0/0/26**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/25
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
  classifiers {
    ieee-802.1 fcoe_p5;
  }
}

user@switch> show configuration class-of-service interfaces xe-0/0/26
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
  classifiers {
    ieee-802.1 fcoe_p5;
  }
}
```

Meaning Both the **show configuration class-of-service interfaces xe-0/0/25** command and the **show configuration class-of-service interfaces xe-0/0/26** command show that the

congestion notification profile **fcoe_p5_cnp** is configured on each interface, and that the IEEE 802.1p classifier associated with each interface is **fcoe_p5**.

Verifying the DCBX Application Configuration

- Purpose** Verify that the DCBX application for FCoE is configured.
- Action** List the DCBX applications by using the configuration mode command **show applications**:
- ```
user@switch# show applications
application fcoe_p5_app {
 ether-type 0x8906;
```
- Meaning** The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **fcoe\_p5\_app** is configured with an EtherType of **0x8906**.

### Verifying the DCBX Application Map Configuration

- Purpose** Verify that the application map is configured.
- Action** List the application maps by using the configuration mode command **show policy-options application-maps**:
- ```
user@switch# show policy-options application-maps
fcoe_p5_app_map {
    application fcoe_p5_app code-points 101;
}
```
- Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that application map **fcoe_p5_app_map** consists of the application named **fcoe_p5_app**, which is mapped to IEEE 802.1p code point 101.

Verifying the DCBX Application Protocol Exchange Interface Configuration

- Purpose** Verify that the application map is applied to the correct interfaces.
- Action** List the application maps on each interface using the configuration mode command **show protocols dcbx**:
- ```
user@switch# show protocols dcbx
interface xe-0/0/25.0 {
 application-map fcoe_p5_app_map;
}
interface xe-0/0/26.0 {
 application-map fcoe_p5_app_map;
}
```
- Meaning** The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that interfaces **xe-0/0/25.0** and **xe-0/0/26.0** use application map **fcoe\_p5\_app\_map**.



- Related Documentation**
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326](#)
  - [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 317](#)
  - [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 340](#)
  - [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)
  - [Example: Configuring Unicast Classifiers on page 273](#)
  - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
  - [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109](#)
  - [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)

## Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface

---

The default system configuration supports FCoE traffic on priority 3 (IEEE 802.1p code point 011). If the FCoE traffic on your converged Ethernet network uses priority 3, the only user configuration required for lossless transport is to enable PFC on code point 011 on the FCoE ingress interfaces.

However, if your converged Ethernet network uses more than one priority for FCoE traffic, you need to configure lossless transport for each FCoE priority. This example shows you how to configure lossless FCoE transport on a converged Ethernet network that uses both priority 3 (IEEE 802.1p code point 011) and priority 5 (IEEE 802.1p code point 101) for FCoE traffic.

- [Requirements on page 317](#)
- [Overview on page 317](#)
- [Configuration on page 320](#)
- [Verification on page 322](#)

### Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch in transit switch (FIP snooping) mode
- Junos OS Release 12.3 or later for the QFX Series

### Overview

Some network topologies support FCoE traffic on more than one IEEE 802.1p priority. For example, a converged Ethernet network might include two separate FCoE networks that use different priorities to identify traffic. Interfaces that carry traffic for both FCoE networks need to support lossless FCoE transport on both priorities.

Supporting lossless behavior for two FCoE traffic classes requires configuring:

- At least one lossless forwarding class for FCoE traffic (this example uses the default **fcoe** forwarding class as one of the lossless FCoE forwarding classes, so we need to explicitly configure only one FCoE forwarding class).
- A behavior aggregate (BA) classifier to map the FCoE forwarding classes to the appropriate IEEE 802.1p code points (priorities).
- A congestion notification profile (CNP) to enable PFC on the FCoE code points at the interface ingress and to configure PFC flow control on the interface egress so that the interface can respond to PFC messages received from the connected peer.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- DCBX applications and an application map to support DCBX application TLV exchange for the lossless FCoE traffic on the configured FCoE priorities. By default, DCBX is enabled on all Ethernet interfaces, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifier, CNP, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

### Topology

This example shows how to configure two lossless FCoE traffic classes on an interface, map them to two different priorities, and configure flow control to ensure lossless behavior. This example uses two Ethernet interfaces, xe-0/0/20 and xe-0/0/21, that are connected to the converged Ethernet network. Both interfaces transport FCoE traffic on priorities 3 (011) and 5 (101), and must support lossless transport of that traffic.

Table 97 on page 318 shows the configuration components for this example.

**Table 97: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology**

| Component | Settings       |
|-----------|----------------|
| Hardware  | QFX3500 switch |

**Table 97: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology (continued)**

| Component               | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding classes      | <p>Name—<b>fcoe1</b><br/> Queue mapping—queue <b>5</b><br/> Packet drop attribute—<b>no-loss</b></p> <p><b>NOTE:</b> A lossless forwarding class can be mapped to any output queue. However, because the <b>fcoe1</b> forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.</p> <p>Name—<b>fcoe</b><br/> This is the default lossless FCoE forwarding class, so no configuration required. The <b>fcoe</b> forwarding class is mapped to priority 3 (IEEE 802.1p code point 011) and to output queue 3 with a packet drop attribute of <b>no-loss</b>.</p>                              |
| BA classifier           | <p>Name—<b>fcoe_classifier</b></p> <p>FCoE priority mapping for forwarding class <b>fcoe</b>—mapped to code point <b>011</b> (IEEE 802.1p priority 3) and a packet loss priority of <b>low</b>.</p> <p>FCoE priority mapping for forwarding class <b>fcoe1</b>—mapped to code point <b>101</b> (IEEE 802.1p priority 5) and a packet loss priority of <b>low</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                           |
| PFC configuration (CNP) | <p>CNP name—<b>fcoe_cnp</b></p> <p>Input CNP code points—<b>011</b> and <b>101</b></p> <p>MRU—2240 bytes</p> <p>Cable length—100 meters</p> <p>Output CNP code points—<b>011</b> and <b>101</b></p> <p>Output CNP flow control queues—<b>3</b> and <b>5</b></p> <p><b>NOTE:</b> When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for PFC pause in the default configuration (queues 3 and 4) are not enabled for PFC pause unless they are included in the explicitly configured output CNP. In this example, because the explicit output CNP overwrites the default output CNP, we must explicitly configure flow control on queue 3.</p> |

**Table 97: Components of the Two Lossless FCoE Priorities on an Interface Configuration Topology (continued)**

| Component                | Settings                                                                                                                                                                                                                                                                                                                    |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DCBX application mapping | Application name— <b>fcoe_app</b><br><br>Application EtherType— <b>0x8906</b><br><br>Application map name— <b>fcoe_app_map</b><br><br>Application map code points— <b>011 and 101</b><br><br><i>NOTE:</i> LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces. |
| Interfaces               | Interfaces <b>xe-0/0/20</b> and <b>xe-0/0/21</b> use the same configuration: <ul style="list-style-type: none"> <li>• Classifier—<b>fcoe_classifier</b></li> <li>• CNP—<b>fcoe_cnp</b></li> <li>• DCBX application map—<b>fcoe_app_map</b></li> </ul>                                                                       |



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

## Configuration

### CLI Quick Configuration

To quickly configure two lossless FCoE forwarding classes that use different priorities on an FCoE transit switch interface, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_classifier forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_classifier forwarding-class fcoe1 loss-priority low
code-points 101set class-of-service interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_classifier
set class-of-service interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_classifier
set class-of-service congestion-notification-profile fcoe_cnp input ieee-802.1 code-point 011 pfc
mru 2240
set class-of-service congestion-notification-profile fcoe_cnp input ieee-802.1 code-point 101 pfc
mru 2240
set class-of-service congestion-notification-profile fcoe_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_cnp output ieee-802.1 code-point 011
pfc flow-control-queue 3
set class-of-service congestion-notification-profile fcoe_cnp output ieee-802.1 code-point 101
pfc flow-control-queue 5
set class-of-service interfaces xe-0/0/20 congestion-notification-profile fcoe_cnp
set class-of-service interfaces xe-0/0/21 congestion-notification-profile fcoe_cnp
set applications application fcoe_app ether-type 0x8906
set policy-options application-maps fcoe_app_map application fcoe_app code-points [011 101]
set protocols dcbx interface xe-0/0/20 application-map fcoe_app_map
set protocols dcbx interface xe-0/0/21 application-map fcoe_app_map
```

**Step-by-Step Procedure** To configure two lossless forwarding classes for FCoE traffic on the same interface, classify FCoE traffic into the forwarding classes, configure CNPs to enable PFC on the FCoE priorities and output queues, and configure DCBX application protocol TLV exchange for traffic on both FCoE priorities:

1. Configure lossless forwarding class **fcoe1** and map it to output queue **5** for FCoE traffic that uses IEEE 802.1p priority 5:

```
[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss
```



**NOTE:** This examples uses the default **fcoe** forwarding class as the other lossless FCoE forwarding class.

2. Configure the ingress classifier. The classifier maps the FCoE priorities (IEEE 802.1p code points 011 and 101) to lossless FCoE forwarding classes **fcoe** and **fcoe1**, respectively:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_classifier forwarding-class fcoe loss-priority low
code-points 011
user@switch# set ieee-802.1 fcoe_classifier forwarding-class fcoe1 loss-priority low
code-points 101
```

3. Apply the classifier to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_classifier
user@switch# set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_classifier
```

4. Configure the CNP. The input stanza enables PFC on the FCoE priorities (IEEE 802.1p code points 011 and 101), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queues 3 and 5 on the FCoE priorities:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_cnp input ieee-802.1 code-point
011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_cnp input ieee-802.1 code-point
101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_cnp output ieee-802.1 code-point
011 pfc flow-control-queue 3
user@switch# set congestion-notification-profile fcoe_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
```

5. Apply the CNP to the interfaces:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 congestion-notification-profile fcoe_cnp
user@switch# set interfaces xe-0/0/21 congestion-notification-profile fcoe_cnp
```

6. Configure a DCBX application for FCoE to map to the Ethernet interfaces, so that DCBX can exchange application protocol TLVs on both of the IEEE 802.1p priorities used for FCoE transport:

```
[edit]
```

```
user@switch# set applications application fcoe_app ether-type 0x8906
```

7. Configure a DCBX application map to map the FCoE application to the correct IEEE 802.1p FCoE priorities:

```
[edit]
```

```
user@switch# set policy-options application-maps fcoe_app_map application fcoe_app
code-points [011 101]
```

8. Apply the application map to the interfaces so that DCBX exchanges FCoE application TLVs on the correct code points:

```
[edit]
```

```
user@switch# set protocols dcbx interface xe-0/0/20 application-map fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/21 application-map fcoe_app_map
```

## Verification

To verify the configuration and proper operation of the lossless forwarding classes and IEEE 802.1p priorities, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 322](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 323](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 323](#)
- [Verifying the Interface Configuration on page 324](#)
- [Verifying the DCBX Application Configuration on page 325](#)
- [Verifying the DCBX Application Map Configuration on page 325](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 325](#)

### Verifying the Forwarding Class Configuration

**Purpose** Verify that the lossless forwarding class **fcoe1** has been created.

**Action** Show the forwarding class configuration by using the operational command **show class-of-service forwarding-class**:

```
user@switch# show class-of-service forwarding-class
```

| Forwarding class | ID | Queue | Policing priority | No-Loss  |
|------------------|----|-------|-------------------|----------|
| best-effort      | 0  | 0     | normal            | Disabled |
| fcoe             | 1  | 3     | normal            | Enabled  |
| no-loss          | 2  | 4     | normal            | Enabled  |
| network-control  | 3  | 7     | normal            | Disabled |
| fcoe1            | 4  | 5     | normal            | Enabled  |
| mcast            | 8  | 8     | normal            | Disabled |

**Meaning** The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue 5 with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

### Verifying the Behavior Aggregate Classifier Configuration

**Purpose** Verify that the three classifiers map the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

**Action** List the classifiers using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
```

Classifier: fcoe\_classifier, Code point type: ieee-802.1, Index: 10964

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 011        | fcoe             | low           |
| 101        | fcoe1            | low           |

**Meaning** The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier.

Classifier **fcoe\_classifier** maps code point **011** to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and maps code point **101** to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

### Verifying the PFC Flow Control Configuration (CNP)

**Purpose** Verify that PFC is enabled on the correct input priorities and that flow control is configured on the correct output queues and priorities.

**Action** List the CNPs using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
Name: fcoe_cnp, Index: 46504
Type: Input
Cable Length: 100 m
 Priority PFC MRU
 000 Disabled
 001 Disabled
 010 Disabled
 011 Enabled 2240
 100 Disabled
 101 Enabled 2240
 110 Disabled
 111 Disabled
Type: Output
 Priority Flow-Control-Queues
 011 3
 101 5
```

**Meaning** The **show class-of-service congestion-notification** command shows the input and output stanzas of the CNP.

The CNP **fcoe\_cnp** input stanza shows that PFC is enabled on code points **011** and **101**, the MRU is **2240** bytes on both priorities, and the interface cable length is **100** meters. The CNP output stanza shows that output flow control is configured on queues **3** and **5** for code points **011** and **101**, respectively.

---

### Verifying the Interface Configuration

**Purpose** Verify that the classifier and congestion notification profile are configured on the interfaces. Both interfaces should show the same configuration.

**Action** List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/20** and **show configuration class-of-service interfaces xe-0/0/21**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/20
congestion-notification-profile fcoe_cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe_classifier;
 }
}

user@switch> show configuration class-of-service interfaces xe-0/0/21
congestion-notification-profile fcoe_cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe_classifier;
 }
}
```



**Meaning** The `show configuration class-of-service interfaces xe-0/0/20` command shows that the congestion notification profile `fcoe_cnp` is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is `fcoe_classifier`.

The `show configuration class-of-service interfaces xe-0/0/21` command shows that the congestion notification profile `fcoe_cnp` is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is `fcoe_classifier`.

### Verifying the DCBX Application Configuration

**Purpose** Verify that the DCBX application for FCoE is configured.

**Action** List the DCBX applications by using the configuration mode command `show applications`:

```
user@switch# show applications
application fcoe_app {
 ether-type 0x8906;
```

**Meaning** The `show applications` configuration mode command shows all of the configured applications. The output shows that the application `fcoe_app` is configured with an EtherType of `0x8906`.

### Verifying the DCBX Application Map Configuration

**Purpose** Verify that the application map is configured.

**Action** List the application maps by using the configuration mode command `show policy-options application-maps`:

```
user@switch# show policy-options application-maps
fcoe_app_map {
 application fcoe_app code-points [011 101];
}
```

**Meaning** The `show policy-options application-maps` configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that application map `fcoe_app_map` consists of the application named `fcoe_app`, which is mapped to IEEE 802.1p code points `011` and `101` (priorities 3 and 5, respectively).

### Verifying the DCBX Application Protocol Exchange Interface Configuration

**Purpose** Verify that the application map is applied to the interfaces.

**Action** List the application maps on each interface using the configuration mode command `show protocols dcbx`:

```
user@switch# show protocols dcbx
interface xe-0/0/20.0 {
 application-map fcoe_app_map;
}
interface xe-0/0/21.0 {
```

```
 application-map fcoe_app_map;
}
```

**Meaning** The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that interfaces **xe-0/0/20.0** and **xe-0/0/21.0** use application map **fcoe\_app\_map**.

- Related Documentation**
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326](#)
  - [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 309](#)
  - [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 340](#)
  - [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)
  - [Example: Configuring Unicast Classifiers on page 273](#)
  - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
  - [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109](#)
  - [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)

---

## Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces

---

Although the default configuration provides two lossless forwarding classes mapped to two different IEEE 802.1p priorities (code points), you can explicitly configure up to six lossless forwarding classes and map them to different priorities. You can support up to six different types of lossless traffic, and you can support the same type of traffic if it uses different priorities in different parts of your converged network.

This example shows you how to configure two lossless forwarding classes for FCoE traffic and map them to two different priorities on an FCoE transit switch.

- [Requirements on page 326](#)
- [Overview on page 327](#)
- [Configuration on page 331](#)
- [Verification on page 334](#)

### Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch in transit switch (FIP snooping) mode
- Junos OS Release 12.3 or later for the QFX Series

## Overview

Some network topologies support FCoE traffic on more than one IEEE 802.1p priority. For example, when the QFX3500 switch acts as a transit switch, it could be connected to two QFX3500 switches in FCoE-FC gateway mode. Each of the gateway switches could connect a set of FCoE clients to a different SAN, and each set of FCoE clients could use a different priority for FCoE traffic to avoid fate sharing and maintain separation of the two FCoE networks. In this case, you need to configure two forwarding classes for FCoE traffic, each mapped to a different output queue and a different priority.

Supporting lossless behavior for two FCoE traffic classes requires configuring:

- At least one lossless forwarding class for FCoE traffic (this example uses the default **fcoe** forwarding class as one of the two lossless FCoE forwarding classes, so we need to explicitly configure only one FCoE forwarding class)
- Behavior aggregate (BA) classifiers to map the FCoE forwarding classes to the appropriate IEEE 802.1p code points (priorities) on each interface
- Congestion notification profiles (CNPs) for each interface to enable PFC on the FCoE code points at the interface ingress and to configure PFC flow control on the interface egress so that the interface can respond to PFC messages received from the connected peer



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- DCBX applications and an application map to support DCBX application TLV exchange for the lossless FCoE traffic on the configured FCoE priorities. By default, DCBX is enabled on all Ethernet interfaces, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifiers, CNPs, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

## Topology

This example shows how to configure two lossless FCoE traffic classes, map them to two different priorities, and configure flow control to ensure lossless behavior for those priorities on the interfaces. This example uses three Ethernet interfaces, xe-0/0/20, xe-0/0/21, and xe-0/0/22:

- Interface xe-0/0/20 connects to an FCoE-FC gateway that connects to Fibre Channel (FC) SAN 1. FCoE traffic to and from FC SAN 1 uses the default **fcoe** forwarding class and the default mapping to priority 3 (IEEE 802.1p code point 011) and output queue 3.
- Interface xe-0/0/21 connects to another FCoE-FC gateway that connects to Fibre Channel (FC) SAN 2. FCoE traffic to and from FC SAN-2 uses an explicitly configured FCoE forwarding class that is mapped to priority 5 (code point 101) and output queue 5.
- Interface xe-0/0/22 connects to FCoE devices on the converged Ethernet network and handles traffic destined for FC SAN 1 and FC SAN 2. Interface xe-0/0/22 must properly handle lossless FCoE traffic of both priorities (both FCoE forwarding classes), including pausing the traffic on ingress or egress as required.

Figure 15 on page 328 shows the topology for this example, and Table 98 on page 328 shows the configuration components for this example.

Figure 15: Topology of the Two Lossless FCoE Priorities Example




Table 98: Components of the Two Lossless FCoE Priorities Configuration Topology

| Component          | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware           | QFX3500 switch                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Forwarding classes | <p>Name—<b>fcoe1</b><br/> Queue mapping—queue 5<br/> Packet drop attribute—<b>no-loss</b></p> <p><b>NOTE:</b> A lossless forwarding class can be mapped to any output queue. However, because the <b>fcoe1</b> forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.</p> <p>Name—<b>fcoe</b><br/> This is the default lossless FCoE forwarding class, so no configuration required. The <b>fcoe</b> forwarding class is mapped to priority 3 (IEEE 802.1p code point 011) and to output queue 3 with a packet drop attribute of <b>no-loss</b></p> |

**Table 98: Components of the Two Lossless FCoE Priorities Configuration Topology (*continued*)**

| Component      | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BA classifiers | <p>Each interface requires a different classifier because each interface handles a different subset of FCoE traffic.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/20 classifier:<br/>Name—<b>fcoe_p3</b><br/>FCoE priority mapping—Forwarding class <b>fcoe</b> mapped to code point <b>011</b> (IEEE 802.1p priority 3) and a packet loss priority of <b>low</b>.</li> <li>Interface xe-0/0/21 classifier:<br/>Name—<b>fcoe_p5</b><br/>FCoE priority mapping—Forwarding class <b>fcoe1</b> mapped to code point <b>101</b> (IEEE 802.1p priority 5) and a packet loss priority of <b>low</b>.</li> <li>Interface xe-0/0/22 classifier:<br/>Name—<b>fcoe_p3_p5</b><br/>FCoE priority mapping—Forwarding class <b>fcoe1</b> mapped to code point <b>101</b> and a packet loss priority of <b>low</b>, and forwarding class <b>fcoe</b> mapped to code point <b>011</b> and a packet loss priority of <b>low</b>.</li> </ul> |

**Table 98: Components of the Two Lossless FCoE Priorities Configuration Topology** (*continued*)

| Component                | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PFC configuration (CNPs) | <p>Each interface requires a different CNP because each interface handles a different subset of FCoE traffic and must pause that traffic on different priorities.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/20 CNP:<br/>CNP name—<b>fcoe_p3_cnp</b><br/>Input CNP code point—<b>011</b><br/>MRU—2240 bytes<br/>Cable length—100 meters</li> </ul> <p><b>NOTE:</b> Because interface xe-0/0/20 uses the default FCoE configuration, output queue 3 is paused by default and you do not need to configure the output stanza of the CNP.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/21 CNP:<br/>CNP name—<b>fcoe_p5_cnp</b><br/>Input CNP code point—<b>101</b><br/>MRU—2240 bytes<br/>Cable length—150 meters<br/>Output CNP code point—<b>101</b><br/>Output CNP flow control queue—<b>5</b></li> <li>Interface xe-0/0/22 CNP:<br/>CNP name—<b>fcoe_p3_p5_cnp</b><br/>Input CNP code points—<b>011</b> and <b>101</b><br/>MRU—2240 bytes (both priorities)<br/>Cable length—100 meters<br/>Output CNP code points—<b>011</b> (for queue 3) and <b>101</b> (for queue 5)<br/>Output CNP flow control queues—<b>3</b> for priority 3 (code point 011) and <b>5</b> for priority 5 (code point 101)</li> </ul> <p><b>NOTE:</b> When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for pause in the default configuration (queues 3 and 4) are not enabled for pause unless they are included in the explicitly configured output CNP.</p> |

**Table 98: Components of the Two Lossless FCoE Priorities Configuration Topology** (*continued*)

| Component                | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DCBX application mapping | <p>Interface xe-0/0/20 does not need an application map because DCBX exchanges application protocol TLVs only on the default FCoE priority (priority 3).</p> <p>Interface xe-0/0/21 requires an application map that enables DCBX application protocol TLV exchange on priority 5 (code point 101) for FCoE traffic. Interface xe-0/0/22 requires an application map that enables DCBX application protocol TLV exchange both on priority 3 (code point 011) and on priority 5 (code point 101) for FCoE traffic.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/21 DCBX application mapping:<br/>Application name—<b>fcoe_p5_app</b><br/>Application ether-type—<b>0x8906</b><br/>Application map name—<b>fcoe_p5_app_map</b><br/>Application map code points—<b>101</b></li> <li>Interface xe-0/0/22 DCBX application mapping:<br/>Application name—<b>fcoe_all_app</b><br/>Application ether-type—<b>0x8906</b><br/>Application map name—<b>fcoe_all_app_map</b><br/>Application map code points—<b>011</b> and <b>101</b></li> </ul> <p><b>NOTE:</b> LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.</p> |



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

## Configuration

### CLI Quick Configuration

To quickly configure two lossless FCoE forwarding classes that use different priorities on an FCoE transit switch, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_p3 forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_p3
set class-of-service interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_p5
set class-of-service interfaces xe-0/0/22 unit 0 classifiers ieee-802.1 fcoe_p3_p5
set class-of-service congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point 011
pfc mru 2240
```

```

set class-of-service congestion-notification-profile fcoe_p3_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 101
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p5_cnp input cable-length 150
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
011 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
011 pfc flow-control-queue 3
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service interfaces xe-0/0/20 congestion-notification-profile fcoe_p3_cnp
set class-of-service interfaces xe-0/0/21 congestion-notification-profile fcoe_p5_cnp
set class-of-service interfaces xe-0/0/22 congestion-notification-profile fcoe_p3_p5_cnp
set applications application fcoe_p5_app ether-type 0x8906
set applications application fcoe_all_app ether-type 0x8906
set policy-options application-maps fcoe_p5_app_map application fcoe_p5_app code-points 101
set policy-options application-maps fcoe_all_app_map application fcoe_all_app code-points [011
101]
set protocols dcbx interface xe-0/0/21 application-map fcoe_p5_app_map
set protocols dcbx interface xe-0/0/22 application-map fcoe_all_app_map

```

### Step-by-Step Procedure

To configure two lossless forwarding classes for FCoE traffic on different interfaces, classify FCoE traffic into the forwarding classes, configure congestion notification profiles to enable PFC on the FCoE priorities and output queues, and configure DCBX application protocol TLV exchange for traffic on both FCoE priorities:

1. Configure lossless forwarding class **fcoe1** and map it to output queue **5** for FCoE traffic that uses IEEE 802.1p priority 5:

```

[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss

```



**NOTE:** This examples uses the default **fcoe** forwarding class as the other lossless FCoE forwarding class.

2. Configure the ingress classifier (**fcoe\_p3**) for interface **xe-0/0/20**. The classifier maps the FCoE priority (IEEE 802.1p code point **011**) to lossless FCoE forwarding class **fcoe**:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p3 forwarding-class fcoe loss-priority low code-points
011

```

3. Configure the ingress classifier (**fcoe\_p5**) for interface **xe-0/0/21**. The classifier maps the FCoE priority (IEEE 802.1p code point **101**) to lossless FCoE forwarding class **fcoe1**:

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p5 forwarding-class fcoe1 loss-priority low code-points
101

```



4. Configure the ingress classifier (**fcoe\_p3\_p5**) for interface **xe-0/0/22**. The classifier maps the two FCoE priorities (IEEE 802.1p code points **011** and **101**) to the two lossless FCoE forwarding classes **fcoe** and **fcoe1**, respectively:

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low code-points
011
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low code-points
101
```

5. Apply each classifier to the appropriate interface:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 unit 0 classifiers ieee-802.1 fcoe_p3
user@switch# set interfaces xe-0/0/21 unit 0 classifiers ieee-802.1 fcoe_p5
user@switch# set interfaces xe-0/0/22 unit 0 classifiers ieee-802.1 fcoe_p3_p5
```

6. Configure the CNP input stanza for interface **xe-0/0/20** to enable PFC on the FCoE priority (IEEE 802.1p code point **011**), set the MRU value (2240 bytes), and set the cable length value (100 meters). No output stanza is needed because queue 3 is paused by default on priority 3, and we are not explicitly configuring output queue flow control for any other queues.

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point
011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_cnp input cable-length 100
```

7. Configure the CNP for interface **xe-0/0/21**. The input stanza enables PFC on the FCoE priority (IEEE 802.1p code point **101**), sets the MRU value (2240 bytes), and sets the cable length value (150 meters). The output stanza configures flow control on output queue 5 on the FCoE priority:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p5_cnp input cable-length 150
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
```

8. Configure the CNP for interface **xe-0/0/22**. The input stanza enables PFC on the FCoE priorities (IEEE 802.1p code points **011** and **101**), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queues 3 and 5 on the FCoE priorities:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 011 pfc flow-control-queue 3
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 101 pfc flow-control-queue 5
```

9. Apply each CNP to the appropriate interface:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/20 congestion-notification-profile fcoe_p3_cnp
```

```

user@switch# set interfaces xe-0/0/21 congestion-notification-profile fcoe_p5_cnp
user@switch# set interfaces xe-0/0/22 congestion-notification-profile fcoe_p3_p5_cnp

```

10. Configure the DCBX FCoE application and application map to apply to interface xe-0/0/21. Interface xe-0/0/21 uses priority 5 (IEEE 802.1p code point 101) for FCoE traffic, which requires DCBX to exchange FCoE application protocol TLVs on priority 5 on interface xe-0/0/21. Configure an application named **fcoe\_p5\_app** for FCoE traffic (EtherType **0x8906**) and configure an application map named **fcoe\_p5\_app\_map** to map the application to code point 101:

```

[edit]
user@switch# set applications application fcoe_p5_app ether-type 0x8906
user@switch# set policy-options application-maps fcoe_p5_app_map application
fcoe_p5_app code-points 101

```



**NOTE:** Interface xe-0/0/20 uses the default FCoE configuration (priority 3). DCBX exchanges protocol TLVs for the FCoE application by default, so you do not need to configure DCBX explicitly on interface xe-0/0/20.

11. Configure the DCBX FCoE application and application map to apply to interface xe-0/0/22. Interface xe-0/0/22 uses both priority 3 (IEEE 802.1p code point 011) and priority 5 for FCoE traffic, which requires DCBX to exchange FCoE application protocol TLVs on both priority 3 and priority 5. Configure an application named **fcoe\_all\_app** for FCoE traffic (EtherType **0x8906**) and configure an application map named **fcoe\_all\_app\_map** to map the application to code points 011 and 101:

```

[edit]
user@switch# set applications application fcoe_all_app ether-type 0x8906
user@switch# set policy-options application-maps fcoe_all_app_map application
fcoe_all_app code-points [011 101]

```

12. Apply the application maps to the interfaces xe-0/0/21 and xe-0/0/22 so that DCBX exchanges FCoE application TLVs on the correct code points on each interface:

```

[edit]
user@switch# set protocols dcbx interface xe-0/0/21 application-map fcoe_p5_app_map
user@switch# set protocols dcbx interface xe-0/0/22 application-map fcoe_all_app_map

```

## Verification

To verify the configuration and proper operation of the lossless forwarding classes and IEEE 802.1p priorities, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 335](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 335](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 336](#)
- [Verifying the Interface Configuration on page 338](#)
- [Verifying the DCBX Application Configuration on page 338](#)
- [Verifying the DCBX Application Map Configuration on page 339](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 339](#)

### Verifying the Forwarding Class Configuration

**Purpose** Verify that the lossless forwarding class **fcoe1** has been created.

**Action** Show the forwarding class configuration by using the operational command **show class-of-service forwarding class**:

```
user@switch# show class-of-service forwarding-class
```

| Forwarding class | ID | Queue | Policing priority | No-Loss  |
|------------------|----|-------|-------------------|----------|
| best-effort      | 0  | 0     | normal            | Disabled |
| fcoe             | 1  | 3     | normal            | Enabled  |
| no-loss          | 2  | 4     | normal            | Enabled  |
| network-control  | 3  | 7     | normal            | Disabled |
| fcoe1            | 4  | 5     | normal            | Enabled  |
| mcast            | 8  | 8     | normal            | Disabled |

**Meaning** The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue **5** with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

### Verifying the Behavior Aggregate Classifier Configuration

**Purpose** Verify that the three classifiers map the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

**Action** List the classifiers configured to support lossless FCoE transport using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
```

Classifier: fcoe\_p3, Code point type: ieee-802.1, Index: 13913

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 011        | fcoe             | low           |

Classifier: fcoe\_p5, Code point type: ieee-802.1, Index: 63065

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 101        | fcoe1            | low           |

Classifier: fcoe\_p3\_p5, Code point type: ieee-802.1, Index: 10964

| Code point | Forwarding class | Loss priority |
|------------|------------------|---------------|
| 011        | fcoe             | low           |
| 101        | fcoe1            | low           |

**Meaning** The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier. The command output shows that there are three classifiers, **fcoe\_p3**, **fcoe\_p5**, and **fcoe\_p3\_p5**.

Classifier **fcoe\_p3** maps code point **011** (priority 3) to default lossless forwarding class **fcoe** and a packet loss priority of **low**.

Classifier **fcoe\_p5** maps code point **101** (priority 5) to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Classifier **fcoe\_p3\_p5** maps code point **011** to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and maps code point **101** to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

### Verifying the PFC Flow Control Configuration (CNP)

**Purpose** Verify that PFC is enabled on the correct input priorities and that flow control is configured on the correct output queues and priorities in each CNP.

**Action** List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Name: fcoe_p3_cnp, Index: 12037
```

```
Type: Input
```

```
Cable Length: 100 m
```

| Priority | PFC      | MRU  |
|----------|----------|------|
| 000      | Disabled |      |
| 001      | Disabled |      |
| 010      | Disabled |      |
| 011      | Enabled  | 2240 |
| 100      | Disabled |      |
| 101      | Disabled |      |
| 110      | Disabled |      |
| 111      | Disabled |      |

```
Type: Output
```

| Priority | Flow-Control-Queues |
|----------|---------------------|
| 000      |                     |
|          | 0                   |
| 001      |                     |
|          | 1                   |
| 010      |                     |
|          | 2                   |
| 011      |                     |
|          | 3                   |
| 100      |                     |
|          | 4                   |
| 101      |                     |
|          | 5                   |
| 110      |                     |
|          | 6                   |
| 111      |                     |
|          | 7                   |

```
Name: fcoe_p3_p5_cnp, Index: 46484
```

```
Type: Input
```

```
Cable Length: 100 m
```

```

Priority PFC MRU
000 Disabled
001 Disabled
010 Disabled
011 Enabled 2240
100 Disabled
101 Enabled 2240
110 Disabled
111 Disabled
Type: Output
Priority Flow-Control-Queues
011 3
101 5

Name: fcoe_p5_cnp, Index: 12133
Type: Input
Cable Length: 150 m
Priority PFC MRU
000 Disabled
001 Disabled
010 Disabled
011 Disabled
100 Disabled
101 Enabled 2240
110 Disabled
111 Disabled
Type: Output
Priority Flow-Control-Queues
101 5

```

**Meaning** The **show class-of-service congestion-notification** command shows the input and output stanzas of the three CNPs. For CNP **fcoe\_p3\_cnp**, the input stanza shows that PFC is enabled on IEEE 802.1p code point **011** (priority 3), the MRU is **2240** bytes, and the cable length is **100** meters. The CNP output stanza shows the default mapping of priorities to output queues.



**NOTE:** By default, only queues 3 and 4 are enabled to respond to pause messages from the connected peer. For queue 3 to respond to pause messages, priority 3 (code point 011) must be enabled for PFC in the input stanza. For queue 4 to respond to pause messages, priority 4 (code point 100) must be enabled for PFC in the input stanza. In this example, only queue 3 responds to pause messages from the connected peer on interfaces that use CNP **fcoe\_p3\_cnp**, because the input stanza enables PFC priority 3 only.

For CNP **fcoe\_p3\_p5\_cnp**, the input stanza shows that PFC is enabled on code points **011** and **101**, the MRU is **2240** bytes on both priorities, and the cable length is **100** meters. The CNP output stanza shows that output flow control is configured on queues **3** and **5** for code points **011** and **101**, respectively.

For CNP **fcoe\_p5\_cnp**, the input stanza shows that PFC is enabled on code point **101** (priority 5), the MRU is **2240** bytes, and the cable length is **150** meters. The CNP output stanza shows that output flow control is configured on queue **5** for code point **101** (priority 5).

---

### Verifying the Interface Configuration

**Purpose** Verify that the correct classifiers and congestion notification profiles are configured on the correct interfaces.

**Action** List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/20**, **show configuration class-of-service interfaces xe-0/0/21**, and **show configuration class-of-service interfaces xe-0/0/22**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/20
congestion-notification-profile fcoe_p3_cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe_p3;
 }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/21
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe_p5;
 }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/22
congestion-notification-profile fcoe_p3_p5_cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe_p3_p5;
 }
}
```

**Meaning** The **show configuration class-of-service interfaces xe-0/0/20** command shows that the congestion notification profile **fcoe\_p3\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_p3**.

The **show configuration class-of-service interfaces xe-0/0/21** command shows that the congestion notification profile **fcoe\_p5\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_p5**.

The **show configuration class-of-service interfaces xe-0/0/22** command shows that the congestion notification profile **fcoe\_p3\_p5\_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe\_p3\_p5**.

---

### Verifying the DCBX Application Configuration

**Purpose** Verify that the two DCBX applications for FCoE are configured.

**Action** List the DCBX applications by using the configuration mode command **show applications**:

```
user@switch# show applications
application fcoe_all_app {
 ether-type 0x8906;

application fcoe_p5_app {
 ether-type 0x8906;
```

**Meaning** The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **fcoe\_all\_app** is configured with an EtherType of **0x8906** (the correct EtherType for FCoE traffic) and that the application **fcoe\_p5\_app** is also configured with an EtherType of **0x8906**.

### Verifying the DCBX Application Map Configuration

**Purpose** Verify that the application maps are configured.

**Action** List the application maps by using the configuration mode command **show policy-options application-maps**:

```
user@switch# show policy-options application-maps
fcoe_all_app_map {
 application fcoe_all_app code-points [011 101];
}
fcoe_p5_app_map {
 application fcoe_p5_app code-points 101;
}
```

**Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that there are two application maps.

Application map **fcoe\_all\_app\_map** consists of the application named **fcoe\_all\_app** mapped to IEEE 802.1p code points **011** (priority 3) and **101** (priority 5).

Application map **fcoe\_p5\_app\_map** consists of the application named **fcoe\_p5\_app** mapped to IEEE 802.1p code point **101** (priority 5).

### Verifying the DCBX Application Protocol Exchange Interface Configuration

**Purpose** Verify that the application maps are applied to the correct interfaces.

**Action** List the application maps on each interface using the configuration mode command **show protocols dcbx**:

```
user@switch# show protocols dcbx
interface xe-0/0/21.0 {
 application-map fcoe_p5_app_map;
}
interface xe-0/0/22.0 {
 application-map fcoe_all_app_map;
}
```

**Meaning** The `show protocols dcbx` configuration mode command lists the application map association with interfaces. The output shows that interface `xe-0/0/21.0` uses application map `fcoe_p5_app_map` and interface `xe-0/0/22.0` uses application map `fcoe_all_app_map`.



**NOTE:** Because interface `xe-0/0/20` uses the default lossless FCoE configuration, you do not configure application mapping to interface `xe-0/0/20`. The default configuration automatically exchanges application protocol TLVs for the default FCoE configuration on priority 3 (IEEE 802.1p code point 011).

**Related Documentation**

- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 317](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 309](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 340](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)
- [Example: Configuring Unicast Classifiers on page 273](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)

---

## Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI)

---

Although the default configuration provides two lossless forwarding classes mapped to two different IEEE 802.1p priorities (code points), you can explicitly configure up to six lossless forwarding classes and map them to different priorities. You can support up to six different types of lossless traffic, and you can support the same type of traffic on different priorities in different parts of your converged network.

This example shows you how to configure two lossless forwarding classes for FCoE traffic and one lossless forwarding class for iSCSI traffic, and map the forwarding classes to three different priorities. (The converged Ethernet network includes two FCoE networks, each of which uses a different priority to identify FCoE traffic, and an iSCSI network.)

- [Requirements on page 341](#)
- [Overview on page 341](#)
- [Configuration on page 345](#)
- [Verification on page 349](#)



## Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch in transit switch (FIP snooping) mode
- Junos OS Release 12.3 or later for the QFX Series

## Overview

Some converged Ethernet networks support FCoE on more than one IEEE 802.1p priority and also require supporting other lossless traffic classes. Interfaces that carry multiple lossless forwarding classes need to support lossless behavior for the priorities mapped to those forwarding classes. To support the two FCoE forwarding classes and the iSCSI forwarding class used in this example, you need to configure:

- At least one lossless forwarding class for FCoE traffic (this example uses the default **fcoe** forwarding class as one of the two lossless FCoE forwarding classes, so we need to explicitly configure only one FCoE forwarding class)
- A lossless forwarding class for iSCSI traffic
- Behavior aggregate (BA) classifiers to map the lossless forwarding classes to the appropriate IEEE 802.1p code points (priorities) on each interface
- Congestion notification profiles (CNPs) for each interface to enable PFC on the FCoE and iSCSI code points at the interface ingress, and to configure PFC flow control on the interface egress so that the interface can respond to PFC messages received from the connected peer



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- DCBX applications and an application map to support DCBX application TLV exchange for the FCoE and iSCSI traffic on the configured lossless priorities. By default, DCBX is enabled on all Ethernet interfaces for FCoE, but only on priority 3 (IEEE 802.1p code point 011). To support DCBX application TLV exchange when you are not using the default configuration, you must configure all of the applications and map them to interfaces and priorities.

The priorities specified in the BA classifiers, CNPs, and DCBX application map must match, or the configuration does not work. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.


## Topology

This example shows how to configure two lossless FCoE traffic classes and one lossless iSCSI traffic class, map them to three different priorities, and configure flow control to ensure lossless behavior for those priorities on the interfaces. This example uses four Ethernet interfaces, xe-0/0/31, xe-0/0/32, xe-0/0/33, and xe-0/0/34:

- Interface xe-0/0/31 handles FCoE traffic on priority 3 (IEEE 802.1p code point 011) and iSCSI traffic on priority 4 (code point 100).
- Interface xe-0/0/32 handles FCoE traffic on priority 5 (code point 101) and iSCSI traffic on priority 4.
- Interface xe-0/0/33 handles FCoE traffic on priority 3 and priority 5.
- Interface xe-0/0/34 handles iSCSI traffic on priority 4.

Figure 16 on page 342 shows the topology for this example, and Table 99 on page 342 shows the configuration components for this example.

**Figure 16: Topology of the Lossless FCoE and iSCSI Priorities Example**



**Table 99: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology**

| Component | Settings       |
|-----------|----------------|
| Hardware  | QFX3500 switch |

**Table 99: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology (*continued*)**

| Component          | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding classes | <p>This example uses one explicitly configured lossless FCoE forwarding class, the default lossless FCoE forwarding class, and one explicitly configured iSCSI forwarding class.</p> <ul style="list-style-type: none"> <li>iSCSI forwarding class:<br/>Name—<b>iscsi</b><br/>Queue mapping—queue 4<br/>Packet drop attribute—<b>no-loss</b></li> <li>FCoE forwarding class (explicitly configured):<br/>Name—<b>fcoe1</b><br/>Queue mapping—queue 5<br/>Packet drop attribute—<b>no-loss</b></li> </ul> <p><b>NOTE:</b> A lossless forwarding class can be mapped to any output queue. However, because the <b>fcoe1</b> forwarding class uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.</p> <ul style="list-style-type: none"> <li>FCoE forwarding class (default)<br/>Name—<b>fcoe</b><br/>The default <b>fcoe</b> forwarding class is mapped to priority 3 (IEEE 802.1p code point 011) and to output queue 3 with a packet drop attribute of <b>no-loss</b>.</li> </ul>                                                                                                                                                                                                                                                                                                                                   |
| BA classifiers     | <p>Each interface requires a different classifier because each interface handles a different subset of FCoE traffic.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/31 classifier:<br/>Name—<b>fcoe_p3_iscsi</b><br/>FCoE priority mapping—Forwarding class <b>fcoe</b> mapped to code point <b>011</b> (IEEE 802.1p priority 3) and a packet loss priority of <b>low</b>.<br/>iSCSI priority mapping—Forwarding class <b>iscsi</b> mapped to code point <b>100</b> (priority 4) and a packet loss priority of <b>low</b>.</li> <li>Interface xe-0/0/32 classifier:<br/>Name—<b>fcoe_p5_iscsi</b><br/>FCoE priority mapping—Forwarding class <b>fcoe1</b> mapped to code point <b>101</b> (IEEE 802.1p priority 5) and a packet loss priority of <b>low</b>.<br/>iSCSI priority mapping—Forwarding class <b>iscsi</b> mapped to code point <b>100</b> (priority 4) and a packet loss priority of <b>low</b>.</li> <li>Interface xe-0/0/33 classifier:<br/>Name—<b>fcoe_p3_p5</b><br/>FCoE priority mapping—Forwarding class <b>fcoe1</b> mapped to code point <b>101</b> (priority 5) and a packet loss priority of <b>low</b>, and forwarding class <b>fcoe</b> mapped to code point <b>011</b> and a packet loss priority of <b>low</b>.</li> <li>Interface xe-0/0/34 classifier:<br/>Name—<b>iscsi_classifier</b><br/>iSCSI priority mapping—Forwarding class <b>iscsi</b> mapped to code point <b>100</b> (priority 4) and a packet loss priority of <b>low</b>.</li> </ul> |

**Table 99: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology (*continued*)**

| Component                | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PFC configuration (CNPs) | <p>Each interface requires a different CNP because each interface handles a different subset of FCoE and iSCSI traffic, and must pause that traffic on different priorities.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/31 CNP:<br/> CNP name—<b>fcoe_p3_cnp</b><br/> Input CNP code points—<b>011</b> and <b>100</b><br/> MRU—2240 bytes for code point <b>011</b>, default value (2500 bytes) for code point <b>100</b><br/> Cable length—100 meters</li> </ul> <p><b>NOTE:</b> On interface xe-0/0/31, the FCoE forwarding class is mapped to queue 3 and priority 3 (code point 011), and the iSCSI forwarding class is mapped to queue 4 and priority 4 (code point 100). Therefore, interface xe-0/0/31 does not require an output CNP configuration because queue 3 and queue 4 are enabled for PFC flow control by default on code points 011 and 100, respectively.</p> <ul style="list-style-type: none"> <li>Interface xe-0/0/32 CNP:<br/> CNP name—<b>fcoe_p5_cnp</b><br/> Input CNP code points—<b>100</b> and <b>101</b><br/> MRU—Default value (2500 bytes) for code point <b>100</b>, <b>2240</b> bytes for code point <b>101</b><br/> Cable length—150 meters<br/> Output CNP code points—<b>100</b> and <b>101</b><br/> Output CNP flow control queues—<b>4</b> and <b>5</b></li> <li>Interface xe-0/0/33 CNP:<br/> CNP name—<b>fcoe_p3_p5_cnp</b><br/> Input CNP code points—<b>011</b> and <b>101</b><br/> MRU—<b>2240</b> bytes (both priorities)<br/> Cable length—100 meters<br/> Output CNP code points—<b>011</b> and <b>101</b><br/> Output CNP flow control queues—<b>3</b> and <b>5</b></li> <li>Interface xe-0/0/34 CNP:<br/> CNP name—<b>iscsi_cnp</b><br/> Input CNP code point—<b>100</b><br/> MRU—<b>2500</b> bytes (default value)<br/> Cable length—100 meters</li> </ul> <p><b>NOTE:</b> On interface xe-0/0/34, the iSCSI forwarding class is mapped to queue 4 and priority 4 (code point 100). Interface xe-0/0/34 does not require an output CNP configuration because queue 4 is enabled for PFC flow control by default on code point 100.</p> <p><b>NOTE:</b> When you apply a CNP with an explicit output queue flow control configuration to an interface, the explicit CNP overwrites the default output CNP. The output queues that are enabled for PFC pause in the default configuration (queues 3 and 4) are not enabled for pause unless they are included in the explicitly configured output CNP.</p> |

**Table 99: Components of the Lossless FCoE and iSCSI Priorities Configuration Topology (*continued*)**

| Component                | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| DCBX application mapping | <p>This example requires configuring applications for FCoE and iSCSI, including them in the same application map, and applying the application map to all four interfaces.</p> <p>Application map name—<b>dcbx_iscsi_fcoe_app_map</b></p> <ul style="list-style-type: none"> <li>FCoE application name—<b>fcoe_app</b><br/>Application ether-type—<b>0x8906</b><br/>Application map code points—<b>011</b> and <b>101</b></li> <li>iSCSI application name—<b>iscsi_app</b><br/>Application protocol type—<b>tcp</b><br/>Application destination port—<b>3260</b><br/>Application map code point—<b>100</b></li> </ul> <p><b>NOTE:</b> LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.</p> |



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or the FIP snooping configuration. This examples focuses only on the lossless FCoE priority configuration.

## Configuration

### CLI Quick Configuration

To quickly configure two lossless FCoE forwarding classes and one lossless iSCSI forwarding class and map them to different priorities, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set class-of-service forwarding-classes class iscsi queue-num 4 no-loss
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_p3_iscsi forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p3_iscsi forwarding-class iscsi loss-priority low
code-points 100
set class-of-service classifiers ieee-802.1 fcoe_p5_iscsi forwarding-class iscsi loss-priority low
code-points 100
set class-of-service classifiers ieee-802.1 fcoe_p5_iscsi forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low
code-points 011
set class-of-service classifiers ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low
code-points 101
set class-of-service classifiers ieee-802.1 iscsi_classifier forwarding-class iscsi loss-priority low
code-points 100
set class-of-service interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe_p3_iscsi
set class-of-service interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe_p5_iscsi
set class-of-service interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe_p3_p5
set class-of-service interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 iscsi_classifier
```

```

set class-of-service congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point 011
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point 100
pfc
set class-of-service congestion-notification-profile fcoe_p3_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 100
pfc
set class-of-service congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point 101
pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p5_cnp input cable-length 150
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
100 pfc flow-control-queue 4
set class-of-service congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
011 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
011 pfc flow-control-queue 3
set class-of-service congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
set class-of-service congestion-notification-profile iscsi_cnp input ieee-802.1 code-point 100 pfc
set class-of-service congestion-notification-profile iscsi_cnp input cable-length 100
set class-of-service interfaces xe-0/0/31 congestion-notification-profile fcoe_p3_cnp
set class-of-service interfaces xe-0/0/32 congestion-notification-profile fcoe_p5_cnp
set class-of-service interfaces xe-0/0/33 congestion-notification-profile fcoe_p3_p5_cnp
set class-of-service interfaces xe-0/0/34 congestion-notification-profile iscsi_cnp
set applications application iscsi_app protocol tcp destination-port 3260
set applications application fcoe_app ether-type 0x8906
set policy-options application-maps dcbx_iscsi_fcoe_app_map application iscsi_app code-points
100
set policy-options application-maps dcbx_iscsi_fcoe_app_map application fcoe_app code-points
[011 101]
set protocols dcbx interface xe-0/0/31 application-map dcbx_iscsi_fcoe_app_map
set protocols dcbx interface xe-0/0/32 application-map dcbx_iscsi_fcoe_app_map
set protocols dcbx interface xe-0/0/33 application-map dcbx_iscsi_fcoe_app_map
set protocols dcbx interface xe-0/0/34 application-map dcbx_iscsi_fcoe_app_map

```

### Step-by-Step Procedure

To configure two lossless forwarding classes for FCoE traffic and one lossless forwarding class for iSCSI traffic, classify the traffic into the three forwarding classes, configure congestion notification profiles to enable PFC on the FCoE priorities and output queues, and configure DCBX application protocol TLV exchange for traffic on both FCoE priorities:

1. Configure lossless forwarding classes **iscsi** for iSCSI traffic and **fcoe1** for FCoE traffic (this example uses the default **fcoe** forwarding class as the other lossless FCoE forwarding class) and map them to output queues:

```

[edit class-of-service]
user@switch# set forwarding-classes class iscsi queue-num 4 no-loss
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss

```

2. Configure the ingress classifier (**fcoe\_p3\_iscsi**) for interface **xe-0/0/31**. The classifier maps the FCoE priority (code point **011**) to lossless FCoE forwarding class **fcoe** and the iSCSI priority (code point **100**) to lossless iSCSI forwarding class **iscsi**:

```

[edit class-of-service classifiers]

```

- ```

user@switch# set ieee-802.1 fcoe_p3_iscsi forwarding-class fcoe loss-priority low
code-points 011
user@switch# set ieee-802.1 fcoe_p3_iscsi forwarding-class iscsi loss-priority low
code-points 100

```
3. Configure the ingress classifier (**fcoe_p5_iscsi**) for interface **xe-0/0/32**. The classifier maps the FCoE priority (code point **101**) to lossless FCoE forwarding class **fcoe1** and the iSCSI priority (code point **100**) to lossless iSCSI forwarding class **iscsi**:


```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p5_iscsi forwarding-class iscsi loss-priority low
code-points 100
user@switch# set ieee-802.1 fcoe_p5_iscsi forwarding-class fcoe1 loss-priority low
code-points 101

```
 4. Configure the ingress classifier (**fcoe_p3_p5**) for interface **xe-0/0/33**. The classifier maps the two FCoE priorities (code points **011** and **101**) to lossless FCoE forwarding classes **fcoe** and **fcoe1**, respectively:


```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe loss-priority low code-points
011
user@switch# set ieee-802.1 fcoe_p3_p5 forwarding-class fcoe1 loss-priority low code-points
101

```
 5. Configure the ingress classifier (**iscsi_classifier**) for interface **xe-0/0/34**. The classifier maps the iSCSI priority (code point **101**) to lossless iSCSI forwarding class **iscsi**:


```

[edit class-of-service classifiers]
user@switch# set ieee-802.1 iscsi_classifier forwarding-class iscsi loss-priority low
code-points 100

```
 6. Apply each classifier to the appropriate interface:


```

[edit class-of-service]
user@switch# set interfaces xe-0/0/31 unit 0 classifiers ieee-802.1 fcoe_p3_iscsi
user@switch# set interfaces xe-0/0/32 unit 0 classifiers ieee-802.1 fcoe_p5_iscsi
user@switch# set interfaces xe-0/0/33 unit 0 classifiers ieee-802.1 fcoe_p3_p5
user@switch# set interfaces xe-0/0/34 unit 0 classifiers ieee-802.1 iscsi_classifier

```
 7. Configure the CNP input stanza for interface **xe-0/0/31** to enable PFC on the FCoE and iSCSI priorities that the interface handles (code points **011** and **100**), set the MRU value for the FCoE traffic (2240 bytes), and set the cable length value (100 meters). No output stanza is needed because queues 3 and 4 are paused by default on priorities 3 and 4, respectively, and we are not explicitly configuring output queue flow control for any other queues.


```

[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point
011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_cnp input ieee-802.1 code-point
100 pfc
user@switch# set congestion-notification-profile fcoe_p3_cnp input cable-length 100

```
 8. Configure the CNP for interface **xe-0/0/32**. The input stanza enables PFC on the FCoE priority (code point **101**), sets the MRU value for FCoE traffic (2240 bytes), enables PFC on the iSCSI priority (code point **100**), and sets the cable length value (150 meters). The output stanza configures flow control on output queue 5 on the FCoE priority and on output queue 4 on the iSCSI priority:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
100 pfc
user@switch# set congestion-notification-profile fcoe_p5_cnp input ieee-802.1 code-point
101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p5_cnp input cable-length 150
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
100 pfc flow-control-queue 4
user@switch# set congestion-notification-profile fcoe_p5_cnp output ieee-802.1 code-point
101 pfc flow-control-queue 5
```

9. Configure the CNP for interface xe-0/0/33. The input stanza enables PFC on the FCoE priorities (IEEE 802.1p code points 011 and 101), sets the MRU value (2240 bytes), and sets the cable length value (100 meters). The output stanza configures flow control on output queues 3 and 5 on the FCoE priorities:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 011 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input ieee-802.1
code-point 101 pfc mru 2240
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp input cable-length 100
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 011 pfc flow-control-queue 3
user@switch# set congestion-notification-profile fcoe_p3_p5_cnp output ieee-802.1
code-point 101 pfc flow-control-queue 5
```

10. Configure the CNP input stanza for interface xe-0/0/34 to enable PFC on the iSCSI priority (code point 100) and set the cable length value (100 meters). No output stanza is needed because queue 4 is paused by default on priority 4, and we are not explicitly configuring output queue flow control for any other queues.

```
[edit class-of-service]
user@switch# set congestion-notification-profile iscsi_cnp input ieee-802.1 code-point
100 pfc
user@switch# set congestion-notification-profile iscsi_cnp input cable-length 100
```

11. Apply each CNP to the appropriate interface:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/31 congestion-notification-profile fcoe_p3_cnp
user@switch# set interfaces xe-0/0/32 congestion-notification-profile fcoe_p5_cnp
user@switch# set interfaces xe-0/0/33 congestion-notification-profile fcoe_p3_p5_cnp
user@switch# set interfaces xe-0/0/34 congestion-notification-profile iscsi_cnp
```

12. Configure the DCBX applications for FCoE and iSCSI to map to the interfaces so that DCBX can exchange application protocol TLVs on the IEEE 802.1p priorities used for FCoE and iSCSI traffic:

```
[edit]
user@switch# set applications application fcoe_app ether-type 0x8906
user@switch# set applications application iscsi_app protocol tcp destination-port 3260
```

13. Configure a DCBX application map to map the FCoE and iSCSI applications to the correct priorities:

```
[edit]
user@switch# set policy-options application-maps dcbx_iscsi_fcoe_app_map application
fcoe_app code-points [011 101]
user@switch# set policy-options application-maps dcbx_iscsi_fcoe_app_map application
iscsi_app code-points 100
```


14. Apply the application map to the interfaces so that DCBX exchanges FCoE application TLVs on the correct code points:

```
[edit]
user@switch# set protocols dcbx interface xe-0/0/31 application-map
dcbx_iscsi_fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/32 application-map
dcbx_iscsi_fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/33 application-map
dcbx_iscsi_fcoe_app_map
user@switch# set protocols dcbx interface xe-0/0/34 application-map
dcbx_iscsi_fcoe_app_map
```

Verification

To verify the configuration and proper operation of the lossless forwarding classes and IEEE 802.1p priorities, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 349](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 350](#)
- [Verifying the PFC Flow Control Configuration \(CNP\) on page 351](#)
- [Verifying the Interface Configuration on page 353](#)
- [Verifying the DCBX Application Configuration on page 354](#)
- [Verifying the DCBX Application Map Configuration on page 355](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 355](#)

Verifying the Forwarding Class Configuration

Purpose Verify that the lossless forwarding classes **iscsi** and **fcoe1** have been created and that the default lossless forwarding class **fcoe** is still enabled for lossless transport.

Action Show the forwarding class configuration by using the operational command **show class-of-service forwarding-class**:

```
user@switch> show class-of-service forwarding-class
```

Forwarding class	ID	Queue	Policing priority	No-Loss
best-effort	0	0	normal	Disabled
fcoe	1	3	normal	Enabled
iscsi	2	4	normal	Enabled
network-control	3	7	normal	Disabled
fcoe1	4	5	normal	Enabled
mcast	8	8	normal	Disabled

Meaning The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **iscsi** and **fcoe1** forwarding classes are configured on output queues 4 and 5, respectively, with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default **fcoe** forwarding class, it remains in its default state (lossless configuration).

Verifying the Behavior Aggregate Classifier Configuration

Purpose Verify that the four classifiers map the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

Action List the classifiers configured to support lossless FCoE transport using the operational mode command **show class-of-service classifier**:

```
user@switch> show class-of-service classifier
Classifier: fcoe_p3_iscsi, Code point type: ieee-802.1, Index: 13915
  Code point    Forwarding class    Loss priority
  011           fcoe                low
  100           iscsi              low

Classifier: fcoe_p5_iscsi, Code point type: ieee-802.1, Index: 62035
  Code point    Forwarding class    Loss priority
  100           iscsi              low
  101           fcoe1             low

Classifier: fcoe_p3_p5, Code point type: ieee-802.1, Index: 17774
  Code point    Forwarding class    Loss priority
  011           fcoe                low
  101           fcoe1             low

Classifier: iscsi_classifier, Code point type: ieee-802.1, Index: 31635
  Code point    Forwarding class    Loss priority
  100           iscsi              low
```

Meaning The **show class-of-service classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in each classifier. The command output shows that there are four classifiers, **fcoe_p3_iscsi**, **fcoe_p5_iscsi**, **fcoe_p3_p5**, and **iscsi_classifier**.

Classifier **fcoe_p3_iscsi** maps code point **011** (priority 3) to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and code point **100** (priority 4) to explicitly configured lossless forwarding class **iscsi**.

Classifier **fcoe_p5_iscsi** maps code point **100** to explicitly configured forwarding class **iscsi** and a packet loss priority of **low**, and code point **101** (priority 5) to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Classifier **fcoe_p3_p5** maps code point **011** to default lossless forwarding class **fcoe** and a packet loss priority of **low**, and maps code point **101** to explicitly configured lossless forwarding class **fcoe1** and a packet loss priority of **low**.

Classifier **iscsi_classifier** maps code point **100** to explicitly configured forwarding class **iscsi** and a packet loss priority of **low**.

Verifying the PFC Flow Control Configuration (CNP)

Purpose Verify that PFC is enabled on the correct input priorities and that flow control is configured on the correct output queues and priorities in each CNP.

Action List the congestion notification profiles using the operational mode command **show class-of-service congestion-notification**:

```
user@switch> show class-of-service congestion-notification
```

```
Name: fcoe_p3_cnp, Index: 12037
```

```
Type: Input
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240
100	Enabled	9216
101	Disabled	
110	Disabled	
111	Disabled	

```
Type: Output
```

Priority	Flow-Control-Queues
000	
001	0
010	1
011	2
100	3
101	4
110	5
111	6
	7

```
Name: fcoe_p3_p5_cnp, Index: 46484
```

```
Type: Input
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240
100	Disabled	
101	Enabled	2240
110	Disabled	
111	Disabled	

```
Type: Output
```

Priority	Flow-Control-Queues
011	
	3
101	
	5

```
Name: fcoe_p5_cnp, Index: 12133
Type: Input
Cable Length: 150 m
  Priority    PFC      MRU
  000        Disabled
  001        Disabled
  010        Disabled
  011        Disabled
  100        Enabled   9216
  101        Enabled   2240
  110        Disabled
  111        Disabled
Type: Output
  100
      4
  101
      5

Name: iscsi_cnp, Index: 19342
Type: Input
Cable Length: 100 m
  Priority    PFC      MRU
  000        Disabled
  001        Disabled
  010        Disabled
  011        Disabled
  100        Enabled   9216
  101        Disabled
  110        Disabled
  111        Disabled
Type: Output
  Priority    Flow-Control-Queues
  000
      0
  001
      1
  010
      2
  011
      3
  100
      4
  101
      5
  110
      6
  111
      7
```

Meaning The **show class-of-service congestion-notification** command shows the input and output stanzas of the four CNPs.

For CNP **fcoe_p3_cnp**, the input stanza shows that PFC is enabled on IEEE 802.1p code point **011** (priority 3) with an MRU of **2240** bytes, and cable length of **100** meters. The input stanza also shows that PFC is enabled on code point **100** (priority 4) with the default MRU value of **9216** bytes. The CNP output stanza shows the default mapping of priorities to output queues because no explicit output CNP is configured.



NOTE: By default, only queues 3 and 4 are enabled respond to pause messages from the connected peer. For queue 3 to respond to pause messages, priority 3 (code point 011) must be enabled for PFC in the input stanza. For queue 4 to respond to pause messages, priority 4 (code point 100) must be enabled for PFC in the input stanza. In this example, only queues 3 and 4 respond to pause messages from the connected peer on interfaces that use CNP `fcoe_p3_cnp` because the input stanza enables PFC only on priorities 3 and 4.

For CNP `fcoe_p3_p5_cnp`, the input stanza shows that PFC is enabled on code points 011 and 101 (priority 5), the MRU is 2240 bytes on both priorities, and the cable length is 100 meters. The CNP output stanza shows that output flow control is configured on queues 3 and 5 for code points 011 and 101, respectively.

For CNP `fcoe_p5_cnp`, the input stanza shows that PFC is enabled on code points 100 and 101. The MRU for code point 101 (FCoE traffic) is 2240 bytes and the MRU for code point 100 is 9216. The interface cable length is 150 meters. The CNP output stanza shows that output flow control is configured on queue 4 for code point 100 and on queue 5 for code point 101.

For CNP `iscsi_cnp`, the input stanza shows that PFC is enabled on code point 100, the MRU value is 9216 bytes, and the interface cable length is 100 meters. The CNP output stanza shows the default mapping of priorities to output queues because no explicit output CNP is configured.

Verifying the Interface Configuration

Purpose Verify that the correct classifiers and congestion notification profiles are configured on the correct interfaces.

Action List the ingress interfaces using the operational mode commands **show configuration class-of-service interfaces xe-0/0/31**, **show configuration class-of-service interfaces xe-0/0/32**, **show configuration class-of-service interfaces xe-0/0/33**, and **show configuration class-of-service interfaces xe-0/0/34**:

```
user@switch> show configuration class-of-service interfaces xe-0/0/31
congestion-notification-profile fcoe_p3_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_p3_iscsi;
    }
}

user@switch> show configuration class-of-service interfaces xe-0/0/32
congestion-notification-profile fcoe_p5_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_p5_iscsi;
    }
}
```

```
user@switch> show configuration class-of-service interfaces xe-0/0/33
congestion-notification-profile fcoe_p3_p5_cnp;
unit 0 {
    classifiers {
        ieee-802.1 fcoe_p3_p5;
    }
}

user@switch> show configuration class-of-service interfaces xe-0/0/34
congestion-notification-profile iscsi_cnp;
unit 0 {
    classifiers {
        ieee-802.1 iscsi_classifier;
    }
}
```

Meaning The **show configuration class-of-service interfaces xe-0/0/31** command shows that the congestion notification profile **fcoe_p3_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe_p3_iscsi**.

The **show configuration class-of-service interfaces xe-0/0/32** command shows that the congestion notification profile **fcoe_p5_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe_p5_iscsi**.

The **show configuration class-of-service interfaces xe-0/0/33** command shows that the congestion notification profile **fcoe_p3_p5_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **fcoe_p3_p5**.

The **show configuration class-of-service interfaces xe-0/0/34** command shows that the congestion notification profile **iscsi_cnp** is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is **iscsi_classifier**.

Verifying the DCBX Application Configuration

Purpose Verify that the DCBX applications for FCoE and iSCSI are configured.

Action List the DCBX applications by using the configuration mode command **show applications**:

```
user@switch# show applications
application iscsi_app {
    protocol tcp;
    destination-port 3260;
}
application fcoe_app {
    ether-type 0x8906;
```

Meaning The **show applications** configuration mode command shows all of the configured applications. The output shows that the application **iscsi_app** is configured with a protocol value of **tcp** and a destination port value of **3260**, and that the application **fcoe_app** is configured with an EtherType of **0x8906** (the correct EtherType for FCoE traffic).

Verifying the DCBX Application Map Configuration

- Purpose** Verify that the application map is configured.
- Action** List the application maps by using the configuration mode command **show policy-options application-maps**:
- ```
user@switch# show policy-options application-maps
dcbx-iscsi-fcoe-app-map {
 application iscsi_app code-points 100;
 application fcoe_app code-points [011 101];
}
```
- Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that there is one application map named **dcbx-iscsi-fcoe\_app\_map**. It consists of the application **iscsi\_app** mapped to code point **100** and the application **fcoe\_app** mapped to code points **011** and **101**.

### Verifying the DCBX Application Protocol Exchange Interface Configuration

- Purpose** Verify that the application maps are applied to the correct interfaces.
- Action** List the application maps on each interface using the configuration mode command **show protocols dcbx**:
- ```
user@switch# show protocols dcbx
interface xe-0/0/31.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}
interface xe-0/0/32.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}
interface xe-0/0/33.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}
interface xe-0/0/34.0 {
    application-map dcbx-iscsi-fcoe-app-map;
}
```
- Meaning** The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that all four interfaces use the application map **dcbx-iscsi-fcoe-app-map**.
- Related Documentation**
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 317](#)
 - [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 309](#)
 - [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326](#)
 - [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)

- [Example: Configuring Unicast Classifiers on page 273](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)

Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway

FCoE traffic typically uses IEEE 802.1p priority 3 (code point 011). However, if your FCoE network uses a different IEEE 802.1p priority than priority 3 for FCoE traffic, then you can use priority remapping to classify FCoE traffic into a lossless forwarding class mapped to that priority. You specify the lossless forwarding class used for the FCoE traffic by configuring a fixed classifier and applying it to the native FC (NP_Port) interface. All traffic received from the FC SAN on that NP_Port interface is classified into the forwarding class specified in the fixed classifier.

When native FC interfaces on the FCoE-FC gateway encapsulate incoming FC traffic in Ethernet to create FCoE frames, by default they assign IEEE 802.1p code point 011 to the FCoE traffic, forward the traffic internally to the gateway Ethernet interfaces, and then forward the traffic to the FCoE network. Setting a rewrite value for the IEEE 802.1p code point configures the gateway native FC interface to assign the rewrite value priority to the FCoE frames when the native FC interface forwards the FCoE frames to the gateway Ethernet interface. Instead of a priority of 3, the FCoE frames use the priority specified in the rewrite value.

You can configure one rewrite value for each local FCoE-FC gateway fabric. All of the native FC interfaces in a particular fabric must use the same rewrite value. Native FC interfaces that belong to different FCoE-FC gateway fabrics can use different rewrite values.

This example shows how to configure FCoE priority remapping for a converged Ethernet network that uses priority 5 (IEEE code point 101) for FCoE traffic. If your network uses priority 3 for FCoE traffic, then you do not need to remap the FCoE priority, because the default configuration supports lossless FCoE transport on priority 3.

- [Requirements on page 356](#)
- [Overview on page 357](#)
- [Configuration on page 360](#)
- [Verification on page 362](#)

Requirements

This example uses the following hardware and software components:

- One Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

Overview

Native FC interfaces on an FCoE-FC gateway receive native FC traffic from the FC SAN and encapsulate it in Ethernet to create FCoE frames. Priority remapping enables you to map the encapsulated FC traffic (the FCoE traffic) to any IEEE 802.1p priority.

To support lossless FCoE traffic flows, you must configure the remapped priority correctly on the native FC interfaces and also on the Ethernet interfaces that connect to the FCoE network. Achieving lossless behavior for FCoE traffic when you remap the FCoE priority requires configuring:

- A lossless forwarding class for FCoE traffic (or using the default **fcoe** forwarding class)
- A behavior aggregate (BA) classifier on the FCoE Ethernet interfaces to map the FCoE forwarding class to the IEEE 802.1p code points (priority) used for FCoE traffic on the FCoE network (the ingress classifier priority for the forwarding class must be the same as the rewrite value priority)
- A fixed classifier on the FCoE-FC gateway FC interface that maps all traffic from the FC network into the lossless FCoE forwarding class (the forwarding class must be lossless)
- A priority rewrite value that remaps the IEEE 802.1p code point on the FCoE-FC gateway FC interface to the priority used for FCoE traffic on the FCoE network
- An input congestion notification profile (CNP) to enable priority-based flow control (PFC) on the FCoE code point (the code point used as the rewrite value) at the Ethernet interface ingress and an output CNP to configure flow control to pause the correct output queue at the Ethernet interface egress



NOTE: Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

- A DCBX application and application map on the Ethernet interface to support DCBX application TLV exchange for the lossless FCoE traffic on the FCoE priority

The priority specified in the BA classifier, CNP, and DCBX application map on the Ethernet ingress interfaces must match the priority specified in the fixed classifier and rewrite value configurations on the FC interfaces. You must specify the same lossless FCoE forwarding class in each configuration and use the same IEEE 802.1p code point (priority) so that the FCoE traffic is properly classified into flows and so that those flows receive lossless treatment.

Topology

This example shows how to configure priority remapping of FCoE traffic on one native FC interface (fc-0/0/2) connected to the FC SAN and on one Ethernet interface (xe-0/0/27) connected to the converged Ethernet (FCoE) network. Both the native FC

interface and the Ethernet interface belong to the same local FC fabric on the FCoE-FC gateway.

The converged Ethernet network uses priority 5 (IEEE 802.1p code point 101) for FCoE traffic. The native FC interface on the FCoE-FC gateway receives FC traffic from the FC SAN. The native FC interface encapsulates the FC traffic in Ethernet to create FCoE frames, tags the frames with the IEEE 802.1p priority value 101, and then forwards the FCoE frames to the FCoE-FC gateway Ethernet interface. Because traffic marked with IEEE 802.1p priority 5 is mapped to a lossless FCoE forwarding class, the traffic receives lossless treatment. The Ethernet interface forwards the FCoE traffic on to the Ethernet network.

FCoE traffic (tagged with priority 5) arriving at the FCoE-FC gateway from the Ethernet network receives lossless treatment and is forwarded to the native FC interface. The native FC interface removes the Ethernet encapsulation from the FCoE frames and forwards the resulting native FC traffic to the FC SAN.

Figure 17 on page 358 shows the topology for this example, and Table 100 on page 358 shows the configuration components for this example.

Figure 17: Topology of the IEEE 802.1p Priority Remapping Example




Table 100: Components of the IEEE 802.1p Priority Remapping Configuration Topology

Component	Settings
Hardware	QFX3500 switch
Forwarding class configuration	Name— fcoe1 Queue mapping—queue 5 Packet drop attribute— no-loss NOTE: The lossless forwarding class can be mapped to any output queue. However, because FCoE uses priority 5 in this example, matching that traffic to a forwarding class that uses queue 5 creates a configuration that is logical and easy to map because the priority and the queue are identified by the same number.
BA classifier (Ethernet interface)	Name— fcoe_gw_classifier Maps code point 101 (IEEE 802.1p priority 5) to the fcoe1 forwarding class and assigns traffic a packet loss priority of low . The classifier is applied to Ethernet interface xe-0/0/27 .

Table 100: Components of the IEEE 802.1p Priority Remapping Configuration Topology (*continued*)

Component	Settings
Fixed classifier (native FC interface)	Forwarding class— fcoe1 The classifier is applied to native FC interface fc-0/0/2
Rewrite value	IEEE 802.1p code point— 101 The rewrite value is applied to native FC interface fc-0/0/2
PFC configuration (CNP on Ethernet interface)	Name— fcoe1_p5_rewrite_cnp Input CNP code point— 101 Output CNP code point— 101 Output CNP flow control queue— 5 Interface— xe-0/0/27
DCBX application mapping	Application name— myfcoe5 Application ether-type— 0x8906 Application map name— myfcoe5_map Application map code points— 101 Interface— xe-0/0/27 NOTE: LLDP and DCBX must be enabled on the interface. By default, LLDP and DCBX are enabled on all Ethernet interfaces.

The priority used to identify FCoE traffic (5, IEEE 802.1p code point 101) is configured for lossless transport across the QFX device on interfaces **xe-0/0/27** and **fc-0/0/2**, which belong to the same local FC fabric on the FCoE-FC gateway.

On the Ethernet interface, the classifier maps priority 5 to a lossless forwarding class (**fcoe1**), the input CNP enables PFC on incoming priority 5 traffic, and the output CNP enables output queue 5 to respond to pause messages received from the peer on traffic tagged with priority 5. On the native FC interface, FC traffic is remapped from priority 3 (the default mapping) to priority 5 and assigned to the same lossless forwarding class, **fcoe1**, because of the fixed classifier configuration. In this way, traffic tagged with priority 5 on interfaces **xe-0/0/27** and **fc-0/0/2** receives lossless treatment.



NOTE: To avoid fate sharing, ensure that the remapped priority is classified only to the forwarding class used in the fixed classifier on all other interfaces. For example, if you configure a fixed classifier on an FC interface that classifies all of the traffic into lossless forwarding class `fcoe1` and remaps the priority to priority 5 (IEEE 802.1p code point 101), then in all other classifier configurations on all other interfaces, priority 5 should always be classified to forwarding class `fcoe1`. If you classify priority 6 on another interface to forwarding class `fcoe1`, then congestion on priority 6 traffic affects priority 5 traffic unfairly.



NOTE: This example does not include scheduling (bandwidth allocation) configuration or the local FC fabric configuration. This examples focuses only on priority remapping.

Configuration

CLI Quick Configuration

To quickly configure IEEE 802.1p priority remapping on an FCoE-FC gateway, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the `[edit]` hierarchy level.

```
set class-of-service forwarding-classes class fcoe1 queue-num 5 no-loss
set class-of-service classifiers ieee-802.1 fcoe_gw_classifier forwarding-class fcoe1 loss-priority
low code-points 101
set class-of-service interfaces xe-0/0/27 unit 0 classifiers ieee-802.1 fcoe_gw_classifier
set class-of-service interfaces fc-0/0/2 forwarding-class fcoe1
set class-of-service interfaces fc-0/0/2 rewrite-value input ieee-802.1p code-point 101
set class-of-service congestion-notification-profile fcoe1_p5_rewrite_cnp input ieee-802.1
code-point 101 pfc
set class-of-service congestion-notification-profile fcoe1_p5_rewrite_cnp output ieee-802.1
code-point 101 flow-control-queue 5
set class-of-service interfaces xe-0/0/27 congestion-notification-profile fcoe1_p5_rewrite_cnp
set applications application myfcoe5 ether-type 0x8906
set policy-options application-maps myfcoe5_app_map application myfcoe5 code-points 101
set protocols dcbx interface xe-0/0/27 application-map myfcoe5_app_map
```

Step-by-Step Procedure

To configure a lossless forwarding class for FCoE traffic, classify FCoE traffic into that forwarding class, configure a rewrite value on the native FC interface for the FCoE traffic, and enable PFC on the Ethernet interface, and configure DCBX application protocol TLV exchange for FCoE traffic:

1. Configure the lossless forwarding class (named `fcoe1` and mapped to output queue 5) for FCoE traffic that uses IEEE 802.1p priority 5:


```
[edit class-of-service]
user@switch# set forwarding-classes class fcoe1 queue-num 5 no-loss
```
2. Configure an ingress classifier named `fcoe_gw_classifier` to map the FCoE priority (IEEE 802.1p code point 101) to the lossless FCoE forwarding class (`fcoe1`):


```
[edit class-of-service classifiers]
```

- ```

user@switch# set ieee-802.1 fcoe_gw_classifier forwarding-class fcoe1 loss-priority low
code-points 101

```
3. Apply the classifier named **fcoe\_gw\_classifier** to Ethernet interface **xe-0/0/27**:
 

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/27 unit 0 classifiers ieee-802.1 fcoe_gw_classifier

```
  4. Configure the fixed classifier on the native FC interface, using the lossless FCoE forwarding class **fcoe1** (all traffic from the FC SAN is classified into the specified forwarding class). The traffic classified into this forwarding class is tagged with the priority value configured in the next step.
 

```

[edit class-of-service]
user@switch# set interfaces fc-0/0/2 forwarding-class fcoe1

```
  5. Configure the rewrite value (IEEE 802.1p code point **101**) applied to all incoming traffic from the FC SAN on the native FC interface. The rewrite value is the IEEE 802.1p priority that the encapsulated FCoE traffic classified into the **fcoe1** forwarding class uses on the converged Ethernet network.
 

```

[edit class-of-service]
user@switch# set interfaces fc-0/0/2 rewrite-value input ieee-802.1p code-point 101

```
  6. Configure the input stanza of the CNP (named **fcoe1\_p5\_rewrite\_cnp**) to enable PFC on the FCoE priority on the Ethernet interface:
 

```

[edit class-of-service]
user@switch# set congestion-notification-profile fcoe1_p5_rewrite_cnp input ieee-802.1
code-point 101 pfc

```
  7. Configure the output stanza of the CNP to enable output queue 5 to respond to pause messages received from the peer on traffic tagged with priority 5:
 

```

[edit class-of-service]
user@switch# set congestion-notification-profile fcoe1_p5_rewrite_cnp output ieee-802.1
code-point 101 flow-control-queue 5

```
  8. Apply the CNP named **fcoe1\_p5\_rewrite\_cnp** to Ethernet interface **xe-0/0/27**:
 

```

[edit class-of-service]
user@switch# set interfaces xe-0/0/27 congestion-notification-profile fcoe1_p5_rewrite_cnp

```
  9. Configure a DCBX application for FCoE to map to the Ethernet interface, so that DCBX can exchange application protocol TLVs on the correct (remapped) IEEE 802.1p FCoE priority:
 

```

[edit]
user@switch# set applications application myfcoe5 ether-type 0x8906

```
  10. Configure a DCBX application map to map the FCoE application to the correct (remapped) IEEE 802.1p FCoE priority:
 

```

[edit]
user@switch# set policy-options application-maps myfcoe5_app_map application myfcoe5
code-points 101

```
  11. Apply the application map to the Ethernet interface so that DCBX exchanges FCoE application TLVs on the correct code point:
 

```

[edit]
user@switch# set protocols dcbx interface xe-0/0/27 application-map myfcoe5_app_map

```

## Verification

To verify the configuration and proper operation of IEEE 802.1p priority remapping on an FCoE-FC gateway, perform these tasks:

- [Verifying the Forwarding Class Configuration on page 362](#)
- [Verifying the Behavior Aggregate Classifier Configuration on page 362](#)
- [Verifying the FC Interface Configuration \(Fixed Classifier, Rewrite Value\) on page 363](#)
- [Verifying the Ethernet Interface PFC Configuration \(CNP\) on page 363](#)
- [Verifying the Ethernet Interface Configuration on page 364](#)
- [Verifying the DCBX Application Configuration on page 364](#)
- [Verifying the DCBX Application Map Configuration on page 365](#)
- [Verifying the DCBX Application Protocol Exchange Interface Configuration on page 365](#)

### Verifying the Forwarding Class Configuration

**Purpose** Verify that the lossless forwarding class **fcoe1** has been created.

**Action** Show the forwarding class configuration by using the operational command **show class-of-service forwarding-class**:

```
user@switch# show class-of-service forwarding-class
```

| Forwarding class | ID | Queue | Policing priority | No-Loss  |
|------------------|----|-------|-------------------|----------|
| best-effort      | 0  | 0     | normal            | Disabled |
| fcoe             | 1  | 3     | normal            | Enabled  |
| no-loss          | 2  | 4     | normal            | Enabled  |
| network-control  | 3  | 7     | normal            | Disabled |
| fcoe1            | 4  | 5     | normal            | Enabled  |
| mcast            | 8  | 8     | normal            | Disabled |

**Meaning** The **show class-of-service forwarding-class** command shows all of the forwarding classes. The command output shows that the **fcoe1** forwarding class is configured on output queue **5** with the no-loss packet drop attribute enabled.

Because we did not explicitly configure the default forwarding classes, they remain in their default state, including the lossless configuration of the **fcoe** and **no-loss** default forwarding classes.

### Verifying the Behavior Aggregate Classifier Configuration

**Purpose** Verify that the classifier maps the forwarding classes to the correct IEEE 802.1p code points (priorities) and packet loss priorities.

**Action** List the classifier configured for priority remapping using the operational mode command **show class-of-service classifier name fcoe\_gw\_classifier**:

```
user@switch> show class-of-service classifier name fcoe_gw_classifier
Classifier: fcoe_gw_classifier, Code point type: ieee-802.1, Index: 13100
 Code point Forwarding class Loss priority
 101 fcoe1 low
```

**Meaning** The **show class-of-service classifier name fcoe\_gw\_classifier** command shows the IEEE 802.1p code points and the loss priorities that are mapped to the forwarding classes in the classifier. The command output shows that the classifier maps forwarding class **fcoe1** to IEEE 802.1p code point **101** (priority 5) with a packet loss priority of **low**.

### Verifying the FC Interface Configuration (Fixed Classifier, Rewrite Value)

**Purpose** Verify that the native FC interface (NP\_Port) classifies incoming traffic into forwarding class **fcoe1** and that the interface rewrite value is priority 5 (IEEE code point 101).

**Action** Display the FC interface configuration using the operational mode command **show configuration class-of-service interfaces fc-0/0/2**:

```
user@switch> show configuration class-of-service interfaces fc-0/0/2
rewrite-value {
 input {
 ieee-802.1 {
 code-point {
 101;
 }
 }
 }
}
forwarding-class fcoe1;
```

**Meaning** The **show configuration class-of-service interfaces fc-0/0/2** command shows that the rewrite value for incoming (input) traffic is IEEE 802.1p code point **101** (priority 5), and that the interface uses forwarding class **fcoe1** as the fixed classifier for all incoming traffic.

### Verifying the Ethernet Interface PFC Configuration (CNP)

**Purpose** Verify that PFC is enabled on the correct priority (IEEE 802.1p code point **101**) for lossless transport and that flow control is enabled on the correct output queue (queue **5**) on the Ethernet interface.

**Action** List the congestion notification profile using the operational mode command **show class-of-service congestion-notification fcoe1\_p5\_rewrite\_cnp**:

```
user@switch> show class-of-service congestion-notification fcoe1_p5_rewrite_cnp
Name: fcoe1_p5_rewrite_cnp, Index: 7061
Type: Input
Cable Length: 100 m
 Priority PFC MRU
 000 Disabled
 001 Disabled
 010 Disabled
 011 Disabled
```

```
100 Disabled
101 Enabled 2500
110 Disabled
111 Disabled
Type: Output
Priority Flow-Control-Queues
101
5
```

**Meaning** The `show class-of-service congestion-notification fcoe1_p5_rewrite_cnp` command shows the input and output stanzas of the CNP. The input stanza shows that PFC is enabled on IEEE 802.1p code point 101 (priority 5). The input stanza also shows that the CNP uses the default values of 100 meters for the cable length value and 2500 bytes for the maximum receive unit (MRU) value.

The output stanza shows that flow control is enabled on output queue 5 for IEEE 802.1p priority code point 101 (priority 5).

---

### Verifying the Ethernet Interface Configuration

**Purpose** Verify that the classifier `fcoe_gw_classifier` and the congestion notification profile `fcoe1_p5_rewrite_cnp` are configured on Ethernet interface `xe-0/0/27`.

**Action** List the ingress interfaces using the operational mode command `show configuration class-of-service interfaces xe-0/0/27`:

```
user@switch> show configuration class-of-service interfaces xe-0/0/27
congestion-notification-profile fcoe1_p5_rewrite_cnp;
unit 0 {
 classifiers {
 ieee-802.1 fcoe_gw_classifier;
 }
}
```

**Meaning** The `show configuration class-of-service interfaces xe-0/0/27` command shows that the congestion notification profile `fcoe1_p5_rewrite_cnp` is configured on the interface, and that the IEEE 802.1p classifier associated with the interface is `fcoe_gw_classifier`.

---

### Verifying the DCBX Application Configuration

**Purpose** Verify that the DCBX application named `myfcoe5` for FCoE is configured.

**Action** List the DCBX applications by using the configuration mode command `show applications`:

```
user@switch# show applications
application myfcoe5 {
 ether-type 0x8906;
}
```

**Meaning** The `show applications` configuration mode command shows all of the configured applications. The output shows that the application `myfcoe5` is configured with an EtherType of `0x8906` (the correct EtherType for FCoE traffic).



### Verifying the DCBX Application Map Configuration

- Purpose** Verify that the application map **myfcoe5\_app\_map** is configured.
- Action** List the application map by using the configuration mode command **show policy-options application-maps**:
- ```
user@switch# show policy-options application-maps
myfcoe5_app_map {
    application myfcoe5 code-points 101;
}
```
- Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The output shows that there is one application map, **myfcoe5_app_map**, which consists of the application named **myfcoe5** mapped to IEEE 802.1p code point 101 (priority 5).

Verifying the DCBX Application Protocol Exchange Interface Configuration

- Purpose** Verify that the application map is applied to the correct interface (**xe-0/0/27**).
- Action** List the application maps using the configuration mode command **show protocols dcbx**:
- ```
user@switch# show protocols dcbx
interface xe-0/0/27.0 {
 application-map myfcoe5_app_map;
}
```
- Meaning** The **show protocols dcbx** configuration mode command lists the application map association with interfaces. The output shows that interface **xe-0/0/27** uses application map **myfcoe5\_app\_map**.
- Related Documentation**
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)
  - [Example: Configuring Unicast Classifiers on page 273](#)
  - [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
  - [Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces \(NP\\_Ports\) on page 423](#)
  - [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 128](#)

## Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic.

This example shows you the recommended configuration of the global shared buffer pool to support a network that carries mostly best-effort (lossy) unicast traffic. The

global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

Use the default shared buffer settings (for a network with a balanced mix of lossless, best effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.

After starting from the recommended configuration, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

- [Requirements on page 366](#)
- [Overview on page 366](#)
- [Configuration on page 368](#)
- [Verification on page 369](#)

## Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

## Overview

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time

the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)

- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multidestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- Multicast buffers—Percentage of shared buffer pool for all multidestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffers to support a network that carries mostly best-effort unicast traffic, more buffer space needs to be allocated to lossy buffers, and less buffer space should be allocated to lossless buffers. This example shows you how to configure the global shared buffer pool allocation that we recommend to support a network that carries mostly unicast traffic.

### Topology

Table 101 on page 367 shows the configuration components for this example.

**Table 101: Components of the Recommended Shared Buffer Configuration for Best-Effort Unicast Network Topologies**

| Component | Settings       |
|-----------|----------------|
| Hardware  | QFX3500 switch |

Table 101: Components of the Recommended Shared Buffer Configuration for Best-Effort Unicast Network Topologies (*continued*)

| Component             | Settings                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Ingress shared buffer | Percentage of available ingress buffer space allocated to the ingress shared buffer: 100%<br><br>Percentage of ingress buffer space allocated to lossless traffic (lossless buffer partition): 5%<br><br>Percentage of ingress buffer space allocated to lossless headroom traffic (lossless-headroom buffer partition): 0%<br><br>Percentage of ingress buffer space allocated to best-effort traffic (lossy buffer partition): 95% |
| Egress shared buffer  | Percentage of available egress buffer space allocated to the egress shared buffer: 100%<br><br>Percentage of egress buffer space allocated to lossless queues (lossless buffer partition): 5%<br><br>Percentage of egress buffer space allocated to best-effort queues (lossy buffer partition): 75%<br><br>Percentage of egress buffer space allocated to multicast traffic (multicast buffer partition): 20%                       |

## Configuration

### CLI Quick Configuration

To quickly configure the recommended shared buffer settings for networks that carry mostly best-effort unicast traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit class-of-service shared-buffer]
set ingress percent 100
set ingress buffer-partition lossless percent 5
set ingress buffer-partition lossless-headroom percent 0
set ingress buffer-partition lossy percent 95
set egress percent 100
set egress buffer-partition lossless percent 5
set egress buffer-partition lossy percent 75
set egress buffer-partition multicast percent 20
```

### Configuring the Global Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic

### Step-by-Step Procedure

To configure the global ingress and egress shared buffer allocations and partitions for a network that carries mostly best-effort unicast traffic:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent 100
```
2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent 5
user@switch# set ingress buffer-partition lossless-headroom percent 0
user@switch# set ingress buffer-partition lossy percent 95
```

3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set egress percent 100
```
4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:  

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent 5
user@switch# set egress buffer-partition lossy percent 75
user@switch# set egress buffer-partition multicast percent 20
```

## Results

Display the results of the configuration:

```
root@dcbg-tp-pa-02> show configuration class-of-service shared-buffer
ingress {
 percent 100;
 buffer-partition lossless {
 percent 5;
 }
 buffer-partition lossy {
 percent 95;
 }
 buffer-partition lossless-headroom {
 percent 0;
 }
}
egress {
 percent 100;
 buffer-partition lossless {
 percent 5;
 }
 buffer-partition lossy {
 percent 75;
 }
 buffer-partition multicast {
 percent 20;
 }
}
```

## Verification

Verify that the shared buffer configuration has been created properly.

### Verifying the Shared Buffer Configuration

**Purpose** Verify that the ingress and egress global shared buffer pools are correctly configured and partitioned among the shared buffer types.

**Action** List the global shared buffer configuration using the operational mode command **show class-of-service shared-buffer**:

```
user@switch> show class-of-service shared-buffer
root@dcbg-tp-pa-02> show class-of-service shared-buffer
Ingress:
 Total Buffer : 9360.00 KB
 Dedicated Buffer : 2158.00 KB
```

```
Shared Buffer : 7202.00 KB
Lossless : 360.10 KB
Lossless Headroom : 0.00 KB
Lossy : 6841.90 KB
```

**Lossless Headroom Utilization:**

| Node Device | Total   | Used    | Free    |
|-------------|---------|---------|---------|
| 0           | 0.00 KB | 0.00 KB | 0.00 KB |

**Egress:**

```
Total Buffer : 9360.00 KB
Dedicated Buffer : 2704.00 KB
Shared Buffer : 6656.00 KB
Lossless : 332.80 KB
Multicast : 1331.20 KB
Lossy : 4992.00 KB
```

**Meaning** The **show class-of-service shared-buffer** operational command shows all of the ingress and egress global shared buffer settings, including the buffer partitioning.

For the ingress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2158 KB. This is the size of the global ingress dedicated buffer pool when you configure the ingress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, ingress dedicated ingress buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the ingress shared buffer pool configured as 100 percent of the available buffers, the total size of the ingress shared buffer pool is 7202 KB.
- The ingress shared buffer pool is partitioned to allocate:
  - 360.10 KB to lossless traffic
  - No space to lossless headroom traffic
  - 6841.90 KB to lossy unicast traffic
- The Lossless Headroom Utilization field shows how much of the buffer space reserved for paused traffic is used. Because the lossless headroom buffer partition is set to 0 (zero) percent, the total amount of lossless headroom buffer space is 0 KB; therefore the amount of used and free lossless headroom buffer space is also 0 KB.

For the egress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2704 KB. This is the size of the global egress dedicated buffer pool when you configure the egress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, egress dedicated buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.

- With the egress shared buffer pool configured as 100 percent of the available buffers, the total size of the egress shared buffer pool is 6656 KB. This is less than the ingress shared buffer pool because the switch reserves more egress dedicated buffer space than ingress dedicated buffer space. (More dedicated buffer space means less shared buffer space, and more shared buffer space means less dedicated buffer space.)
- The egress shared buffer pool is partitioned to allocate:
  - 332.80 KB to lossless traffic
  - 1331.20 KB to multicast traffic
  - 4992 KB to lossy unicast traffic



**NOTE:** The output values are valid for QFX3500 and QFX3600 switches. QFX5100 switches have larger buffers (12MB instead of 9MB), so the total buffer size and the sizes of each buffer partition are larger on QFX5100 switches.

#### Related Documentation

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 371](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 377](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 383](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 425](#)
- [Understanding CoS Buffer Configuration on page 131](#)

### Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic.

This example shows you the recommended configuration of the global shared buffer pool to support a network that carries mostly best-effort (lossy) traffic on links with Ethernet PAUSE (IEEE 802.3X) enabled. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

Use the default shared buffer settings (for a network with a balanced mix of lossless, best effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.

After starting from the recommended configuration, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

- [Requirements on page 372](#)
- [Overview on page 372](#)
- [Configuration on page 374](#)
- [Verification on page 375](#)

## Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

## Overview

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)
- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multidestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.



The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- Multicast buffers—Percentage of shared buffer pool for all multidestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffers to support a network that carries mostly best-effort traffic on links enabled for Ethernet PAUSE, more buffer space needs to be allocated to ingress dedicated port buffers, and less buffer space should be allocated to ingress shared buffers. Also, more buffer space needs to be allocated to lossless-headroom buffers, and less space to ingress lossy buffers. This example shows you how to configure the global shared buffer pool allocation that we recommend to support a network that carries mostly best-effort traffic on links enabled for Ethernet PAUSE.

### Topology

Table 102 on page 373 shows the configuration components for this example.

**Table 102: Components of the Recommended Shared Buffer Configuration for Best-Effort Network Topologies with Links Enabled for Ethernet PAUSE**

| Component             | Settings                                                                                                                                                                                                                                                                                                                                                                                                                             |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware              | QFX3500 switch                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Ingress shared buffer | Percentage of available ingress buffer space allocated to the ingress shared buffer: 70%<br><br>Percentage of ingress buffer space allocated to lossless traffic (lossless buffer partition): 5%<br><br>Percentage of ingress buffer space allocated to lossless headroom traffic (lossless-headroom buffer partition): 80%<br><br>Percentage of ingress buffer space allocated to best-effort traffic (lossy buffer partition): 15% |

**Table 102: Components of the Recommended Shared Buffer Configuration for Best-Effort Network Topologies with Links Enabled for Ethernet PAUSE** (*continued*)

| Component            | Settings                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Egress shared buffer | <p>Percentage of available egress buffer space allocated to the egress shared buffer: 100%</p> <p>Percentage of egress buffer space allocated to lossless queues (lossless buffer partition): 5%</p> <p>Percentage of egress buffer space allocated to best-effort queues (lossy buffer partition): 75%</p> <p>Percentage of egress buffer space allocated to multicast traffic (multicast buffer partition): 20%</p> |

## Configuration

**CLI Quick Configuration** To quickly configure the recommended shared buffer settings for networks that carry mostly best-effort unicast traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit class-of-service shared-buffer]
set ingress percent 70
set ingress buffer-partition lossless percent 5
set ingress buffer-partition lossless-headroom percent 80
set ingress buffer-partition lossy percent 15
set egress percent 100
set egress buffer-partition lossless percent 5
set egress buffer-partition lossy percent 75
set egress buffer-partition multicast percent 20
```

### Configuring the Global Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links Enabled for Ethernet PAUSE

**Step-by-Step Procedure** To configure the global ingress and egress shared buffer allocations and partitions:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent 70
```
2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent 5
user@switch# set ingress buffer-partition lossless-headroom percent 80
user@switch# set ingress buffer-partition lossy percent 15
```
3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:

```
[edit class-of-service shared-buffer]
user@switch# set egress percent 100
```
4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent 5
user@switch# set egress buffer-partition lossy percent 75
```

```
user@switch# set egress buffer-partition multicast percent 20
```

## Results

Display the results of the configuration:

```
root@dcbg-tp-pa-02> show configuration class-of-service shared-buffer
ingress {
 percent 70;
 buffer-partition lossless {
 percent 5;
 }
 buffer-partition lossy {
 percent 15;
 }
 buffer-partition lossless-headroom {
 percent 80;
 }
}
egress {
 percent 100;
 buffer-partition lossless {
 percent 5;
 }
 buffer-partition lossy {
 percent 75;
 }
 buffer-partition multicast {
 percent 20;
 }
}
```

## Verification

Verify that the shared buffer configuration has been created properly.

### Verifying the Shared Buffer Configuration

**Purpose** Verify that the ingress and egress global shared buffer pools are correctly configured and partitioned among the shared buffer types.

**Action** List the global shared buffer configuration using the operational mode command **show class-of-service shared-buffer**:

```
user@switch> show class-of-service shared-buffer
root@dcbg-tp-pa-02> show class-of-service shared-buffer
Ingress:
Total Buffer : 9360.00 KB
Dedicated Buffer : 4318.60 KB
Shared Buffer : 5041.40 KB
 Lossless : 252.07 KB
 Lossless Headroom : 4033.12 KB
 Lossy : 756.21 KB

Egress:
Total Buffer : 9360.00 KB
Dedicated Buffer : 2704.00 KB
Shared Buffer : 6656.00 KB
 Lossless : 332.80 KB
```

```
Multicast : 1331.20 KB
Lossy : 4992.00 KB
```

**Meaning** The **show class-of-service shared-buffer** operational command shows all of the ingress and egress global shared buffer settings, including the buffer partitioning.

For the ingress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 4318.6 KB. This is the size of the global ingress dedicated buffer pool when you configure the ingress shared buffer pool as 70 percent of the available (user-configurable) buffer space.
- With the ingress shared buffer pool configured as 70 percent of the available buffers, the total size of the ingress shared buffer pool is 5041.4 KB.
- The ingress shared buffer pool is partitioned to allocate:
  - 252.07 KB to lossless traffic
  - 4033.12 KB to lossless headroom traffic
  - 756.21 KB to lossy unicast traffic

For the egress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2704 KB. This is the size of the global egress dedicated buffer pool when you configure the egress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, egress dedicated buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the egress shared buffer pool configured as 100 percent of the available buffers, the total size of the egress shared buffer pool is 6656 KB. This is less than the ingress shared buffer pool because the switch reserves more egress dedicated buffer space than ingress dedicated buffer space. (More dedicated buffer space means less shared buffer space, and more shared buffer space means less dedicated buffer space.)
- The egress shared buffer pool is partitioned to allocate:
  - 332.80 KB to lossless traffic
  - 1331.20 KB to multicast traffic
  - 4992 KB to lossy unicast traffic



**NOTE:** The output values are valid for QFX3500 and QFX3600 switches. QFX5100 switches have larger buffers (12MB instead of 9MB), so the total buffer size and the sizes of each buffer partition are larger on QFX5100 switches.

---

- Related Documentation**
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 365](#)
  - [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 377](#)
  - [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 383](#)
  - [Configuring Global Ingress and Egress Shared Buffers on page 425](#)
  - [Understanding CoS Buffer Configuration on page 131](#)

## Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic.

This example shows you the recommended configuration of the global shared buffer pool to support a network that carries mostly multicast traffic. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on all ports until buffer reprogramming is complete.

Use the default shared buffer settings (for a network with a balanced mix of lossless, best effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.

After starting from the recommended configuration, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

- [Requirements on page 377](#)
- [Overview on page 378](#)
- [Configuration on page 379](#)
- [Verification on page 381](#)

## Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch

- Junos OS Release 12.3 or later for the QFX Series

## Overview

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)
- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multdestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- Multicast buffers—Percentage of shared buffer pool for all multdestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than

100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffers to support a network that carries mostly multicast traffic, more buffer space needs to be allocated to lossy buffers, less buffer space should be allocated to lossless buffers, and more space needs to be allocated to egress multicast buffers. This example shows you how to configure the global shared buffer pool allocation that we recommend to support a network that carries mostly multicast traffic.

### Topology

Table 103 on page 379 shows the configuration components for this example.

**Table 103: Components of the Recommended Shared Buffer Configuration for Multicast Network Topologies**

| Component             | Settings                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware              | QFX3500 switch                                                                                                                                                                                                                                                                                                                                                                                                           |
| Ingress shared buffer | Percentage of available ingress buffer space allocated to the ingress shared buffer: 100%<br>Percentage of ingress buffer space allocated to lossless traffic (lossless buffer partition): 5%<br>Percentage of ingress buffer space allocated to lossless headroom traffic (lossless-headroom buffer partition): 0%<br>Percentage of ingress buffer space allocated to best-effort traffic (lossy buffer partition): 95% |
| Egress shared buffer  | Percentage of available egress buffer space allocated to the egress shared buffer: 100%<br>Percentage of egress buffer space allocated to lossless queues (lossless buffer partition): 5%<br>Percentage of egress buffer space allocated to best-effort queues (lossy buffer partition): 20%<br>Percentage of egress buffer space allocated to multicast traffic (multicast buffer partition): 75%                       |

### Configuration

#### CLI Quick Configuration

To quickly configure the recommended shared buffer settings for networks that carry mostly multicast traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit class-of-service shared-buffer]
set ingress percent 100
set ingress buffer-partition lossless percent 5
set ingress buffer-partition lossless-headroom percent 0
set ingress buffer-partition lossy percent 95
set egress percent 100
set egress buffer-partition lossless percent 5
set egress buffer-partition lossy percent 20
set egress buffer-partition multicast percent 75
```

## Configuring the Global Shared Buffer Pool for Networks with Mostly Multicast Traffic

**Step-by-Step Procedure** To configure the global ingress and egress shared buffer allocations and partitions for a network that carries mostly multicast traffic:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent 100
```
2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent 5
user@switch# set ingress buffer-partition lossless-headroom percent 0
user@switch# set ingress buffer-partition lossy percent 95
```
3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set egress percent 100
```
4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:  

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent 5
user@switch# set egress buffer-partition lossy percent 20
user@switch# set egress buffer-partition multicast percent 75
```

## Results

Display the results of the configuration:

```
root@dcbg-tp-pa-02> show configuration class-of-service shared-buffer
ingress {
 percent 100;
 buffer-partition lossless {
 percent 5;
 }
 buffer-partition lossy {
 percent 95;
 }
 buffer-partition lossless-headroom {
 percent 0;
 }
}
egress {
 percent 100;
 buffer-partition lossless {
 percent 5;
 }
 buffer-partition lossy {
 percent 20;
 }
 buffer-partition multicast {
 percent 75;
 }
}
```



```
 }
}
```

## Verification

Verify that the shared buffer configuration has been created properly.

### Verifying the Shared Buffer Configuration

**Purpose** Verify that the ingress and egress global shared buffer pools are correctly configured and partitioned among the shared buffer types.

**Action** List the global shared buffer configuration using the operational mode command **show class-of-service shared-buffer**:

```
user@switch> show class-of-service shared-buffer
root@dcbg-tp-pa-02> show class-of-service shared-buffer
Ingress:
 Total Buffer : 9360.00 KB
 Dedicated Buffer : 2158.00 KB
 Shared Buffer : 7202.00 KB
 Lossless : 360.10 KB
 Lossless Headroom : 0.00 KB
 Lossy : 6841.90 KB

 Lossless Headroom Utilization:
 Node Device Total Used Free
 0 0.00 KB 0.00 KB 0.00 KB

Egress:
 Total Buffer : 9360.00 KB
 Dedicated Buffer : 2704.00 KB
 Shared Buffer : 6656.00 KB
 Lossless : 332.80 KB
 Multicast : 4992.00 KB
 Lossy : 1331.20 KB
```

**Meaning** The **show class-of-service shared-buffer** operational command shows all of the ingress and egress global shared buffer settings, including the buffer partitioning.

For the ingress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2158 KB. This is the size of the global ingress dedicated buffer pool when you configure the ingress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, ingress dedicated ingress buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the ingress shared buffer pool configured as 100 percent of the available buffers, the total size of the ingress shared buffer pool is 7202 KB.
- The ingress shared buffer pool is partitioned to allocate:
  - 360.10 KB to lossless traffic

- No space to lossless headroom traffic
- 6841.90 KB to lossy unicast traffic
- The Lossless Headroom Utilization field shows how much of the buffer space reserved for paused traffic is used. Because the lossless headroom buffer partition is set to 0 (zero) percent, the total amount of lossless headroom buffer space is 0 KB; therefore the amount of used and free lossless headroom buffer space is also 0 KB.

For the egress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2704 KB. This is the size of the global egress dedicated buffer pool when you configure the egress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, egress dedicated buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the egress shared buffer pool configured as 100 percent of the available buffers, the total size of the egress shared buffer pool is 6656 KB. This is less than the ingress shared buffer pool because the switch reserves more egress dedicated buffer space than ingress dedicated buffer space. (More dedicated buffer space means less shared buffer space, and more shared buffer space means less dedicated buffer space.)
- The egress shared buffer pool is partitioned to allocate:
  - 332.80 KB to lossless traffic
  - 4992 KB to multicast traffic
  - 1331.20 KB to lossy unicast traffic



**NOTE:** The output values are valid for QFX3500 and QFX3600 switches. QFX5100 switches have larger buffers (12MB instead of 9MB), so the total buffer size and the sizes of each buffer partition are larger on QFX5100 switches.

---

**Related  
Documentation**

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 365](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 371](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 383](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 425](#)
- [Understanding CoS Buffer Configuration on page 131](#)

## Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic.

This example shows you the recommended configuration of the global shared buffer pool to support a network that carries mostly lossless traffic. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

Use the default shared buffer settings (for a network with a balanced mix of lossless, best effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.



**NOTE:** When we discuss lossless buffers, we mean buffers that handle traffic on which you enable priority-based flow control (PFC) to ensure lossless transport. The lossless buffers are not used for best-effort traffic on a link on which you enable Ethernet PAUSE (IEEE 802.3x).

After starting from the recommended configuration, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

- [Requirements on page 383](#)
- [Overview on page 384](#)
- [Configuration on page 385](#)
- [Verification on page 387](#)

### Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX3500 Switch
- Junos OS Release 12.3 or later for the QFX Series

## Overview

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)
- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multdestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- Multicast buffers—Percentage of shared buffer pool for all multdestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffers to support a network that carries mostly lossless traffic, more buffer space needs to be allocated to lossless buffers, and less buffer space should be allocated to lossy buffers. This example shows you how to configure the global shared buffer pool allocation that we recommend to support a network that carries mostly lossless traffic.

### Topology

Table 104 on page 385 shows the configuration components for this example.

**Table 104: Components of the Recommended Shared Buffer Configuration for Lossless Network Topologies**

| Component             | Settings                                                                                                                                                                                                                                                                                                                                                                                                                  |
|-----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware              | QFX3500 switch                                                                                                                                                                                                                                                                                                                                                                                                            |
| Ingress shared buffer | Percentage of available ingress buffer space allocated to the ingress shared buffer: 100%<br>Percentage of ingress buffer space allocated to lossless traffic (lossless buffer partition): 15%<br>Percentage of ingress buffer space allocated to lossless headroom traffic (lossless headroom buffer partition): 80%<br>Percentage of ingress buffer space allocated to best-effort traffic (lossy buffer partition): 5% |
| Egress shared buffer  | Percentage of available egress buffer space allocated to the egress shared buffer: 100%<br>Percentage of egress buffer space allocated to lossless queues (lossless buffer partition): 90%<br>Percentage of egress buffer space allocated to best-effort queues (lossy buffer partition): 5%<br>Percentage of egress buffer space allocated to multicast traffic (multicast buffer partition): 5%                         |

### Configuration

#### CLI Quick Configuration

To quickly configure the recommended shared buffer settings for networks that carry mostly lossless traffic, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level:

```
[edit class-of-service shared-buffer]
set ingress percent 100
set ingress buffer-partition lossless percent 15
set ingress buffer-partition lossless-headroom percent 80
set ingress buffer-partition lossy percent 5
set egress percent 100
set egress buffer-partition lossless percent 90
set egress buffer-partition lossy percent 5
set egress buffer-partition multicast percent 5
```

## Configuring the Global Shared Buffer Pool for Networks with Mostly Lossless Traffic

**Step-by-Step Procedure** To configure the global ingress and egress shared buffer allocations and partitions for a network that carries mostly lossless traffic:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent 100
```
2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent 15
user@switch# set ingress buffer-partition lossless-headroom percent 80
user@switch# set ingress buffer-partition lossy percent 5
```
3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set egress percent 100
```
4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:  

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent 90
user@switch# set egress buffer-partition lossy percent 5
user@switch# set egress buffer-partition multicast percent 5
```

## Results

Display the results of the configuration:

```
rroot@dcbg-tp-pa-02> show configuration class-of-service shared-buffer
ingress {
 percent 100;
 buffer-partition lossless {
 percent 15;
 }
 buffer-partition lossy {
 percent 5;
 }
 buffer-partition lossless-headroom {
 percent 80;
 }
}
egress {
 percent 100;
 buffer-partition lossless {
 percent 90;
 }
 buffer-partition lossy {
 percent 5;
 }
 buffer-partition multicast {
 percent 5;
 }
}
```

```
 }
}
```

## Verification

Verify that the shared buffer configuration has been created properly.

### Verifying the Shared Buffer Configuration

**Purpose** Verify that the ingress and egress global shared buffer pools are correctly configured and partitioned among the shared buffer types.

**Action** List the global shared buffer configuration using the operational mode command **show class-of-service shared-buffer**:

```
user@switch> show class-of-service shared-buffer
root@dcbg-tp-pa-02> show class-of-service shared-buffer
Ingress:
 Total Buffer : 9360.00 KB
 Dedicated Buffer : 2158.00 KB
 Shared Buffer : 7202.00 KB
 Lossless : 1080.30 KB
 Lossless Headroom : 5761.60 KB
 Lossy : 360.10 KB

 Lossless Headroom Utilization:
 Node Device Total Used Free
 0 5761.60 KB 0.00 KB 5761.60 KB

Egress:
 Total Buffer : 9360.00 KB
 Dedicated Buffer : 2704.00 KB
 Shared Buffer : 6656.00 KB
 Lossless : 5990.40 KB
 Multicast : 332.80 KB
 Lossy : 332.80 KB
```

**Meaning** The **show class-of-service shared-buffer** operational command shows all of the ingress and egress global shared buffer settings, including the buffer partitioning.

For the ingress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2158 KB. This is the size of the global ingress dedicated buffer pool when you configure the ingress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, ingress dedicated ingress buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the ingress shared buffer pool configured as 100 percent of the available buffers, the total size of the ingress shared buffer pool is 7202 KB.
- The ingress shared buffer pool is partitioned to allocate:
  - 1080 KB to lossless traffic

- 5761.60 KB to lossless headroom traffic
- 360.10 KB to lossy unicast traffic
- The Lossless Headroom Utilization field shows how much of the buffer space reserved for paused traffic is used. Of the total available lossless headroom buffer space of 5761.60 KB, currently no buffer space is being used, so all 5761.60 KB of buffer space is free.

For the egress shared buffers, the command output shows:

- The total switch buffer pool is 9360 KB (9 MB).
- The dedicated buffer pool is 2704 KB. This is the size of the global egress dedicated buffer pool when you configure the egress shared buffer pool as 100 percent of the available (user-configurable) buffer space. This is the minimum size of the reserved, egress dedicated buffer pool (not user-configurable). If you configure the shared buffer as less than 100 percent of the available buffer pool, the remaining buffer space is added to the dedicated buffer pool.
- With the egress shared buffer pool configured as 100 percent of the available buffers, the total size of the egress shared buffer pool is 6656 KB. This is less than the ingress shared buffer pool because the switch reserves more egress dedicated buffer space than ingress dedicated buffer space. (More dedicated buffer space means less shared buffer space, and more shared buffer space means less dedicated buffer space.)
- The egress shared buffer pool is partitioned to allocate:
  - 5990.40 KB to lossless traffic
  - 332.80 KB to multicast traffic
  - 332.80 KB to lossy unicast traffic



**NOTE:** The output values are valid for QFX3500 and QFX3600 switches. QFX5100 switches have larger buffers (12MB instead of 9MB), so the total buffer size and the sizes of each buffer partition are larger on QFX5100 switches.

---

**Related  
Documentation**

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 365](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 371](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 377](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 425](#)
- [Understanding CoS Buffer Configuration on page 131](#)



## Example: Configuring DCBX Application Protocol TLV Exchange

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of connected peers by exchanging application configuration information. DCBX detects feature misconfiguration and mismatches and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP). LLDP must remain enabled on every interface on which you want to use DCBX.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX application protocol exchange for Layer 2 and Layer 4 applications such as the Internet Small Computer System Interface (iSCSI). You specify applications by EtherType (for Layer 2 applications) or by the destination port and protocol (for Layer 4 applications; the protocol can be either TCP or UDP).

The QFX Series handles Fibre Channel over Ethernet (FCoE) application protocol exchange differently than other protocols in some cases:

- If FCoE is the only application for which you want to enable DCBX application protocol TLV exchange on an interface, you do not have to explicitly configure the FCoE application or an application map. By default, the QFX Series exchanges FCoE application protocol TLVs on all interfaces that carry FCoE traffic (traffic mapped to the **fcoe** forwarding class) and have priority-based flow control (PFC) enabled on the FCoE priority (the FCoE IEEE 802.1p code point). The default priority mapping for the FCoE application is IEEE 802.1p code point 011 (the default **fcoe** forwarding class code point).
- If you want an interface to use DCBX to exchange application protocol TLVs for any other applications in addition to FCoE, you must configure the applications (including FCoE), define an application map (including FCoE), and apply the application map to the interface. If you apply an application map to an interface, you must explicitly configure the FCoE application, or the interface does not exchange FCoE application protocol TLVs.

This example shows how to configure interfaces to exchange both Layer 2 and Layer 4 applications by configuring one interface to exchange iSCSI and FCoE application protocol information and configuring another interface to exchange iSCSI and Precision Time Protocol (PTP) application protocol information.

- [Requirements on page 390](#)
- [Overview on page 390](#)
- [Configuration on page 393](#)
- [Verification on page 395](#)

## Requirements

This example uses the following hardware and software components:

- Juniper Networks QFX Series device
- Junos OS Release 12.1 or later for the QFX Series

## Overview

The switch supports DCBX application protocol exchange for:

- Layer 2 applications, defined by EtherType
- Layer 4 applications, defined by destination port and protocol



**NOTE:** DCBX also advertises PFC and enhanced transmission selection (ETS) information. See [“Configuring DCBX Autonegotiation” on page 431](#) for how DCBX negotiates and advertises configuration information for these features and for the applications.

DCBX is configured on a per-interface basis for each supported feature or application. For applications that you want to enable for DCBX application protocol exchange, you must:

- Define the application name and configure the EtherType or the destination port and protocol (TCP or UDP) of the application. Use the EtherType for Layer 2 applications, and use the destination port and protocol for Layer 4 protocols.
- Map the application to an IEEE 802.1p code point in an application map.
- Add the application map to DCBX interface.

In addition, for all applications (including FCoE, even when you do not use an application map), you either must create an IEEE 802.1p classifier and apply it to the appropriate ingress interfaces or use the default classifier. A classifier maps the code points of incoming traffic to a forwarding class and a loss priority so that ingress traffic is assigned to the correct class of service (CoS). The forwarding class determines the output queue on the egress interface.

If you do not create classifiers, trunk and tagged-access ports use the unicast IEEE 802.1 default trusted classifier. [Table 105 on page 391](#) shows the default mapping of IEEE 802.1 code-point values to unicast forwarding classes and loss priorities for ports in trunk mode or tagged-access mode. [Table 106 on page 391](#) shows the default untrusted classifier IEEE 802.1 code-point values to unicast forwarding class mapping for ports in access mode.

**Table 105: Default IEEE 802.1 Classifiers for Trunk Ports and Tagged-Access Ports (Default Trusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| be (000)   | best-effort      | low           |
| be1 (001)  | best-effort      | low           |
| ef (010)   | best-effort      | low           |
| ef1 (011)  | fcoe             | low           |
| af11 (100) | no-loss          | low           |
| af12 (101) | best-effort      | low           |
| nc1 (110)  | network-control  | low           |
| nc2 (111)  | network-control  | low           |

**Table 106: Default IEEE 802.1 Unicast Classifiers for Access Ports (Default Untrusted Classifier)**

| Code Point | Forwarding Class | Loss Priority |
|------------|------------------|---------------|
| 000        | best-effort      | low           |
| 001        | best-effort      | low           |
| 010        | best-effort      | low           |
| 011        | best-effort      | low           |
| 100        | best-effort      | low           |
| 101        | best-effort      | low           |
| 110        | best-effort      | low           |
| 111        | best-effort      | low           |

### Topology

This example shows how to configure DCBX application protocol exchange for three protocols (iSCSI, PTP, and FCoE) on two interfaces. One interface exchanges iSCSI and FCoE application protocol information, and the other interface exchanges iSCSI and PTP application protocol information.



**NOTE:** You must map FCoE traffic to the interfaces on which you want to forward FCoE traffic. You must also enable PFC on the FCoE interfaces and create an ingress classifier for FCoE traffic, or else use the default classifier.

Table 107 on page 392 shows the configuration components for this example.

**Table 107: Components of DCBX Application Protocol Exchange Configuration Topology**

| Component                   | Settings                                                                                                                                                                                                                                                                                                                                                     |
|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Hardware                    | QFX Series device                                                                                                                                                                                                                                                                                                                                            |
| LLDP                        | Enabled by default on Ethernet interfaces                                                                                                                                                                                                                                                                                                                    |
| DCBX                        | Enabled by default on Ethernet interfaces                                                                                                                                                                                                                                                                                                                    |
| iSCSI application (Layer 4) | Application name— <b>iscsi</b><br>protocol— <b>TCP</b><br>destination-port— <b>3260</b><br>code-points— <b>111</b>                                                                                                                                                                                                                                           |
| PTP application (Layer 2)   | Application name— <b>ptp</b><br>ether-type— <b>0x88F7</b><br>code-points— <b>001, 101</b>                                                                                                                                                                                                                                                                    |
| FCoE application (Layer 2)  | Application name— <b>fcoe</b><br>ether-type— <b>0x8906</b><br>code-points— <b>011</b><br><br><b>NOTE:</b> You explicitly configure the FCoE application because you are applying an application map to the interface. When you apply an application map to an interface, all applications must be explicitly configured and included in the application map. |
| Application maps            | <b>dcbx-iscsi-fcoe-app-map</b> —Maps the iSCSI and FCoE applications to IEEE 802.1p code points<br><br><b>dcbx-iscsi-ptp-app-map</b> —Maps iSCSI and PTP applications to IEEE 802.1p code points                                                                                                                                                             |
| Interfaces                  | <b>xe-0/0/10</b> —Configured to exchange FCoE and iSCSI application TLVs (uses application map <b>dcbx-iscsi-fcoe-app-map</b> , carries FCoE traffic, and has PFC enabled on the FCoE priority)<br><br><b>xe-0/0/11</b> —Configured to exchange iSCSI and PTP application TLVs (uses application map <b>dcbx-iscsi-ptp-app-map</b> )                         |

**Table 107: Components of DCBX Application Protocol Exchange Configuration Topology (*continued*)**

| Component                                                                                                         | Settings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PFC congestion notification profile for FCoE application exchange                                                 | <b>fcoe-cnp:</b> <ul style="list-style-type: none"> <li>• Code point—<b>011</b></li> <li>• Interface—<b>xe-0/0/10</b></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Behavior aggregate classifiers (map forwarding classes to incoming packets by the packet's IEEE 802.1 code point) | <b>fcoe-iscsi-cl1:</b> <ul style="list-style-type: none"> <li>• Maps the <b>fcoe</b> forwarding class to the IEEE 802.1p code point used for the FCoE application (<b>011</b>) and a loss priority of <b>high</b></li> <li>• Maps the <b>network-control</b> forwarding class to the IEEE 802.1p code point used for the iSCSI application (<b>111</b>) and a loss priority of <b>high</b></li> <li>• Applied to interface <b>xe-0/0/10</b></li> </ul> <b>iscsi-ntp-cl2:</b> <ul style="list-style-type: none"> <li>• Maps the <b>network-control</b> forwarding class to the IEEE 802.1p code point used for the iSCSI application (<b>111</b>) and a loss priority of <b>low</b></li> <li>• Maps the <b>best-effort</b> forwarding class to the IEEE 802.1p code points used for the PTP application (<b>001</b> and <b>101</b>) and a loss priority of <b>low</b></li> <li>• Applied to interface <b>xe-0/0/11</b></li> </ul> |



**NOTE:** This example does not include scheduling (bandwidth allocation) configuration or lossless configuration for the iSCSI forwarding class.

## Configuration

### CLI Quick Configuration

To quickly configure DCBX application protocol exchange, copy the following commands, paste them in a text file, remove line breaks, change variables and details to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set applications application iSCSI protocol tcp destination-port 3260
set applications application FCoE ether-type 0x8906
set applications application PTP ether-type 0x88F7
set policy-options application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011
set policy-options application-maps dcbx-iscsi-ntp-app-map application iSCSI code-points 111
set policy-options application-maps dcbx-iscsi-ntp-app-map application PTP code-points [001 101]
set protocols dcbx interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
set protocols dcbx interface xe-0/0/11 application-map dcbx-iscsi-ntp-app-map
set class-of-service congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011 pfc
set class-of-service interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class fcoe
loss-priority high code-points 011
set class-of-service classifiers ieee-802.1 fcoe-iscsi-cl1 import default forwarding-class
network-control loss-priority high code-points 111
```

```

set class-of-service classifiers ieee-802.1p iscsi-ntp-cl2 import default forwarding-class
network-control loss-priority low code-points 111
set class-of-service classifiers ieee-802.1p iscsi-ntp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001101]
set class-of-service interfaces xe-0/0/10 unit 0 classifiers ieee-802.1p fcoe-iscsi-cl1
set class-of-service interfaces xe-0/0/11 unit 0 classifiers ieee-802.1p iscsi-ntp-cl2

```

### Configuring DCBX Application Protocol TLV Exchange

#### Step-by-Step Procedure

To define the applications, map the applications to IEEE 802.1p code points, apply the applications to interfaces, and create classifiers for DCBX application protocol exchange:

1. Define the iSCSI application by specifying its protocol and destination port, and define the FCoE and PTP applications by specifying their EtherTypes.  

```

[edit applications]
user@switch# set application iSCSI protocol tcp destination-port 3260
user@switch# set application FCoE ether-type 0x8906
user@switch# set application PTP ether-type 0x88f7

```
2. Define an application map that maps the iSCSI and FCoE applications to IEEE 802.1p code points.  

```

[edit policy-options]
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-fcoe-app-map application FCoE code-points 011

```
3. Define the application map that maps the iSCSI and PTP applications to IEEE 802.1p code points.  

```

[edit policy-options]
user@switch# set application-maps dcbx-iscsi-ntp-app-map application iSCSI code-points 111
user@switch# set application-maps dcbx-iscsi-ntp-app-map application PTP code-points [001101]

```
4. Apply the iSCSI and FCoE application map to interface **xe-0/0/10**, and apply the iSCSI and PTP application map to interface **xe-0/0/11**.  

```

[edit protocols dcbx]
user@switch# set interface xe-0/0/10 application-map dcbx-iscsi-fcoe-app-map
user@switch# set interface xe-0/0/11 application-map dcbx-iscsi-ntp-app-map

```
5. Create the congestion notification profile to enable PFC on the FCoE code point (**011**), and apply the congestion notification profile to interface **xe-0/0/10**.  

```

[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1p code-point 011 pfc
user@switch# set interfaces xe-0/0/10 congestion-notification-profile fcoe-cnp

```
6. Configure the classifier to apply to the interface that exchanges iSCSI and FCoE application information.  

```

[edit class-of-service classifiers]
user@switch# set ieee-802.1p fcoe-iscsi-cl1 import default forwarding-class fcoe loss-priority high code-points 011
user@switch# set ieee-802.1p fcoe-iscsi-cl1 import default forwarding-class network-control loss-priority high code-points 111

```

7. Configure the classifier to apply to the interface that exchanges iSCSI and PTP application information.
 

```
[edit class-of-service classifiers]
user@switch# set ieee-802.1 iscsi-ptp-cl2 import default forwarding-class network-control
loss-priority low code-points 111
user@switch# set ieee-802.1 iscsi-ptp-cl2 import default forwarding-class best-effort
loss-priority low code-points [001 101]
```
8. Apply the classifiers to the appropriate interfaces.
 

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/10 unit 0 classifiers ieee-802.1 fcoe-iscsi-cl1
user@switch# set interfaces xe-0/0/11 unit 0 classifiers ieee-802.1 iscsi-ptp-cl2
```

## Verification

To verify that DCBX application protocol exchange configuration has been created and is operating properly, perform these tasks:

- [Verifying the Application Configuration on page 395](#)
- [Verifying the Application Map Configuration on page 396](#)
- [Verifying DCBX Application Protocol Exchange Interface Configuration on page 396](#)
- [Verifying the PFC Configuration on page 397](#)
- [Verifying the Classifier Configuration on page 397](#)

### Verifying the Application Configuration

**Purpose** Verify that DCBX applications have been configured.

**Action** List the applications by using the configuration mode command **show applications**:

```
user@switch# show applications
application iSCSI {
 protocol tcp;
 destination-port 3260;
}

application fcoe {
 ether-type 0x8906;
}

application ptp {
 ether-type 0x88F7;
}
```

**Meaning** The **show applications** configuration mode command lists all of the configured applications and either their protocol and destination port (Layer 4 applications) or their EtherType (Layer 2 applications). The command output shows that the iSCSI application is configured with the **tcp** protocol and destination port **3260**, the FCoE application is configured with the EtherType **0x8906**, and that the PTP application is configured with the EtherType **0x88F7**.

### Verifying the Application Map Configuration

---

- Purpose** Verify that the application maps have been configured.
- Action** List the application maps by using the configuration mode command **show policy-options application-maps**:
- ```
user@switch# show policy-options application-maps
dcbx-iscsi-fcoe-app-map {
    application iSCSI code-points 111;
    application FCoE code-points 011;
}

dcbx-iscsi-ptp-app-map {
    application iSCSI code-points 111;
    application PTP code-points [001 101];
}
```
- Meaning** The **show policy-options application-maps** configuration mode command lists all of the configured application maps and the applications that belong to each application map. The command output shows that there are two application maps, **dcbx-iscsi-fcoe-app-map** and **dcbx-iscsi-ptp-app-map**.
- The application map **dcbx-iscsi-fcoe-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point 111, and the FCoE application, which is mapped to IEEE 802.1p code point 011.
- The application map **dcbx-iscsi-ptp-app-map** consists of the iSCSI application, which is mapped to IEEE 802.1p code point 111, and the PTP application, which is mapped to IEEE 802.1p code points 001 and 101.

Verifying DCBX Application Protocol Exchange Interface Configuration

- Purpose** Verify that the application maps have been applied to the correct interfaces.
- Action** List the application maps by using the configuration mode command **show protocols dcbx**:
- ```
user@switch# show protocols dcbx
interface xe-0/0/10.0 {
 application-map dcbx-iscsi-fcoe-app-map;
}

interface xe-0/0/11.0 {
 application-map dcbx-iscsi-ptp-app-map;
}
```
- Meaning** The **show protocols dcbx** configuration mode command lists whether the interfaces are enabled for DCBX and lists the application map applied to each interface. The command output shows that interfaces **xe-0/0/10.0** and **xe-0/0/11.0** are enabled for DCBX, and that interface **xe-0/0/10.0** uses application map **dcbx-iscsi-fcoe-app-map**, and interface **xe-0/0/11.0** uses application map **dcbx-iscsi-ptp-app-map**.



### Verifying the PFC Configuration

**Purpose** Verify that PFC has been enabled on the FCoE code point and applied to the correct interface.

**Action** Display the PFC configuration to verify that PFC is enabled on the FCoE code point (011) in the congestion notification profile **fcoe-cnp** by using the configuration mode command **show class-of-service congestion-notification-profile**:

```
user@switch# show class-of-service congestion-notification-profile
fcoe-cnp {
 input {
 ieee-802.1 {
 code-point 011 {
 pfc;
 }
 }
 }
}
```

Display the class-of-service (CoS) interface information to verify that the correct interface has PFC enabled for the FCoE application by using the configuration mode command **show class-of-service interfaces**:

```
user@switch# show class-of-service interfaces
xe-0/0/10 {
 congestion-notification-profile fcoe-cnp;
}
```



**NOTE:** The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the PFC configuration.

**Meaning** The **show class-of-service congestion-notification-profile** configuration mode command lists the configured congestion notification profiles. The command output shows that the congestion notification profile **fcoe-cnp** has been configured and has enabled PFC on the IEEE 802.1p code point **011** (the default FCoE code point).

The **show class-of-service interfaces** configuration mode command shows the interface CoS configuration. The command output shows that the congestion notification profile **fcoe-cnp**, which enables PFC on the FCoE code point, is applied to interface **xe-0/0/10**.

### Verifying the Classifier Configuration

**Purpose** Verify that the classifiers have been configured and applied to the correct interfaces.

**Action** Display the classifier configuration by using the configuration mode command **show class-of-service**:

```
user@switch# show class-of-service
```

```

classifiers {
 ieee-802.1 fcoe-iscsi-cl1 {
 import default;
 forwarding-class network-control {
 loss-priority high code-points 111;
 }
 forwarding-class fcoe {
 loss-priority high code-points 011;
 }
 }
 ieee-802.1 iscsi-ntp-cl2 {
 import default;
 forwarding-class network-control {
 loss-priority low code-points 111;
 }
 forwarding-class best-effort {
 loss-priority low code-points [001 101];
 }
 }
}
interfaces {
 xe-0/0/10 {
 congestion-notification-profile fcoe-cnp;
 unit 0 {
 classifiers {
 ieee-802.1 fcoe-iscsi-cl1;
 }
 }
 }
 xe-0/0/11 {
 unit 0 {
 classifiers {
 ieee-802.1 iscsi-ntp-cl2;
 }
 }
 }
}

```



**NOTE:** The sample output does not include all of the information this command can show. The output is abbreviated to focus on verifying the classifier configuration.

**Meaning** The **show class-of-service** configuration mode command lists the classifier and CoS interface configuration, as well as other information not shown in this example. The command output shows that there are two classifiers configured, **fcoe-iscsi-cl1** and **iscsi-ntp-cl2**.

Classifier **fcoe-iscsi-cl1** uses the **default** classifier as a template and edits the template as follows:

- The forwarding class **network-control** is set to a loss priority of **high** and is mapped to code point 111 (the code point mapped to the iSCSI application).

- The forwarding class **fcoe** is set to a loss priority of **high** and is mapped to code point **011** (the code point mapped by default to the FCoE application).

Classifier **iscsi-ptp-cl2** uses the **default** classifier as a template and edits the template as follows:

- The forwarding class **network-control** is set to a loss priority of **low** and is mapped to IEEE 802.1p code point **111** (the code point mapped to the iSCSI application).
- The forwarding class **best-effort** is set to a loss priority of **low** and is mapped to IEEE 802.1p code points **001** and **101** (the code points mapped by default to the PTP application).

The command output also shows that classifier **fcoe-iscsi-cl1** is mapped to interface **xe-0/0/10.0** and that classifier **iscsi-ptp-cl2** is mapped to interface **xe-0/0/11.0**.

#### Related Documentation

- [Example: Configuring Unicast Classifiers on page 273](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 434](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 436](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 437](#)
- [Configuring DCBX Autonegotiation on page 431](#)
- [show dcbx on page 624](#)
- [show dcbx neighbors on page 625](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 179](#)
- *Using DCBX Protocol to Lower Costs*



## CHAPTER 6

# Configuration Tasks

- [Configuring CoS on page 402](#)
- [Defining CoS Code-Point Aliases on page 404](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 404](#)
- [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 406](#)
- [Configuring CoS WRED Drop Profiles on page 407](#)
- [Configuring CoS Drop Profile Maps on page 408](#)
- [Defining CoS Forwarding Classes on page 408](#)
- [Defining CoS Forwarding Class Sets on page 410](#)
- [Defining CoS Queue Schedulers on page 410](#)
- [Defining CoS Queue Scheduling Priority on page 414](#)
- [Changing the Host Outbound Traffic Default Queue Mapping on page 415](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 416](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 420](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 421](#)
- [Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces \(NP\\_Ports\) on page 423](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 425](#)
- [Defining CoS Rewrite Rules on page 427](#)
- [Assigning CoS Components to Interfaces on page 429](#)
- [Configuring the DCBX Mode on page 430](#)
- [Configuring DCBX Autonegotiation on page 431](#)
- [Disabling the ETS Recommendation TLV on page 434](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 434](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 436](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 437](#)

## Configuring CoS

---

The class-of-service topics describe how to configure the Junos CoS components for the QFX Series. You can configure forwarding classes for transmitting packets, define which packets are placed into each output queue, and schedule the transmission service level for each queue. After defining the CoS components, you assign classifiers to the required physical and logical interfaces.

You can configure various CoS components individually or in combination to define CoS services.



**NOTE:** When you change or when you deactivate and then reactivate the class-of-service configuration, the system experiences packet drops because the system momentarily blocks traffic to change the mapping of incoming traffic to input queues.

The following topics describe how to configure CoS components :

- [Example: Configuring Unicast Classifiers on page 273](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 277](#)
- [Example: Configuring WRED Drop Profiles on page 279](#)
- [Example: Configuring Drop Profile Maps on page 281](#)
- [Example: Configuring Forwarding Classes on page 283](#)
- [Example: Configuring Forwarding Class Sets on page 286](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring Queue Scheduling Priority on page 295](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 300](#)
- [Example: Configuring Maximum Output Bandwidth on page 305](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 241](#)
- [Example: Configuring CoS for FCoE Transit Switch Traffic Across an MC-LAG on page 249](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 356](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326](#)

- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 309](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 317](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 340](#)
- [Defining CoS Code-Point Aliases on page 404](#)
- [Defining CoS Rewrite Rules on page 427](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
- [Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces \(NP\\_Ports\) on page 423](#)
- [Assigning CoS Components to Interfaces on page 429](#)
- [Changing the Host Outbound Traffic Default Queue Mapping on page 415](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 416](#)
- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 420](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 421](#)
- [Configuring the DCBX Mode on page 430](#)
- [Configuring DCBX Autonegotiation on page 431](#)
- [Disabling the ETS Recommendation TLV on page 434](#)
- [Defining an Application for DCBX Application Protocol TLV Exchange on page 434](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 436](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 437](#)

## Defining CoS Code-Point Aliases

---

You can use code-point aliases to streamline the process of configuring CoS features on your switch. A code-point alias assigns a name to a pattern of code-point bits. You can use this name instead of the bit pattern when you configure other CoS components such as classifiers, drop-profile maps, and rewrite rules.

You can configure code-point aliases for the following CoS marker types:

- DSCP or DSCP IPv6—Handles incoming IPv4 or IPv6 packets.
- IEEE 802.1p—Handles Layer 2 CoS.

To configure a code-point alias:

1. Specify a CoS marker type (IEEE 802.1 or DSCP).
2. Assign an alias.
3. Specify the code point that corresponds to the alias.

```
[edit class-of-service code-point-aliases]
user@switch# set (dscp | dscp-ipv6 | ieee-802.1) alias-name code-point-bits
```

For example, to configure a code-point alias for an IEEE 802.1 CoS marker type that has the alias name `fcoe1` and maps to the code-point bits 011:

```
[edit class-of-service code-point-aliases]
user@switch# set ieee-802.1 fcoe1 011
```

### Related Documentation

- [Monitoring CoS Value Aliases on page 550](#)
- [Understanding CoS Code-Point Aliases on page 49](#)

## Defining CoS Unicast BA Classifiers (DSCP, DSCP IPv6, IEEE 802.1p)

---

Packet classification associates incoming packets with a particular CoS servicing level. Behavior aggregate (BA) classifiers examine the Differentiated Services code point (DSCP or DSCP IPv6) value, the IEEE 802.1p CoS value, or the MPLS EXP value in the packet header to determine the CoS settings applied to the packet. (See *Configuring a Global MPLS EXP Classifier* for how to define EXP classifiers for MPLS traffic.) BA classifiers allow you to set the forwarding class and loss priority of a packet based on the incoming CoS value.

Unicast traffic must use different classifiers than multidestination (multicast, broadcast, and destination lookup fail) traffic.

To configure a unicast DSCP, DSCP IPv6, or IEEE 802.1p BA classifier using the CLI:

1. Create a unicast BA classifier:



- To create a unicast DSCP, DSCP IPv6, or IEEE 802.1p BA classifier based on the default classifier, import the default DSCP, DSCP IPv6, or IEEE 802.1p classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1) classifier-name import default forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

- To create a unicast BA classifier that is not based on the default classifier, create a DSCP, DSCP IPv6, or IEEE 802.1p classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1) classifier-name forwarding-class
forwarding-class-name loss-priority level code-points [aliases] [bit-patterns]
```

2. Apply the unicast classifier to a specific 10-Gigabit Ethernet interface or to all 10-Gigabit Ethernet interfaces or to all Fibre Channel interfaces on the switch.

- To apply the classifier to a specific interface:

```
[edit class-of-service interfaces]
user@switch# set interface-name unit unit classifiers (dscp | ieee-802.1) classifier-name
```

- To apply the classifier to all 10-Gigabit Ethernet interfaces on the switch, use wildcards for the interface name and the logical-interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set xe-* unit * classifiers (dscp | ieee-802.1) classifier-name
```

#### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Unicast Classifiers on page 273](#)
- [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 406](#)
- [Configuring a Global MPLS EXP Classifier](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers](#)
- [Monitoring CoS Classifiers on page 545](#)
- [Understanding CoS Classifiers on page 51](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57](#)

## Defining CoS Multidestination (Multicast, Broadcast, DLF) BA Classifiers

Packet classification associates incoming packets with a particular CoS servicing level. Behavior aggregate (BA) classifiers examine the Differentiated Services code point (DSCP) value or IEEE 802.1p CoS value in the packet header to determine the CoS settings applied to the packet. BA classifiers allow you to set the forwarding class and loss priority of a packet based on the incoming CoS value.



**NOTE:** DSCP IPv6 multidestination classifiers are not supported. IPv6 multidestination traffic uses the DSCP multidestination classifier.

Multidestination classifiers apply to all of the switch interfaces and handle multicast, broadcast, and destination lookup fail (DLF) traffic. You cannot apply a multidestination classifier to a single interface or to a range of interfaces.

Unicast and multidestination traffic must use different classifiers.

To configure a multidestination BA classifier using the CLI:

1. Create a DSCP or IEEE 802.1p classifier and associate it with a forwarding class, a loss priority, and a code point:

```
[edit class-of-service classifiers]
user@switch# set (dscp | ieee-802.1) classifier-name forwarding-class forwarding-class-name
loss-priority level code-points [aliases] [bit-patterns]
```

2. Configure the classifier as a multidestination classifier:

```
[edit class-of-service]
user@switch# set multi-destination classifiers (dscp | ieee-802.1) classifier-name
```

### Related Documentation

- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 277](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 404](#)
- [Monitoring CoS Classifiers on page 545](#)
- [Understanding CoS Classifiers on page 51](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57](#)

## Configuring CoS WRED Drop Profiles

You can configure an interpolated weighted random early detection (WRED) profile to control packet drop characteristics for different traffic loss priorities.



**NOTE:** You cannot enable WRED on multidestination (multicast) queues. You can enable WRED only on unicast queues.

Also, do not enable WRED on lossless traffic flows. Use priority-based flow control (PFC) to prevent packet loss on lossless forwarding classes.

*Interpolated* means that the switch creates a smooth drop curve from a drop start point to a drop end point, with a maximum drop rate that is reached at the drop end point.

The drop start point is the average queue fill level when the WRED algorithm starts to drop packets. Before the drop start point, no packets are scheduled to drop. Specify the drop start point using the first of two **fill-level** statements.

The drop end point is the average queue fill level at which all subsequently arriving packets are dropped. When the queue fill levels falls below the drop end point, packets begin to be forwarded again. (At the drop end point, the packet drop probability becomes 100 percent.) Specify the drop end point using the second of two **fill-level** statements.

The minimum drop rate is always 0. Specify the minimum drop rate using the first of two **drop-probability** statements. The maximum drop rate is the drop probability when the average queue fill level reaches the drop end point. Specify the maximum drop rate using the second of two **drop-probability** statements.

The drop rate is zero until the queue fill level reaches the drop start point. As the queue continues to fill, packets drop in smooth linear curve until the queue reaches the drop end point, when packets drop at the maximum drop rate. If the queue fills beyond the drop end point, all packets that match the drop profile are dropped.

To configure a WRED profile using the CLI:

- Name the drop profile and set the drop start point, drop end point, minimum drop rate, and maximum drop rate for the drop profile:

```
[edit class-of-service]
user@switch# set drop-profile drop-profile-name interpolate fill-level percentage fill-level
percentage drop-probability 0 drop-probability percentage
```

### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring WRED Drop Profiles on page 279](#)
- [Defining CoS Queue Schedulers on page 410](#)
- [Configuring CoS Drop Profile Maps on page 408](#)
- [Understanding CoS WRED Drop Profiles on page 149](#)

## Configuring CoS Drop Profile Maps

---

A drop-profile map associates a WRED profile for traffic of a specified loss priority with a scheduler. When you use a scheduler map to map a scheduler to a forwarding class, the drop profile map associated with the scheduler applies the specified WRED profile to traffic in the forwarding class that matches the specified loss priority.

Drop profile maps enable you to configure different drop profiles for traffic of different loss priorities within the same scheduler. You can associate different drop profiles with low-priority, medium-high priority, and high-priority traffic within a single scheduler, and then map that scheduler to a forwarding class. This applies the appropriate drop profile to traffic of each loss priority in a forwarding class. Drop profile maps apply to all traffic protocols.

To configure a drop-profile map using the CLI:

- For the desired scheduler, configure the traffic loss priority and specify the drop profile you want to use to control the drop characteristics for traffic of that loss priority:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name drop-profile-map loss-priority level protocol
any drop-profile drop-profile-name
```

### Related Documentation

- [Example: Configuring Drop Profile Maps on page 281](#)
- [Configuring CoS WRED Drop Profiles on page 407](#)
- [Defining CoS Queue Schedulers on page 410](#)
- [Understanding CoS WRED Drop Profiles on page 149](#)

## Defining CoS Forwarding Classes

---

Forwarding classes allow you to group packets for transmission. The switch supports a total of 12 forwarding classes. In order to forward traffic, you map (assign) the forwarding classes to unicast or multdestination (multicast, broadcast, and destination lookup fail) output queues.

The switch has 12 output queues. Queues 0 through 7 are for unicast traffic and queues 8 through 11 are for multicast traffic. Forwarding classes mapped to unicast queues must carry unicast traffic, and forwarding classes mapped to multdestination queues must carry multdestination traffic. There are four default unicast forwarding classes and one default multdestination forwarding class.

The default unicast forwarding classes are:

- **best-effort**—Best-effort traffic
- **fcoe**—Guaranteed delivery for FCoE traffic
- **no-loss**—Guaranteed delivery for TCP no-loss traffic
- **network-control**—Network control traffic

The default multidestination forwarding class is:

- **mcast**—Multidestination traffic

Map forwarding classes to queues using the **class** statement, which enables you to configure up to 12 forwarding classes. You can map more than one forwarding class to a single queue, but all forwarding classes mapped to a particular queue must be of the same type, either unicast or multicast. In addition, all forwarding classes mapped to a particular queue must be either lossless or lossy. You cannot mix lossless and lossy forwarding classes (traffic) on the same queue. Also, you cannot mix unicast and multicast forwarding classes on the same queue.

[edit class-of-service forwarding-classes]

```
user@switch# class class-name queue-num queue-number <no-loss>
```



**NOTE:** If you are using Junos OS Release 12.2 or later, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best-effort) traffic and does *not* receive lossless treatment unless you include the optional no-loss packet drop attribute introduced in Junos OS Release 12.3 in the forwarding class configuration..



**NOTE:** Junos OS Release 11.3R1 and earlier supported an alternate method of mapping forwarding classes to queues that allowed you to map only one forwarding class to a queue using the statement:

[edit class-of-service forwarding-classes]

```
user@switch# queue queue-number class-name
```

The **queue** statement has been deprecated and is no longer valid in Junos OS Release 11.3R2 and later. If you have a configuration that uses the **queue** statement to map forwarding classes to queues, edit the configuration to replace the **queue** statement with the **class** statement.

#### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Forwarding Classes on page 283](#)
- [Monitoring CoS Forwarding Classes on page 546](#)
- [Understanding CoS Forwarding Classes on page 68](#)

## Defining CoS Forwarding Class Sets

---

A forwarding class set is a priority group for enhanced transmission selection (ETS) traffic control. Each forwarding class set consists of one or more forwarding classes (priorities, which can also be considered as output queues).

You can configure up to three unicast forwarding class sets and one multicast forwarding class set.

To configure a forwarding class set using the CLI:

1. Assign one or more forwarding classes to the forwarding class set:

```
[edit class-of-service]
user@switch# set forwarding-class-sets forwarding-class-set-name class
forwarding-class-name
```

2. Map the forwarding class set to an interface:

```
[edit class-of-service]
user@switch# set interfaces interface-name forwarding-class-set forwarding-class-set-name
```

### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Forwarding Class Sets on page 286](#)
- [Defining CoS Queue Schedulers on page 410](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 416](#)
- [Understanding CoS Forwarding Class Sets \(Priority Groups\) on page 73](#)

## Defining CoS Queue Schedulers

---

Schedulers define the CoS properties of output queues. These properties include the amount of interface bandwidth assigned to the queue, the priority of the queue, whether explicit congestion notification (ECN) is enabled on the queue, the WRED packet drop profiles associated with the queue, and the queue buffer size.

The parameters you configure in a scheduler define the following characteristics for the queues mapped to the scheduler:

- **transmit-rate**—Minimum bandwidth, also known as the committed information rate (CIR), set as a percentage rate or as an absolute value in bits per second. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue.



**NOTE:** Include the preamble bytes and interframe gap (IFG) bytes as well as the data bytes in your bandwidth calculations.



**NOTE:** You cannot configure a transmit rate for strict-high priority queues. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.

- **shaping-rate**—Maximum bandwidth, also known as the peak information rate (PIR), set as a percentage rate or as an absolute value in bits per second.



**NOTE:** Include the preamble bytes and interframe gap (IFG) bytes as well as the data bytes in your bandwidth calculations.

- **priority**—One of two bandwidth priorities that queues associated with a scheduler can receive:

- **low**—The scheduler has low priority.
- **strict-high**—The scheduler has strict-high priority. You can configure only one queue as a strict-high priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.

We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- **drop-profile-map**—Drop profile mapping to a loss priority and protocol to apply WRED to the scheduler.
- **buffer-size**—Size of the queue buffer as a percentage of the dedicated buffer space on the port, or as a proportional share of the dedicated buffer space on the port that remains after the explicitly configured queues are served.
- **explicit-congestion-notification**—Enables ECN on a best-effort queue. ECN enables end-to-end congestion notification between two ECN-enabled endpoints on TCP/IP based networks. ECN must be enabled on both endpoints and on all of the intermediate devices between the endpoints for ECN to work properly. ECN is disabled by default.



**NOTE:** Ingress port congestion can occur during periods of egress port congestion if an ingress port forwards traffic to more than one egress port, and at least one of those egress ports experiences congestion. If this occurs, the congested egress port can cause the ingress port to exceed its fair allocation of ingress buffer resources. When the ingress port exceeds its buffer resource allocation, frames are dropped at the ingress. Ingress port frame drop affects not only the congested egress ports, but also all of the egress ports to which the congested ingress port forwards traffic.

If a congested ingress port drops traffic that is destined for one or more uncongested egress ports, configure a weighted random early detection (WRED) drop profile and apply it to the egress queue that is causing the congestion. The drop profile prevents the congested egress queue from affecting egress queues on other ports by dropping frames at the egress instead of causing congestion at the ingress port.



**NOTE:** Do not configure drop profiles for the fcoe and no-loss forwarding classes. FCoE and other lossless traffic queues require lossless behavior. Use priority-based flow control (PFC) to prevent frame drop on lossless priorities.

To apply scheduling properties to traffic, map schedulers to forwarding classes using a scheduler map, and then associate the scheduler map with the interfaces. This applies the configured scheduling to the traffic in the specified forwarding class on the associated interface. Using different scheduler maps, you can map different schedulers to the same traffic (the same forwarding class) to apply different scheduling to that traffic on different interfaces.

To configure a scheduler using the CLI:

1. Name the scheduler and define the minimum guaranteed bandwidth for the queue:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name transmit-rate (rate | percent percentage)
```

2. Define the maximum bandwidth for the queue:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set shaping-rate (rate | percent percentage)
```

3. Define the queue priority:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set priority level
```

4. Define the drop profile using a drop profile map:

```
[edit class-of-service schedulers scheduler-name]
user@switch# set drop-profile-map loss-priority (low | medium-high | high) protocol protocol
drop-profile drop-profile-name
```

5. Configure the size of the port dedicated buffer space for the queue:

```
[edit class-of-service schedulers scheduler-name]
```



```
user@switch# set buffer-size percent 20
```

6. Enable ECN, if desired (queue should handle best-effort traffic):

```
[edit class-of-service schedulers scheduler-name]
user@switch# set explicit-congestion-notification
```

7. Configure a scheduler map to map the scheduler to a forwarding class, which applies the scheduler's properties to the traffic in that forwarding class:

```
[edit class-of-service]
user@switch# set scheduler-maps scheduler-map-name forwarding-class
forwarding-class-name scheduler scheduler-name
```

8. Assign the scheduler map and its associated schedulers to one or more interfaces using hierarchical scheduling. See [“Example: Configuring CoS Hierarchical Port Scheduling \(ETS\)” on page 220](#) for a detailed example of hierarchical scheduling.

```
[edit class-of-service]
user@switch# set traffic-control-profiles tcp-name scheduler-map scheduler-map-name
user@switch# set interfaces interface-name forwarding-class-set fc-set-name
output-traffic-control-profile tcp-name
```

#### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 300](#)
- [Example: Configuring Maximum Output Bandwidth on page 305](#)
- [Example: Configuring ECN](#)
- [Defining CoS Queue Scheduling Priority on page 414](#)
- [Configuring CoS WRED Drop Profiles on page 407](#)
- [Monitoring CoS Scheduler Maps on page 549](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Understanding CoS Priority Group Scheduling on page 95](#)
- [Understanding CoS Buffer Configuration on page 131](#)
- [Understanding CoS Explicit Congestion Notification](#)

## Defining CoS Queue Scheduling Priority

---

You can configure the scheduling priority of individual queues by specifying the priority in a scheduler, and then associating the scheduler with a queue by using a scheduler map. Queues can have one of two bandwidth priorities:

- **strict-high** —The scheduler has strict-high priority. You can configure only one queue as a strict-high priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.

We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.

- **low**—Low priority. Traffic with this priority is serviced after any queue that has a **strict-high** priority.
- To configure queue priority using the CLI:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name priority level
```

### Related Documentation

- [Example: Configuring Queue Scheduling Priority on page 295](#)
- [Defining CoS Queue Schedulers on page 410](#)
- [Monitoring CoS Scheduler Maps on page 549](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)

## Changing the Host Outbound Traffic Default Queue Mapping

If you do not want to use the default mapping of host Routing Engine and CPU outbound traffic to queues, you can change the default output queue. You can also change the default DSCP bits used in the type of service (ToS) field of packets generated by the Routing Engine.

Configuring a queue for host outbound traffic maps all traffic that the host generates to one forwarding class (queue). The configuration is global and applies to all host-generated traffic on the switch. Configuring a forwarding class for host outbound traffic does not affect transit or incoming traffic.



**NOTE:** Fibre Channel over Ethernet (FCoE) Initialization Protocol (FIP) packets generated by the CPU are always transmitted on the `fcoe` queue (queue 3), even if you configure a queue for host outbound traffic. This helps to ensure lossless behavior for FCoE traffic. QFabric systems classify FIP control packets into the same traffic class (`fcoe`) across the Interconnect device (`fabric`) and the egress Node device.

To change the host outbound traffic egress queue by including the `host-outbound-traffic` statement at the `[edit class-of-service]` hierarchy level:

```
[edit class-of-service]
host-outbound-traffic {
 forwarding-class class-name;
 dscp-code-point code-point;
}
```

For example, to map host outbound traffic to queue 7 (the network control forwarding class) and set the DSCP code point value to 101010:

```
[edit class-of-service]
host-outbound-traffic {
 forwarding-class network-control;
 dscp-code-point 101010
}
```

### Related Documentation

- [Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 47](#)

## Defining CoS Traffic Control Profiles (Priority Group Scheduling)

A traffic control profile defines the output bandwidth and scheduling characteristics of forwarding class sets (priority groups). The forwarding classes (queues) contained in a forwarding class set share the bandwidth resources that you configure in the traffic control profile. A scheduler map associates forwarding classes with schedulers to define how the individual queues in a forwarding class set share the bandwidth allocated to that forwarding class set.

The parameters you configure in a traffic control profile define the following characteristics for the priority group:

- **guaranteed-rate**—Minimum bandwidth, also known as the committed information rate (CIR). The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.



**NOTE:** You cannot configure a guaranteed rate for a forwarding class set (priority group) that includes strict-high priority queues. If the traffic control profile is for a forwarding class set that contains strict-high priority queues, do not configure a guaranteed rate.

- **shaping-rate**—Maximum bandwidth, also known as the peak information rate (PIR).
- **scheduler-map**—Bandwidth and scheduling characteristics for the queues, defined by mapping forwarding classes to schedulers. (The queue scheduling characteristics represent amounts or percentages of the priority group bandwidth, not the amounts or percentages of total link bandwidth.)



**NOTE:** Because a port can have more than one priority group, when you assign resources to a priority group, keep in mind that the total port bandwidth must serve all of the queues associated with that port.

To configure a traffic control profile using the CLI:

1. Name the traffic control profile and define the minimum guaranteed bandwidth for the priority group:
 

```
[edit class-of-service]
user@switch# set traffic-control-profiles traffic-control-profile-name guaranteed-rate (rate | percent percentage)
```
2. Define the maximum bandwidth for the priority group:
 

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
user@switch# set shaping-rate (rate | percent percentage)
```
3. Attach a scheduler map to the traffic control profile:
 

```
[edit class-of-service traffic-control-profiles traffic-control-profile-name]
user@switch# set scheduler-map scheduler-map-name
```

### Related Documentation

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
- [Example: Configuring Minimum Guaranteed Output Bandwidth on page 300](#)
- [Example: Configuring Maximum Output Bandwidth on page 305](#)
- [Defining CoS Queue Schedulers on page 410](#)
- [Understanding CoS Traffic Control Profiles on page 98](#)

## Configuring CoS PFC (Congestion Notification Profiles)

A congestion notification profile (CNP) enables priority-based flow control (PFC) on specified IEEE 802.1p priorities (code points). A CNP has two components:

- Input CNP:
  - Enable PFC on a specified priority.
  - Configure the maximum receive unit (MRU) on an interface for traffic that matches the PFC priority (optional).
  - Specify the length of the attached cable on the ingress interface (optional)
- Output CNP (optional): Configure flow control to enable PFC pause on specific output queues for specified priorities.



**NOTE:** By default, output queues 3 and 4 (which are mapped to default lossless forwarding classes `fcoe` and `no-loss`, respectively) are configured to respond to PFC pause messages received from the connected peer on priorities 3 and 4 (code points 011 and 100, respectively). If you explicitly configure flow control on any output queue, you must configure flow control on every output queue that you want to respond to pause messages. (The explicit configuration overrides the default configuration.)

To achieve lossless behavior, the output queue priorities on which you enable PFC flow control must match the PFC priorities on which you enable PFC on the input interfaces. For example, if you program output queues to pause priorities 3 (011) and 5 (101) in the output component of the CNP, then you must also enable pause on priorities 3 and 5 on the input component of the CNP. (In addition, the forwarding classes mapped to the paused output queues must be lossless forwarding classes.)

Associating a CNP with an interface enables PFC on the ingress traffic that matches the priority specified in the input CNP, and programs the queues listed in the output CNP to pause when the interface receives a PFC pause message from the connected peer. Configure PFC on a priority end to end along the entire data path to create a lossless lane of traffic on the network.



**NOTE:** You must enable PFC on the priority used by FCoE traffic on ingress interfaces (input CNP). Enable PFC on the FCoE priority on every interface that carries FCoE traffic. By convention, FCoE traffic uses priority 3 (code point 011), which maps to queue 3. If your network uses priority 3 for FCoE traffic, the default forwarding class and classifier configuration support lossless transport, but you must still configure a CNP and apply it to the correct ingress interfaces to enable PFC and achieve lossless transport.

If your network does not use priority 3 for FCoE traffic, you need to configure a classifier that classifies FCoE traffic into a lossless forwarding class, based on the priority your network uses for FCoE traffic. If you are not using the default lossless forwarding class configuration, then you also need to ensure that the output queue mapped to the lossless FCoE forwarding class is programmed to pause.

---

You can attach only one CNP to an interface. There is no limit to the total number of CNPs you can create.

Configuring a CNP consists of:

- Naming the CNP.
- Specifying the IEEE 802.1 code point (priority) on which you want to enable PFC on ingress interfaces (input CNP).
- Optionally, specifying the MRU and the length of the attached cable on ingress interfaces (input CNP).
- Optionally, configuring flow control (PFC pause) on specified output queues if you want queues other than queues 3 and 4 to respond to pause messages received from the connected peer (output CNP).
- Mapping the CNP to an interface.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

---

1. Enable PFC on the desired priority in the input CNP and optionally configure the interface MRU for traffic on that priority:

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name input ieee-802.1 code-point
code-point bits pfc mru mru-value
```

For example, to configure a CNP named **fcoe-cnp** that enables PFC on IEEE 802.1 code point 011 and configures an MRU value of 2240:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input ieee-802.1 code-point 011
pfc mru 2240
```

2. Configure the length of the cable attached to the ingress interface (optional):

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name input cable-length
cable-length-value
```

For example, to configure a CNP named **fcoe-cnp** that sets the length of the ingress interface cable to 100 meters:

```
[edit class-of-service]
user@switch# set congestion-notification-profile fcoe-cnp input cable-length 100
```

3. (Optional) Configure flow control on output queues:

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name output ieee-802.1 code-point
code-point-bits flow-control-queue [queue | list-of-queues]
```

For example, to configure a CNP named **fcoe-cnp** that enables PFC pause flow control on output queues 3 and 5 for FCoE traffic that uses priority 3 (code point 011) and on output queue 4 for traffic that uses priority 4 (code point 100):

```
[edit class-of-service]
user@switch# set congestion-notification-profile cnp-name output ieee-802.1 code-point
011 flow-control-queue [3 5]
user@switch# set congestion-notification-profile cnp-name output ieee-802.1 code-point
100 flow-control-queue 4
```

4. Map the CNP to an interface:

```
[edit class-of-service]
user@switch# set interfaces interface congestion-notification-profile cnp-name
```

For example, to map the CNP **fcoe-cnp** to the interface **xe-0/0/7**:

```
[edit class-of-service]
user@switch# set interfaces xe-0/0/7 congestion-notification-profile fcoe-cnp
```

#### Related Documentation

- [Example: Configuring CoS PFC for FCoE Traffic on page 241](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 356](#)
- [Assigning CoS Components to Interfaces on page 429](#)
- [Monitoring Interfaces That Have CoS Components on page 547](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)

- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109](#)

## Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control

---

Ethernet PAUSE flow control is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link, including Ethernet links that belong to Ethernet link aggregated (LAG) interfaces. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to PAUSE messages it receives from the connected peer to stop sending traffic.

Symmetric flow control means that an interface has the same PAUSE configuration in both directions. The PAUSE generation and PAUSE response functions are both configured as enabled, or they are both disabled.

Asymmetric flow control allows you to configure the PAUSE functionality in each direction independently on an interface. The configuration for generating PAUSE messages and for responding to PAUSE messages does not have to be the same. It can be enabled in both directions, disabled in both directions, or enabled in one direction and disabled in the other direction. If you do not want to PAUSE all of the traffic on a link, you can use priority-based flow control (PFC) to selectively pause traffic based on its IEEE 802.1p code point.

On any particular interface, symmetric and asymmetric flow control are mutually exclusive. If you attempt to configure both features, the switch returns a commit error. Ethernet PAUSE and PFC are also mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.

By default, all flow control features are disabled. You enable symmetric flow control on the interfaces on which you want to PAUSE all of the traffic on a link.

- To enable symmetric flow control on an interface:

```
[edit interfaces interface-name ether-options]
user@switch# set flow-control
```

- To disable symmetric flow control on an interface:

```
[edit interfaces interface-name ether-options]
user@switch# set no-flow-control
```

### Related Documentation

- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 421](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)



---

## Configuring CoS Asymmetric Ethernet PAUSE Flow Control

---

Ethernet PAUSE flow control is a congestion relief feature that works by providing link-level flow control for all traffic on a full-duplex Ethernet link, including Ethernet links that belong to link aggregated (LAG) interfaces. Ethernet PAUSE works in both directions on the link. In one direction, an interface generates and sends PAUSE messages to stop the connected peer from sending more traffic. In the other direction, the interface responds to PAUSE messages it receives from the connected peer to stop sending traffic.

Asymmetric flow control allows you to configure the PAUSE functionality in each direction independently on an interface. The configuration for generating PAUSE messages and for responding to PAUSE messages does not have to be the same. It can be enabled in both directions, disabled in both directions, or enabled in one direction and disabled in the other direction.

Symmetric flow control means that the interface has the same configuration in both directions. The PAUSE generation and PAUSE response functions are both configured as enabled or they are both disabled. If you do not want to PAUSE all of the traffic on a link, you can use priority-based flow control (PFC) to selectively pause traffic based on its IEEE 802.1p code point.

Asymmetric flow control provides the ability to configure the receive buffer and transmit buffer Ethernet PAUSE actions independently on an interface. The buffers perform the following actions:

- The receive buffers generate and send PAUSE messages to the connected peer to ask the peer to stop sending traffic for a time period specified in the PAUSE frame. The peer interface's buffers may store outgoing frames until the PAUSE period elapses and the interface can resume sending traffic.
- The transmit buffers respond to PAUSE messages received from the connected peer to stop sending traffic to the peer. The transmit buffer may store outgoing frames until the PAUSE period elapses and the interface can resume sending traffic.

Asymmetric flow control enables you to specify independently whether or not the interface receive buffer generates and sends PAUSE messages to stop the connected peer from transmitting traffic, and whether or not the interface transmit buffer responds to PAUSE messages it receives from the connected peer and stops transmitting traffic. The receive buffer configuration determines if the interface transmits PAUSE messages, and the transmit buffer configuration determines if the interface receives and responds to PAUSE messages:

- Receive buffers on—Enable PAUSE transmission (generate and send PAUSE frames)
- Transmit buffers on—Enable PAUSE reception (respond to received PAUSE frames)

You must explicitly set both the receive buffer and the transmit buffer to configure asymmetric flow control.

- To configure asymmetric flow control on an interface:  
`[edit interfaces interface-name ether-options]`

```
user@switch# set configured-flow-control rx-buffers (on | off) tx-buffers (on | off)
```

For example, to configure interface **xe-0/0/24** to generate and send PAUSE messages but not to respond to received PAUSE messages:

```
set interfaces xe-0/0/24 ether-options configured-flow-control rx-buffers on tx-buffers off
```

For example, to configure interface **xe-0/0/30** to respond to received PAUSE messages but not to generate and send PAUSE messages:

```
set interfaces xe-0/0/30 ether-options configured-flow-control rx-buffers off tx-buffers on
```



**NOTE:** If you configure both buffers to be on, that is equivalent to symmetric flow control. If you configure both buffers to be off, there is no flow control (flow control is disabled).

---

**Related  
Documentation**

- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 420](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)

## Configuring CoS Fixed Classifier Rewrite Values for Native FC Interfaces (NP\_Ports)

Fibre Channel over Ethernet (FCoE) traffic typically uses IEEE 802.1p priority 3 (code point 011). When Fibre Channel (FC) traffic arrives on a native FC interface (NP\_Port) on an FCoE-FC gateway, the interface encapsulates the FC traffic in Ethernet to create FCoE frames. By default, the native FC interface assigns priority 3 to the FCoE traffic. The traffic is then forwarded internally to the gateway Ethernet interfaces, and then forwarded to the FCoE network.

If your FCoE network uses priority 3 for FCoE traffic, you do not need to use a rewrite value to remap the FCoE priority on native FC interfaces, because the default configuration maps priority 3 to the FCoE forwarding class.

However, if the FCoE network uses a different priority than priority 3 for FCoE traffic, then you can configure a rewrite value to remap incoming traffic from the FC SAN to that priority after the interface encapsulates the FC packets in Ethernet. Setting a rewrite value for the IEEE 802.1p code point (priority) configures the gateway native FC interface to assign the rewrite value to the encapsulated FCoE frames before forwarding the FCoE frames to the gateway Ethernet interface. Instead of a priority of 3, the FCoE frames use the priority specified in the rewrite value.

Traffic coming from the FC SAN is classified into a lossless forwarding class, and that lossless forwarding class is mapped to the rewrite value (the priority used for FCoE traffic on the converged Ethernet network). You specify the lossless forwarding class used for FCoE traffic on a native FC interface by configuring a fixed classifier and applying it to the native FC interface. (The same forwarding class must also be mapped to the rewrite value priority in the ingress classifier applied to the FCoE Ethernet interfaces.) All traffic received from the FC SAN on that FC interface is encapsulated in Ethernet, classified into the forwarding class specified in the fixed classifier, and assigned the rewrite value priority.

Configuring a rewrite value consists of:

- Configuring a fixed classifier on the native FC interface. The fixed classifier assigns all the traffic that arrives at the interface from the connected peer in the FC SAN to one fixed forwarding class. The forwarding class must be a lossless forwarding class and must be classified to the rewrite value priority in the ingress classifier configuration on the FCoE Ethernet interfaces.
- Specifying an IEEE 802.1p rewrite value for the native FC interface. The traffic mapped to the forwarding class in the fixed classifier is marked with the priority you specify in the rewrite value when the traffic is encapsulated in Ethernet. The rewrite value must be the IEEE 802.1p priority used for FCoE traffic in your converged Ethernet network.

You can configure one rewrite value for each local FCoE-FC gateway fabric. All of the native FC interfaces in a particular fabric must use the same rewrite value. Native FC interfaces that belong to different FCoE-FC gateway fabrics can use different rewrite values.

1. Configure a fixed classifier on the native FC interface:

[edit [class-of-service](#)]

```
user@switch# set interfaces fc-interface-name forwarding-class
lossless-forwarding-class-name
```

For example, to configure a fixed classifier on native FC interface **fc-0/0/2** that specifies the lossless forwarding class **fcoe1**:

```
[edit class-of-service]
user@switch# set interfaces fc-0/0/2 forwarding-class fcoe1
```

2. Configure a rewrite value for the traffic classified into the fixed classifier (this must be the IEEE 802.1p priority used for the traffic on your converged Ethernet network):

```
[edit class-of-service]
user@switch# set interfaces fc-interface-name rewrite-value input ieee-802.1 code-point
code-point-bits
```

For example, to configure a rewrite value on native FC interface **fc-0/0/2** that specifies an IEEE 802.1p priority of **101** (the lossless forwarding class specified in the fixed classifier must be classified to this priority in the ingress classifier configuration on the FCoE Ethernet interfaces):

```
[edit class-of-service]
user@switch# set interfaces fc-0/0/2 rewrite-value input ieee-802.1 code-point 101
```

In the example, all traffic from the FC SAN that arrives at FCoE-FC gateway interface **fc-0/0/2** is encapsulated in Ethernet, classified into the lossless **fcoe1** forwarding class, and tagged with the IEEE 802.1p priority 5 (code point 101). In this example, we assume that the converged Ethernet network uses priority 5 for FCoE traffic, and that the **fcoe1** forwarding class is mapped to priority 5 in the ingress classifier configuration on the Ethernet interfaces. To achieve lossless transport, you must also enable PFC on priority 5 on the Ethernet interfaces that connect the FCoE traffic to the Ethernet network.

#### Related Documentation

- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 356](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 128](#)

## Configuring Global Ingress and Egress Shared Buffers

Although the switch reserves some buffer space to ensure a minimum memory allocation for ports and queues, you can configure how the system uses the rest of the buffer space to optimize the buffer allocation for your particular mix of network traffic. The global shared buffer pool is memory space that all of the ports on the switch share dynamically as they need buffers. You can allocate global shared memory space to different types of ingress and egress buffers to better support different mixes of network traffic.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

Use the default shared buffer settings (for a network with a balanced mix of lossless, best-effort, and multicast traffic) or one of the recommended shared buffer configurations for your mix of network traffic (mostly best-effort unicast traffic, mostly best-effort traffic on links enabled for Ethernet PAUSE, mostly multicast traffic, or mostly lossless traffic). Either the default configuration or one of the recommended configurations provides a buffer allocation that satisfies the needs of most networks.

After starting from one of the recommended configurations, you can fine-tune the shared buffer settings, but do so with caution to prevent traffic loss due to buffer misconfiguration.

You can configure the percentage of available (user-configurable) buffer space allocated to the global shared buffers. Any space that you do not allocate to the global shared buffer pool is added to the dedicated buffer pool. The default configuration allocates 100 percent of the available buffer space to the global shared buffers.

You can partition the ingress and egress shared buffer pools to allocate more buffers to the types of traffic your network predominantly carries, and fewer buffers to other traffic. From the buffer space allocated to the ingress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless ingress traffic. The minimum value for the lossless buffers is 5 percent.
- Lossless headroom buffers—Percentage of shared buffer pool for packets received while a pause is asserted. If Ethernet PAUSE is configured on a port or if priority-based flow control (PFC) is configured on priorities on a port, when the port sends a pause message to the connected peer, the port uses the headroom buffers to store the packets that arrive between the time the port sends the pause message and the time the last packet arrives after the peer pauses traffic. The minimum value for the lossless headroom buffers is 0 (zero) percent. (Lossless headroom buffers are the only buffers that can have a minimum value of less than 5 percent.)
- Lossy buffers—Percentage of shared buffer pool for all best-effort ingress traffic (best-effort unicast, multidestination, and strict-high priority traffic). The minimum value for the lossy buffers is 5 percent.

The combined percentage values of the ingress lossless, lossless headroom, and lossy buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All ingress buffer partitions must be explicitly configured, even when the lossless headroom buffer partition has a value of 0 (zero) percent.

From the buffer space allocated to the egress shared buffer pool, you can allocate space to:

- Lossless buffers—Percentage of shared buffer pool for all lossless egress queues. The minimum value for the lossless buffers is 5 percent.
- Lossy buffers—Percentage of shared buffer pool for all best-effort egress queues (best-effort unicast, and strict-high priority queues). The minimum value for the lossy buffers is 5 percent.
- Multicast buffers—Percentage of shared buffer pool for all multidestination (multicast, broadcast, and destination lookup fail) egress queues. The minimum value for the multicast buffers is 5 percent.

The combined percentage values of the egress lossless, lossy, and multicast buffer partitions must total exactly 100 percent. If the buffer percentages total more than 100 percent or less than 100 percent, the switch returns a commit error. All egress buffer partitions must be explicitly configured and must have a value of at least 5 percent.

To configure the shared buffer allocation and partitioning using the CLI:

1. Configure the percentage of available (nonreserved) buffers used for the ingress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress percent percent
```
2. Configure the global ingress buffer partitions for lossless, lossless-headroom, and lossy traffic:  

```
[edit class-of-service shared-buffer]
user@switch# set ingress buffer-partition lossless percent percent
user@switch# set ingress buffer-partition lossless-headroom percent percent
user@switch# set ingress buffer-partition lossy percent percent
```
3. Configure the percentage of available (nonreserved) buffers used for the egress global shared buffer pool:  

```
[edit class-of-service shared-buffer]
user@switch# set egress percent percent
```
4. Configure the global egress buffer partitions for lossless, lossy, and multicast queues:  

```
[edit class-of-service shared-buffer]
user@switch# set egress buffer-partition lossless percent percent
user@switch# set egress buffer-partition lossy percent percent
user@switch# set egress buffer-partition multicast percent percent
```

**Related  
Documentation**

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 365](#)

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Traffic on Links with Ethernet PAUSE Enabled on page 371](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 377](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 383](#)
- [Understanding CoS Buffer Configuration on page 131](#)

## Defining CoS Rewrite Rules

You configure rewrite rules to alter CoS values in outgoing packets on the outbound interfaces of a switch to match the policies of a targeted peer. Policy matching allows the downstream routing platform or switch in a neighboring network to classify each packet into the appropriate service group.

To configure a CoS rewrite rule, create the rule by giving it a name and associating it with a forwarding class, loss priority, and code point. This creates a rewrite table. After the rewrite rule is created, enable it on an interface (EXP rewrite rules can only be enabled on **family mpls** logical interfaces, not on physical interfaces). You can also apply an existing rewrite rule on an interface.



**NOTE:** On each physical interface, either all forwarding classes that are being used on the interface must have rewrite rules configured, or no forwarding classes that are being used on the interface can have rewrite rules configured. On any physical port, do not mix forwarding classes with rewrite rules and forwarding classes without rewrite rules.



**NOTE:** To replace an existing rewrite rule on the interface with a new rewrite rule of the same type, first explicitly remove the existing rewrite rule and then apply the new rule.



**NOTE:** For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rule rewrites only the outer VLAN tag.

To create rewrite rules and enable them on interfaces:

- To create an 802.1p rewrite rule named **customup-rw** in the rewrite table for all Layer 2 interfaces:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority low code-point
000
user@switch# set ieee-802.1 customup-rw forwarding-class be loss-priority high code-point
001
```

```

user@switch# set ieee-802.1p customup-rw forwarding-class be loss-priority low code-point
010
user@switch# set ieee-802.1p customup-rw forwarding-class fc0e loss-priority low code-point
011
user@switch# set ieee-802.1p customup-rw forwarding-class ef-no-loss loss-priority low
code-point 100
user@switch# set ieee-802.1p customup-rw forwarding-class ef-no-loss loss-priority high
code-point 101
user@switch# set ieee-802.1p customup-rw forwarding-class nc loss-priority low code-point
110
user@switch# set ieee-802.1p customup-rw forwarding-class nc loss-priority high code-point
111

```

- To enable an 802.1p rewrite rule named **customup-rw** on a Layer 2 interface:

```

[edit]
user@switch# set class-of-service interfaces xe-0/0/7 unit 0 rewrite-rules ieee-802.1
customup-rw

```



**NOTE:** All forwarding classes assigned to port xe-0/0/7 must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same physical interface.

- To enable an 802.1p rewrite rule named **customup-rw** on all 10-Gigabit Ethernet interfaces on the switch, use wildcards for the interface name and logical interface (unit) number:

```

[edit]
user@switch# set class-of-service interfaces xe-* unit * rewrite-rules customup-rw

```



**NOTE:** In this case, *all* forwarding classes assigned to *all* 10-Gigabit Ethernet ports must have rewrite rules. Do not mix forwarding classes that have rewrite rules with forwarding classes that do not have rewrite rules on the same physical interface.

#### Related Documentation

- [Monitoring CoS Rewrite Rules on page 548](#)
- [Configuring Rewrite Rules for MPLS EXP Classifiers](#)
- [Configuring a Global MPLS EXP Classifier](#)
- [Understanding CoS Rewrite Rules on page 154](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57](#)



## Assigning CoS Components to Interfaces

After you define the following CoS components, you assign them to physical or logical interfaces. Components that you assign to physical interfaces are valid for all of the logical interfaces configured on the physical interface. Components that you assign to a logical interface are valid only for that logical interface.

- Classifiers—Assign only to logical interfaces.
- Congestion notification profiles—Assign only to physical interfaces.
- Forwarding classes—Assign to interfaces by mapping to forwarding class sets.
- Forwarding class sets—Assign only to physical interfaces.
- Output traffic control profiles—Assign only to physical interfaces (with a forwarding class set).
- Rewrite rules—Assign only to logical interfaces.

You can assign a CoS component to a single interface or to multiple interfaces using wildcards. You can also assign a congestion notification profile or a forwarding class set globally to all interfaces.

To assign CoS components to interfaces:

Assign CoS components to a single interface by associating a CoS component (for example a forwarding class set named **san-priority-group**) with an interface:

```
[edit class-of-service interfaces]
user@switch# set xe-0/0/7 forwarding-class-set san-priority-group
```

Assign a CoS component to multiple interfaces by associating a CoS component (for example, a rewrite rule named **customup-rw**) to all 10-Gigabit Ethernet interfaces on the switch, use wildcard characters for the interface name and logical interface (unit) number:

```
[edit class-of-service interfaces]
user@switch# set xe-* unit * rewrite-rules ieee-802.1 customup-rw
```

Assign a congestion notification profile or a forwarding class set globally to all interfaces using the **set class-of-service interfaces all** statement. For example, to assign a forwarding class set named **be\_fcset** to all interfaces:

```
[edit class-of-service interfaces]
user@switch# set all forwarding-class-set be_fcset
```



**NOTE:** If there is an existing CoS configuration of any type on an interface, the global configuration is not applied to that particular interface. The global configuration is applied to all interfaces that do not have an existing CoS configuration.

For example, if you configure a rewrite rule, assign it to interfaces `xe-0/0/20.0` and `xe-0/0/22.0`, and then configure a congestion notification profile and apply it to all interfaces, the congestion notification profile is applied to every interface except `xe-0/0/20` and `xe-0/0/22`.

#### Related Documentation

- [Monitoring Interfaces That Have CoS Components on page 547](#)
- [Understanding Junos CoS Components on page 30](#)
- [Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57](#)

## Configuring the DCBX Mode

You can configure the DCBX mode that an interface uses to communicate with the connected peer. QFX Systems support three DCBX modes:

- Autonegotiation—The interface negotiates with the connected peer to determine the DCBX mode. This is the default DCBX mode.
- IEEE DCBX—The interface uses IEEE DCBX type, length, and value (TLV) to exchange DCBX information with the connected peer. QFX3500 Node devices come up with IEEE DCBX enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.
- DCBX Version 1.01—The interface uses Converged Enhanced Ethernet (CEE) DCBX version 1.01 TLVs to exchange DCBX information with the connected peer. QFabric Node devices come up with DCBX version 1.01 enabled by default and then autonegotiate with the connected peer to determine the final DCBX mode.



**NOTE:** QFX Systems do not support pre-CEE (pre-DCB) versions of DCBX such as DCBX version 1.00. If a QFX Series interface receives an LLDP frame with pre-CEE DCBX TLVs, the system drops the frame.

Configure the DCBX mode by specifying the mode for one interface or for all interfaces.

- To configure the DCBX mode, specify the interface and the mode:

```
[edit protocols dcbx]
user@switch# set interface interface-name mode (auto-negotiate | ieee-dcbx |
dcbx-version-1.01)
```

For example, to configure DCBX version 1.01 on interface `xe-0/0/21`:

```
user@switch# set protocols dcbx interface xe-0/0/21 mode dcbx-version-1.01
```

To configure IEEE DCBX on all interfaces:

```
user@switch# set protocols dcbx interface all mode ieee-dcbx
```

#### Related Documentation

- [Configuring DCBX Autonegotiation on page 431](#)
- [Disabling the ETS Recommendation TLV on page 434](#)
- [Understanding DCBX on page 170](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 179](#)
- [show dcbx neighbors on page 625](#)

## Configuring DCBX Autonegotiation

Data Center Bridging Capability Exchange protocol (DCBX) discovers the data center bridging (DCB) capabilities of peers by exchanging feature configuration information. DCBX also detects feature misconfiguration and mismatches, and can configure DCB on peers. DCBX is an extension of the Link Layer Discovery Protocol (LLDP), and LLDP must remain enabled on every interface for which you want to use DCBX. If you attempt to enable DCBX on an interface on which LLDP is disabled, the configuration commit operation fails.



**NOTE:** LLDP and DCBX are enabled by default on all interfaces.

The switch supports DCBX autonegotiation for:

- Priority-based flow control (PFC) configuration
- Layer 2 and Layer 4 applications such as Fibre Channel over Ethernet (FCoE) and Internet Small Computer System Interface (iSCSI)
- Enhanced transmission selection (ETS) advertisement

DCBX autonegotiation is configured on a per-interface basis for each supported feature or application. The PFC and application DCBX exchanges use autonegotiation by default. The default autonegotiation behavior is:

- DCBX is enabled on the interface if the connected peer device also supports DCBX.
- DCBX is disabled on the interface if the connected peer device does not support DCBX.

You can override the default behavior for each feature by turning off autonegotiation to force an interface to enable or disable the feature.

Autonegotiation of ETS means that when ETS is enabled on an interface (priority groups are configured), the interface advertises its ETS configuration to the peer device. In this case, priorities (forwarding classes) that are not part of a priority group (forwarding class set) receive no bandwidth and are advertised in an automatically generated default forwarding class. If ETS is not enabled on an interface (no priority groups are configured),

all of the priorities are advertised in one automatically generated default priority group that receives 100 percent of the port bandwidth.

Disabling ETS autonegotiation prevents the interface from sending the Recommendation TLV or the Configuration TLV to the connected peer.

On interfaces that use IEEE DCBX mode to exchange DCBX parameters, you can disable autonegotiation of the enhanced transmission selection (ETS) Recommendation TLV to the peer if you want an asymmetric ETS configuration between the peers. DCBX still exchanges the ETS Configuration TLV if you disable the ETS Recommendation TLV.

Autonegotiation of PFC means that when PFC is enabled on an interface, if the peer device connected to the interface supports PFC and is provisioned compatibly with the switch, DCBX sets the PFC operational state to enabled. If the peer device connected to the interface does not support PFC or is not provisioned compatibly with the switch, DCBX sets the operational state to disabled.

In addition, if the peer advertises that it is “willing” to learn its PFC configuration from the switch, DCBX pushes the switch’s PFC configuration to the peer and does not check the peer’s administrative state. The switch does not learn PFC configuration from peers (the switch does not advertise its state as “willing”).

Disabling PFC autonegotiation prevents the interface from exchanging PFC configuration information with the peer. It forces the interface to enable PFC if PFC is configured on the interface or to disable PFC if PFC is not configured on the interface. If you disable PFC autonegotiation, the assumption is that the peer is also configured manually.

Autonegotiation of applications depends on whether or not you apply an application map to an interface. If you apply an application map to an interface, the interface autonegotiates DCBX for each application in the application map. PFC must be enabled on the FCoE priority (the FCoE IEEE 802.1p code point) for the interface to advertise the FCoE application. The interface only advertises applications that are included in the application map.

For example, if you apply an application map to an interface and the application map does not include the FCoE application, then that interface does not perform DCBX advertisement of FCoE.

If you do not apply an application map to an interface, DCBX does not advertise applications on that interface, with the exception of FCoE, which is handled differently than other applications.



**NOTE:** If you do not apply an application map to an interface, the interface performs autonegotiation of FCoE if the interface carries traffic in the FCoE forwarding class and also has PFC enabled on the FCoE priority. On such interfaces, if DCBX detects that the peer device connected to the interface supports FCoE, the switch advertises its FCoE capability and IEEE 802.1p code point on that interface. If DCBX detects that the peer device connected to the interface does not support FCoE, DCBX marks that interface as “FCoE down” and disables FCoE on the interface.

---

When DCBX marks an interface as “FCoE down,” the behavior of the switch depends on how you use it in the network:

- When the switch acts as an FCoE-FC gateway, it does not send or receive FCoE Initialization Protocol (FIP) packets.
- When the switch acts as an FCoE transit switch, the interface drops all of the FIP packets it receives. In addition, FIP packets received from an FCoE forwarder (FCF) are not forwarded to interfaces marked as “FCoE down.”

Disabling autonegotiation prevents the interface from exchanging application information with the peer. In this case, the assumption is that the peer is also configured manually.

To disable DCBX autonegotiation of PFC, applications (including FCoE), and ETS using the CLI:

1. Turn off autonegotiation for PFC.

```
[edit]
user@switch# set protocols dcbx interface interface-name priority-flow-control
no-auto-negotiation
```

2. Turn off autonegotiation for applications.

```
[edit]
user@switch# set protocols dcbx interface interface-name applications no-auto-negotiation
```

3. Turn off autonegotiation for ETS.

```
[edit]
user@switch# set protocols dcbx interface interface-name enhanced-transmission-selection
no-auto-negotiation
```

To disable autonegotiation of the ETS Recommendation TLV so that DCBX exchanges only the ETS Configuration TLV:

- [edit protocols dcbx interface *interface-name*]  
user@switch# set enhanced-transmission-selection no-recommendation-tlv

#### Related Documentation

- [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 241](#)
- [Disabling the ETS Recommendation TLV on page 434](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 179](#)

## Disabling the ETS Recommendation TLV

---

The enhanced transmission selection (ETS) Recommendation TLV communicates the ETS settings that the switch wants the connected peer interface to use. If the peer interface is “willing,” the peer interface changes its configuration to match the configuration in the ETS Recommendation TLV. By default, the switch interfaces send the ETS Recommendation TLV to the peer. The settings communicated are the egress ETS settings defined by configuring hierarchical scheduling on the interface.

We recommend that you use the same ETS settings on the connected peer that you use on the switch interface and that you leave the ETS Recommendation TLV enabled. However, on interfaces that use IEEE DCBX as the DCBX mode, if you want an asymmetric configuration between the switch interface and the connected peer, you can disable the ETS Recommendation TLV.



**NOTE:** Disabling the ETS Recommendation TLV on interfaces that use DCBX version 1.01 as the DCBX mode has no effect and does not change DCBX behavior.

If you disable the ETS Recommendation TLV, the switch still sends the ETS Configuration TLV to the connected peer. The result is that the connected peer is informed about the switch DCBX ETS configuration, but even if the peer is “willing,” the peer does not change its configuration to match the switch configuration. This is asymmetric configuration—the two interfaces can have different parameter values for the ETS attribute.

To disable the ETS Recommendation TLV:

- [edit protocols dcbx interface *interface-name*]  
user@switch# **set enhanced-transmission-selection no-recommendation-tlv**

### Related Documentation

- [Configuring the DCBX Mode on page 430](#)
- [Configuring DCBX Autonegotiation on page 431](#)
- [Understanding DCBX on page 170](#)
- [Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches](#)

## Defining an Application for DCBX Application Protocol TLV Exchange

---

Define each application for which you want DCBX to exchange application protocol information. You can define Layer 2 and Layer 4 applications. After you define applications, you map them to IEEE 802.1p code points, and then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to configure application maps and apply them to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

Define Layer 2 applications by mapping an application name to an EtherType. Define Layer 4 applications by mapping an application name to a protocol (TCP or UDP) and a destination port.

- To define a Layer 2 application, specify the name of the application and its EtherType:

```
[edit applications]
user@switch# set application application-name ether-type ether-type
```

For example, to configure an application named **PTP** (for Precision Time Protocol) that uses the EtherType **0x88F7**:

```
user@switch# set applications application ptp ether-type 0x88F7
```

- To define a Layer 4 application, specify the name of the application, its protocol (TCP or UDP), and its destination port:

```
[edit]
user@switch# set applications application application-name protocol (tcp | udp)
destination-port port-value
```

For example, to configure an application named **iscsi** (for Internet Small Computer System Interface) that uses the protocol **TCP** and the destination port **3260**:

```
user@switch# set applications application iscsi protocol tcp destination-port 3260
```

#### Related Documentation

- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 436](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 437](#)
- [Configuring DCBX Autonegotiation on page 431](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 179](#)
- [show dcbx neighbors on page 625](#)

## Configuring an Application Map for DCBX Application Protocol TLV Exchange

After you define applications for which you want to exchange DCBX application protocol information, map the applications to IEEE 802.1p code points. The IEEE 802.1p code points identify incoming traffic and allow you to map that traffic to the desired application. You then apply the application map to the interfaces on which you want DCBX to exchange application protocol information with connected peers. (See *Related Documentation* for how to define applications and apply the application map to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

Configure an application map by creating an application map name and mapping an application to one or more IEEE 802.1p code points.

- To define an application map, specify the name of the application map, the name of the application, and the IEEE 802.1p code points of the incoming traffic that you want to associate with the application in the application map:

```
[edit policy-options]
user@switch# set application-maps application-map-name application application-name
code-points [aliases] [bit-patterns]
```

For example, to configure an application map named **ptp-app-map** that includes an application named **PTP** (for Precision Time Protocol) and map the application to IEEE 802.1p code points **001** and **101**:

```
user@switch# set policy-options application-maps ptp-app-map application ptp code points
[001 101]
```

### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 434](#)
- [Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 437](#)
- [Configuring DCBX Autonegotiation on page 431](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 625](#)



## Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange

After you define applications and map them to IEEE 802.1p code points in an application map, apply the application map to the interfaces on which you want DCBX to exchange the application protocol information with connected peers. (See *Related Documentation* for how to define applications and configure application maps to interfaces, and for an example of the entire procedure that also includes classifier configuration.)



**NOTE:** In Junos OS Release 12.1, the FCoE application was configured by default, so you did not need to configure it in an application map. In Junos OS Release 12.2, if you want DCBX to advertise the FCoE application on an interface and you apply an application map to that interface, you must explicitly configure FCoE in the application map. You also must enable priority-based flow control (PFC) on the FCoE code point on all interfaces that you want to advertise FCoE. If you apply an application map to an interface, the interface sends DCBX TLVs only for the applications configured in the application map.

- To apply an application map to a DCBX interface, specify the DCBX interface and the application map name:

[edit protocols]

```
user@switch# set dcbx interface interface-name application-map application-map-name
```

For example, to apply an application map named **ptp-app-map** on interface **xe-0/0/11**:

```
user@switch# set protocols dcbx interface xe-0/0/11 application-map ptp-app-map
```

### Related Documentation

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 434](#)
- [Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 436](#)
- [Configuring DCBX Autonegotiation on page 431](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [show dcbx neighbors on page 625](#)



## CHAPTER 7

# Configuration Statements

- [application \(Application Maps\) on page 442](#)
- [application \(Applications\) on page 443](#)
- [application-map on page 444](#)
- [application-maps on page 445](#)
- [applications \(Applications\) on page 446](#)
- [applications \(DCBX\) on page 447](#)
- [buffer-partition \(Egress\) on page 448](#)
- [buffer-partition \(Ingress\) on page 450](#)
- [buffer-size on page 452](#)
- [cable-length \(Congestion Notification\) on page 454](#)
- [class-of-service on page 455](#)
- [class \(Forwarding Classes\) on page 459](#)
- [class \(Forwarding Class Sets\) on page 460](#)
- [classifiers on page 461](#)
- [code-point \(Fibre Channel Interfaces\) on page 463](#)
- [code-point \(Input Congestion Notification\) on page 464](#)
- [code-point \(Output Congestion Notification\) on page 465](#)
- [code-point \(Rewrite Rules\) on page 466](#)
- [code-point-aliases on page 466](#)
- [code-points \(Application Maps\) on page 467](#)
- [code-points \(CoS\) on page 467](#)
- [configured-flow-control on page 468](#)
- [congestion-notification-profile on page 469](#)
- [dcbx on page 471](#)
- [dcbx-version on page 472](#)
- [destination-port \(Applications\) on page 473](#)
- [disable \(DCBX\) on page 474](#)
- [drop-probability on page 475](#)

- [drop-profile on page 476](#)
- [drop-profile-map on page 476](#)
- [drop-profiles on page 477](#)
- [dscp on page 478](#)
- [dscp-ipv6 on page 480](#)
- [dscp-code-point on page 481](#)
- [egress \(Buffer Configuration\) on page 482](#)
- [enhanced-transmission-selection on page 483](#)
- [ether-type on page 484](#)
- [fill-level on page 485](#)
- [flow-control on page 486](#)
- [flow-control-queue \(Output Congestion Notification\) on page 487](#)
- [forwarding-class on page 489](#)
- [forwarding-class \(Fibre Channel Interfaces\) on page 491](#)
- [forwarding-class \(Host Outbound Traffic\) on page 492](#)
- [forwarding-class-set on page 493](#)
- [forwarding-class-sets on page 493](#)
- [forwarding-classes on page 494](#)
- [guaranteed-rate on page 496](#)
- [host-outbound-traffic on page 497](#)
- [ieee-802.1 on page 498](#)
- [ieee-802.1 \(Fibre Channel Interfaces\) on page 500](#)
- [ieee-802.1 \(Input Congestion Notification\) on page 501](#)
- [ieee-802.1 \(Output Congestion Notification\) on page 502](#)
- [import on page 503](#)
- [ingress \(Buffer Configuration\) on page 504](#)
- [input \(Congestion Notification\) on page 505](#)
- [input \(Fibre Channel Interfaces\) on page 506](#)
- [interface \(DCBX\) on page 507](#)
- [interfaces \(Class of Service\) on page 508](#)
- [interpolate on page 509](#)
- [loss-priority \(Classifiers\) on page 510](#)
- [loss-priority \(Drop Profiles\) on page 511](#)
- [loss-priority \(Rewrite Rules\) on page 512](#)
- [multi-destination on page 513](#)
- [mru on page 514](#)
- [output \(Congestion Notification\) on page 515](#)

- [output-traffic-control-profile](#) on page 516
- [pfc](#) (Input Congestion Notification) on page 517
- [policy-options](#) on page 518
- [priority](#) (Schedulers) on page 519
- [priority-flow-control](#) on page 520
- [protocol](#) (Applications) on page 521
- [protocol](#) (Drop Profile Map) on page 522
- [queue-num](#) on page 523
- [recommendation-tlv](#) on page 524
- [rewrite-rules](#) on page 525
- [rewrite-value](#) (Fibre Channel Interfaces) on page 526
- [rx-buffers](#) on page 528
- [scheduler](#) on page 529
- [scheduler-map](#) on page 529
- [scheduler-maps](#) on page 530
- [schedulers](#) on page 531
- [shaping-rate](#) on page 532
- [shared-buffer](#) on page 534
- [traceoptions](#) (Class of Service) on page 536
- [traffic-control-profiles](#) on page 538
- [transmit-rate](#) on page 539
- [tx-buffers](#) on page 541
- [unit](#) on page 542

## application (Application Maps)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>application <i>application-name</i> {<br/>    <i>code-points</i> [ <i>aliases</i> ] [ <i>bit-patterns</i> ];<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit policy-options <b>application-maps</b> <i>application-map-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Add an application to an application map and define the application's code points.                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b><i>application-name</i></b>—Name of the application.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 436</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 389</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 179</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## application (Applications)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> application <i>application-name</i> {     <i>destination-port</i> <i>port-value</i>;     <i>protocol</i> (tcp   udp);     <i>ether-type</i> <i>type</i>; } </pre>                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit applications]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Configure properties to define an application.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <p><i>application-name</i>—Name of the application.</p> <p>The statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 434</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 389</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 179</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## application-map

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>application-map <i>application-map-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit protocols <code>dcbx interface interface-name</code> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify an application map to apply to an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <i>application-map-name</i> —Name of the application map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 625</a></li><li>• <a href="#">Applying an Application Map to an Interface for DCBX Application Protocol TLV Exchange on page 437</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 389</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 179</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |



## application-maps

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> application-maps <i>application-map-name</i> {   application <i>application-name</i> {     code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];   } } </pre>                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit policy-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 12.1 for EX Series switches.</p> <p>Statement introduced in Junos OS Release 12.1 for the QFX Series.</p>                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Define an application map by specifying the applications that belong to the application map.                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | <p><b><i>application-map-name</i></b>—Name of the application map.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 436</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 389</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 179</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## applications (Applications)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>applications {<br/>  application application-name {<br/>    destination-port port-value;<br/>    protocol (tcp   udp);<br/>    ether-type type;<br/>  }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Define applications that DCBX advertises.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 434</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 389</a></li><li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 179</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li></ul> |

## applications (DCBX)

---

|                                 |                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>applications {<br/>    no-auto-negotiation;<br/>}</pre>                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">dcbx</a> <a href="#">interface</a> <i>interface-name</i> ]                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 12.1 for the EX Series                                                |
| <b>Description</b>              | Configure Data Center Bridging Capability Exchange protocol (DCBX) applications on an interface.                                                                                    |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                  |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 625</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 156</a></li></ul> |

## buffer-partition (Egress)

**Syntax** `buffer-partition (lossless | lossy | multicast) {  
percent percent;  
}`

**Hierarchy Level** [edit [class-of-service shared-buffer egress](#)]

**Release Information** Statement introduced in Junos OS Release 12.3 for the QFX Series.

**Description** The egress shared buffer pool is divided into three partitions. Each partition reserves a percentage of the available shared buffer pool for a type of traffic, so that the switch provides enough resources to support a mix of best-effort, lossless, and multicast traffic (multicast also includes broadcast and destination lookup fail traffic). To better support the mix of traffic on your network, you can optimize the allocation of egress shared buffers to different types of traffic by fine-tuning the shared buffer partitions.

The percentages you configure for the three egress shared buffer partitions must total exactly 100 percent. If the total of the three shared buffer percentages is not 100 percent, the system returns a commit error and does not commit the configuration. You can configure any partition to 0 (zero) percent as long as the allocation to other partitions totals 100 percent.

This is a global allocation that applies to all ports. All ports on the switch receive the same allocation of egress shared buffers.

If you do not configure buffer partitions, the switch uses the default partitioning.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

**Default** The default egress buffer partition shown in [Table 108 on page 448](#) supports networks with a balanced mix of best-effort, multicast, and lossless traffic. It is the recommended configuration if you are using the default configuration with two lossless forwarding classes.

**Table 108: Default Egress Shared Buffer Partitioning**

| Lossless Partition | Lossy Partition | Multicast Partition |
|--------------------|-----------------|---------------------|
| 50%                | 31%             | 19%                 |

The sum of the default percentages configured for each partition is 100 percent. The sum of the partition percentages must always total 100 percent.

**Options** **lossless**—Shared buffer space reserved for all lossless egress traffic.

**lossy**—Shared buffer space for best-effort unicast egress traffic.

**multicast**—Shared buffer space reserved for all multicast (including broadcast and destination lookup fail) egress traffic.

**percent percent**—The percentage of buffer space to allocate to the specified buffer partition (lossless, lossy, or multicast buffers). The sum of the percentages for the three buffer partitions must total 100 percent.

|                           |                                                               |
|---------------------------|---------------------------------------------------------------|
| <b>Required Privilege</b> | interfaces—To view this statement in the configuration.       |
| <b>Level</b>              | interface-control—To add this statement to the configuration. |

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 365</a></li><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 377</a></li><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 383</a></li><li>• <a href="#">Configuring Global Ingress and Egress Shared Buffers on page 425</a></li><li>• <a href="#">Understanding CoS Buffer Configuration on page 131</a></li></ul> |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## buffer-partition (Ingress)

**Syntax** `buffer-partition (lossless | lossless-headroom | lossy) {  
percent percentage;  
}`

**Hierarchy Level** [edit [class-of-service shared-buffer ingress](#)]

**Release Information** Statement introduced in Junos OS Release 12.3 for the QFX Series.

**Description** The ingress shared buffer pool is divided into three partitions. Each partition reserves a percentage of the available shared buffer pool for a type of traffic, so that the switch provides enough resources to support a mix of best effort (best-effort unicast and multicast) and lossless traffic. To better support the mix of traffic on your network, you can optimize the allocation of ingress shared buffers to different types of traffic by fine-tuning the shared buffer partitions.

The percentages you configure for the three ingress shared buffer partitions must total exactly 100 percent. If the total of the three shared buffer percentages is not 100 percent, the system returns a commit error and does not commit the configuration. You can configure any partition to 0 (zero) percent as long as the allocation to other partitions totals 100 percent.

This is a global allocation that applies to all ingress traffic. All ports on the switch receive the same allocation of ingress shared buffers.

If you do not configure buffer partitions, the switch uses the default partitioning.



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until buffer reprogramming is complete.

**Default** The default ingress buffer partition shown in [Table 109 on page 450](#) supports networks with a balanced mix of best-effort, multicast, and lossless traffic. It is the recommended configuration if you are using the default configuration with two lossless forwarding classes.

**Table 109: Default Ingress Shared Buffer Partitioning**

| Lossless Partition | Lossless-Headroom Partition | Lossy Partition |
|--------------------|-----------------------------|-----------------|
| 9%                 | 45%                         | 46%             |

The sum of the default percentages configured for each partition is 100 percent. The sum of the partition percentages always must total 100 percent.

**Options** **lossless**—Shared buffer space reserved for all lossless ingress traffic.

**lossless-headroom**—Shared buffer space reserved to store packets received while either an 802.3x Ethernet PAUSE or a priority-based flow control (PFC) pause is asserted. (When an ingress interface pauses traffic, it must have the buffer space to store all of the packets currently in the buffer, and also all of the packets received before the connected peer stops sending traffic and the wire is cleared of packets.)

**lossy**—Shared buffer space for best-effort ingress traffic.

**percent *percent***—The percentage of buffer space to allocate to the specified buffer partition (lossless, lossless-headroom, or lossy buffers). The sum of the percentages for the three buffer partitions must total 100 percent.

|                           |                                                               |
|---------------------------|---------------------------------------------------------------|
| <b>Required Privilege</b> | interfaces—To view this statement in the configuration.       |
| <b>Level</b>              | interface-control—To add this statement to the configuration. |

- |                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 365</a></li><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 377</a></li><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 383</a></li><li>• <a href="#">Configuring Global Ingress and Egress Shared Buffers on page 425</a></li><li>• <a href="#">Understanding CoS Buffer Configuration on page 131</a></li></ul> |
|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## buffer-size

---

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>buffer-size (percent <i>percent</i>   remainder);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>     | [edit <code>class-of-service schedulers <i>scheduler-name</i></code> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>         | <p>Set the dedicated buffer size of the egress queue that you bind the scheduler to in the scheduler map configuration. The switch allocates space from the global dedicated buffer pool to ports and queues in a hierarchical manner. The switch allocates an equal number of dedicated buffers to each egress port, so each egress port receives the same amount of dedicated buffer space. The amount of dedicated buffer space per port is not configurable.</p> <p>However, the <b>buffer-size</b> statement allows you to control the way each port allocates its share of dedicated buffers to its queues. For example, if a port only uses two queues to forward traffic, you can configure the port to allocate all of its dedicated buffer space to those two ports and avoid wasting buffer space on queues that are not in use. We recommend that the buffer size should be the same size as the minimum guaranteed transmission rate (the <b>transmit-rate</b>).</p> <p>You configure the proportion of port dedicated buffers allocated to a particular output queue using the following process:</p> <ol style="list-style-type: none"><li>1. Configure a scheduler and set the <b>buffer-size</b> option to match the scheduler <b>transmit-rate</b> value.</li><li>2. Use a scheduler map to map the scheduler to the forwarding class that is mapped to the queue to which you want to apply the buffer size.</li></ol> <p>For example, suppose that you want to change the dedicated buffer allocation for FCoE traffic. FCoE traffic is mapped to the <code>fcoe</code> forwarding class, and the <code>fcoe</code> forwarding class is mapped to queue 3 (this is the default configuration). To use default FCoE traffic mapping, in the scheduler map configuration, map the scheduler to the <b>fcoe</b> forwarding class.</p> <ol style="list-style-type: none"><li>3. Associate the scheduler map with the traffic control profile you want to use on the egress ports that carry FCoE traffic.</li><li>4. Associate the traffic control profile that includes the scheduler map with the desired egress ports. For this example, you associate the traffic control profile with the ports that carry FCoE traffic.</li></ol> <p>Queue 3, which is mapped to the <code>fcoe</code> forwarding class and therefore to the FCoE traffic, receives the dedicated buffer allocation specified in the <b>buffer-size</b> statement.</p> |



**NOTE:** The total of all of the explicitly configured buffer size percentages for all of the queues on a port cannot exceed 100 percent.

---



**Default** The port allocates dedicated buffers to queues that have an explicitly configured scheduler buffer size. If you do not explicitly configure a scheduler buffer size for a queue, the port serves the explicitly configured queues first. Then the port divides the remaining dedicated buffers equally among the queues that have an explicitly attached scheduler *without* an explicitly configured buffer size configuration. (If you configure a scheduler, but you do not configure the buffer size parameter, the default is equivalent to configuring the buffer size with the **remainder** option.)

If you use the default scheduler and scheduler map on a port (no explicit scheduler configuration), then the port allocates its dedicated buffer pool to queues based on the default scheduling, as shown in [Table 110 on page 453](#). The default buffer size is the same as the default transmit rate for each default queue:

**Table 110: Default Output Queue Buffer Sizes**

| Queue Number | Forwarding Class | Transmit Rate | Buffer Size |
|--------------|------------------|---------------|-------------|
| 0            | best-effort      | 5%            | 5%          |
| 3            | fcoe             | 35%           | 35%         |
| 4            | no-loss          | 35%           | 35%         |
| 7            | network-control  | 5%            | 5%          |
| 8            | mcast            | 20%           | 20%         |

Because the default scheduler includes only five forwarding classes, only the queues mapped to those forwarding classes receive dedicated buffers from the port buffer pool. (Buffers are not wasted on queues that do not carry traffic.)

**Options** **percent percent**—Percentage of the port dedicated buffer pool allocated to the queue (or queues) mapped to the scheduler.

**remainder**—Remaining dedicated buffer pool after the port satisfies the needs of the explicitly configured buffers. The port divides the remaining buffers equally among the queues that are explicitly attached to a scheduler but that do not have an explicit buffer size configuration (or are configured with **remainder** as the buffer size).

**Required Privilege** interfaces—To view this statement in the configuration.


**Level** interface-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 365](#)
  - [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 377](#)
  - [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 383](#)
  - [Understanding CoS Buffer Configuration on page 131](#)

---

## cable-length (Congestion Notification)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>cable-length <i>cable-length-value</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service congestion-notification-profile <i>profile-name</i></a> <a href="#">input</a> ]                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Specify the length of the cable between the interface and its peer interface in meters. The system uses the cable length and the maximum receive unit (MRU) to calculate the amount of buffer headroom reserved to support priority-based flow control (PFC). The the shorter the cable length and lower the MRU, the less headroom buffer space is required for PFC.</p>                                                                                                                         |
|                                 | <div> <b>NOTE:</b> You can also set a maximum transmission unit (MTU) value (the largest packet size the interface sends) for interfaces by including the <code>mtu</code> statement at the [edit <a href="#">interfaces <i>interface-name</i></a>] hierarchy level.</div>                                                                                                                                        |
| <b>Default</b>                  | The default cable length value is 100 meters (approximately 328 feet).                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <code><i>cable-length-value</i></code> —Length of the cable in meters. (Generally from 1 to 300 meters, but there is no configuration restriction.)                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | <code>interfaces</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 417</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326</a></li><li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 160</a></li><li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109</a></li></ul> |

## class-of-service

```
Syntax class-of-service {
 classifiers {
 (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name {
 import (classifier-name | default);
 forwarding-class class-name {
 loss-priority level {
 code-points [aliases] [bit-patterns];
 }
 }
 }
 }
 code-point-aliases {
 (dscp | dscp-ipv6 | ieee-802.1) {
 alias-name bits;
 }
 }
 congestion-notification-profile profile-name {
 input {
 ieee-802.1 {
 code-point [code-point-bits] {
 pfc {
 mru mru-value;
 }
 }
 }
 cable-length cable-length-value;
 }
 output {
 ieee-802.1 {
 code-point [code-point-bits] {
 flow-control-queue [queue | list-of-queues];
 }
 }
 }
 }
 drop-profiles {
 profile-name {
 interpolate {
 fill-level low-value fill-level high-value drop-probability 0 drop-probability high-value;
 }
 }
 }
 forwarding-class class-name {
 loss-priority level {
 code-points [aliases] [bit-patterns];
 }
 }
 forwarding-class class-name {
 scheduler scheduler-name;
 }
 forwarding-class-sets forwarding-class-set-name {
 class class-name;
 }
 }
```

```

}
forwarding-classes {
 class {
 class-name {
 queue-num queue-number <no-loss>;
 }
 }
}
host-outbound-traffic {
 forwarding-class class-name;
 dscp-code-point code-point;
}
interfaces {
 interface-name {
 congestion-notification-profile profile-name {
 }
 forwarding-class lossless-forwarding-class-name;
 forwarding-class-set forwarding-class-set-name {
 output-traffic-control-profile profile-name;
 }
 rewrite-value {
 input {
 ieee-802.1 {
 code-point code-point-bits;
 }
 }
 }
 unit logical-unit-number {
 classifiers {
 (dscp | dscp-ipv6 | ieee-802.1 exp) (classifier-name | default);
 }
 forwarding-class class-name;
 rewrite-rules {
 (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);
 }
 }
 }
}
multi-destination {
 classifiers {
 (dscp | ieee-802.1) classifier-name;
 }
}
rewrite-rules {
 (dscp | dscp-ipv6 | ieee-802.1 | exp) classifier-name {
 import (rewrite-name | default);
 forwarding-class class-name {
 loss-priority priority code-point (alias | bits);
 }
 }
}
scheduler-maps {
 map-name {
 forwarding-class class-name scheduler scheduler-name;
 }
}

```

```

schedulers {
 scheduler-name {
 buffer-size (percent percentage | remainder);
 drop-profile-map loss-priority (low | medium-high | high) protocol protocol drop-profile
 drop-profile-name;
 explicit-congestion-notification;
 priority priority;
 shaping-rate (rate | percent percentage);
 transmit-rate (percent percentage);
 }
}
shared-buffer {
 egress {
 percent percent;
 buffer-partition (lossless | lossy | multicast) {
 percent percent
 }
 }
 ingress {
 percent percent;
 buffer-partition (lossless | lossless-headroom | lossy) {
 percent percent
 }
 }
}
system-defaults {
 classifiers exp classifier-name;
}
traffic-control-profiles profile-name {
 guaranteed-rate(rate| percent percentage);
 scheduler-map map-name;
 shaping-rate (rate| percent percentage);
}
}

```

Hierarchy Level [\[edit\]](#)

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure class-of-service parameters on the switch.

The remaining statements are explained separately.


**Default** If you do not configure any CoS features, the default CoS settings are used.

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related  
Documentation**

- [Assigning CoS Components to Interfaces on page 429](#)
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 404](#)
- [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 406](#)
- [\*Configuring a Global MPLS EXP Classifier\*](#)
- [Defining CoS Code-Point Aliases on page 404](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
- [\*Example: Configuring ECN\*](#)
- [Configuring CoS Drop Profile Maps on page 408](#)
- [Defining CoS Forwarding Class Sets on page 410](#)
- [Defining CoS Forwarding Classes on page 408](#)
- [\*Configuring Rewrite Rules for MPLS EXP Classifiers\*](#)
- [Defining CoS Rewrite Rules on page 427](#)
- [Defining CoS Queue Schedulers on page 410](#)
- [Configuring CoS WRED Drop Profiles on page 407](#)
- [Defining CoS Traffic Control Profiles \(Priority Group Scheduling\) on page 416](#)
- [Overview of Junos OS CoS for the QFX Series on page 4](#)

## class (Forwarding Classes)

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |                                                                                                                                                                                                                                                                                                                                        |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <pre>class {   class-name {     queue-num queue-number &lt;no-loss&gt;;   } }</pre>                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | [edit <a href="#">class-of-service forwarding-classes</a> ]                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>No-loss option introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                            |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Map one or more forwarding classes to a single queue. You can map unicast forwarding classes to a unicast queue (0 through 7) and multdestination forwarding classes to a multicast queue (8 through 11). The queue to which you map a forwarding class determines if the forwarding class is a unicast or multicast forwarding class. |
| <div>  <p><b>NOTE:</b> If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does <i>not</i> receive lossless treatment.</p> <p>If you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the fcoe and no-loss forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the no-loss option. If you do not specify the no-loss option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the fcoe forwarding class and you do not include the no-loss option, the fcoe forwarding class is lossy, not lossless.</p> </div> |                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <p><b>class-name</b> —Name of the forwarding class.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                    |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Forwarding Classes on page 283</a></li> <li>• <a href="#">Understanding CoS Forwarding Classes on page 68</a></li> </ul>                                                                                                                                     |

## class (Forwarding Class Sets)

---

|                                 |                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>class <i>class-name</i>;</code>                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service forwarding-class-sets</a> <i>forwarding-class-set-name</i> ]                                                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Group forwarding classes into sets of forwarding classes (priority groups). You can group some or all of the configured forwarding classes into up to three unicast forwarding class sets and one multidestination forwarding class set.                                                                                  |
| <b>Options</b>                  | <i>class-name</i> —Name of the forwarding class.                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220</a></li><li>• <a href="#">Example: Configuring Forwarding Class Sets on page 286</a></li><li>• <a href="#">Understanding CoS Forwarding Class Sets (Priority Groups) on page 73</a></li></ul> |



## classifiers

|                                                                                  |                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                                            | <a href="#">Syntax (BA Classifiers) on page 461</a><br><a href="#">Syntax (Multidestination BA Classifiers) on page 461</a><br><a href="#">Syntax (Interface Classifier Association: DSCP, DSCP IPv6, IEEE) on page 461</a><br><a href="#">Syntax (Global EXP Interface Classifier Association with Interfaces) on page 461</a> |
| <b>Syntax (BA Classifiers)</b>                                                   | <pre> classifiers {   (dscp   dscp-ipv6   ieee-802.1   exp) classifier-name {     import (classifier-name   default);     forwarding-class class-name {       loss-priority level {         code-points [ aliases ] [ bit-patterns ];       }     }   } } </pre>                                                                |
| <b>Syntax (Multidestination BA Classifiers)</b>                                  | <pre> classifiers {   (dscp   ieee-802.1) classifier-name; } </pre>                                                                                                                                                                                                                                                             |
| <b>Syntax (Interface Classifier Association: DSCP, DSCP IPv6, IEEE)</b>          | <pre> classifiers {   (dscp   dscp-ipv6   ieee-802.1) (default   classifier-name); } </pre>                                                                                                                                                                                                                                     |
| <b>Syntax (Global EXP Interface Classifier Association with Interfaces)</b>      | <pre> classifiers {   exp classifier-name; } </pre>                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level (BA Classifiers)</b>                                          | [edit <a href="#">class-of-service</a> ],                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level (Multidestination BA Classifiers)</b>                         | [edit <a href="#">class-of-service multi-destination</a> ],                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level (Interface Classifier Association: DSCP, DSCP IPv6, IEEE)</b> | [edit <a href="#">class-of-service interfaces</a> interface-name unit logical-unit-number]                                                                                                                                                                                                                                      |
| <b>Hierarchy Level (Global EXP Classifier)</b>                                   | [edit <a href="#">class-of-service</a> system-defaults]                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>                                                       | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>EXP statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                      |
| <b>Description</b>                                                               | Define a unicast or multidestination CoS behavior aggregate (BA) classifier.                                                                                                                                                                                                                                                    |
| <b>Options</b>                                                                   | The statements are explained separately.                                                                                                                                                                                                                                                                                        |

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 404](#)
- [Configuring a Global MPLS EXP Classifier](#)
- [Example: Configuring Unicast Classifiers on page 273](#)
- [Example: Configuring Multidestination \(Multicast, Broadcast, DLF\) Classifiers on page 277](#)
- [Understanding CoS Classifiers on page 51](#)
- [Understanding CoS MPLS EXP Classifiers and Rewrite Rules](#)

## code-point (Fibre Channel Interfaces)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>code-point <i>code-point-bits</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service interfaces</a> <i>fibre-channel-interface-name</i> <a href="#">rewrite-value</a> <a href="#">input ieee-802.1</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Configure the IEEE 802.1p code point value assigned to all traffic received from the Fibre Channel (FC) network on the specified FC interface (NP_Port). When native FC traffic from the FC SAN arrives at the NP_Port interface, the NP_Port interface encapsulates it in Ethernet to create FCoE packets before forwarding the traffic onto the FCoE network. Instead of using the default value of priority 3 (code point 011) for the FCoE traffic, the interface rewrites the IEEE 802.1p code point to the value specified in the rewrite value code points.</p> <p>After the code point value is rewritten, the interface forwards the traffic to the Ethernet (FCoE) network. This works in conjunction with configuring a fixed classifier on the FC interface. The fixed classifier maps all traffic from the FC network into one lossless forwarding class (the lossless forwarding class must be mapped to the code point specified in the rewrite value). Traffic mapped to the lossless forwarding class uses the IEEE 802.1p priority specified by the code point bits in the rewrite value.</p> <p>FCoE traffic typically uses priority 3 (IEEE code point 011). The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. Rewriting the code point value enables you to change the IEEE 802.1p priority of the FCoE traffic if the Ethernet network uses a different priority than priority 3 (code point 011).</p> <p>The system supports only one IEEE 802.1p code point value per FC interface. You cannot configure more than one IEEE 802.1p rewrite value per FC interface. In addition, you can specify only one rewrite value per local FCoE-FC gateway fabric; all interfaces in the local fabric must use the same rewrite value. Attempting to configure FC interfaces in the same local fabric with different rewrite values generates a commit error. You can specify different rewrite values for interfaces that belong to different local FCoE-FC gateway fabrics.</p> |
| <b>Options</b>                  | <i>code-point-bits</i> —Value of the code-point bits, in decimal form.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">forwarding-class (Fibre Channel Interfaces) on page 491</a></li> <li>• <a href="#">Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 356</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 128</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 51</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## code-point (Input Congestion Notification)

---

|                                 |                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>code-point [<i>code-point-bits</i>] {<br/>    pfc {<br/>        mru <i>mru-value</i>;<br/>    }<br/>}</pre>                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service congestion-notification-profile</b> <i>profile-name</i> <b>input</b> <b>ieee-802.1</b> ]                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Enable priority-based flow control (PFC) on an IEEE 802.1p code point (priority).                                                                                                                                                                                                                                      |
| <b>Options</b>                  | <p><b><i>code-point-bits</i></b>—3-bit value in decimal form.</p> <p>The remaining statements are described separately.</p>                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 241</a></li><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 417</a></li><li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 160</a></li></ul> |

## code-point (Output Congestion Notification)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>code-point [ <i>code-point-bits</i> ] {<br/>    <i>flow-control-queue</i> [ <i>queue</i>   <i>list-of-queues</i> ];<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <i>class-of-service congestion-notification-profile profile-name output ieee-802.1</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Specify the IEEE 802.1p code point bits that identify the traffic you want to enable for priority-based flow control (PFC) pause.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Default</b>                  | <p>By default, IEEE 802.1p priorities 3 and 4 (code points 011 and 100, respectively) are enabled for PFC pause on all Ethernet interfaces. If you explicitly configure priorities to pause and the output queues on which to enable pause, the explicit configuration overrides the default configuration. When you apply an explicit output congestion notification profile to an interface, only the priorities and queues specified in the output congestion notification profile are enabled for pause on that interface.</p> <p>For example, if you configure an output congestion notification profile that specifies priority 2 (code point 010), then traffic with IEEE 802.1p priority 2 is paused on the configured output queue during periods of congestion. However, traffic with priority 3 and priority 4 is not programmed to pause, because the explicit configuration overwrites the default configuration, and the explicit configuration does not pause priority 3 and priority 4. If you configure an explicit output congestion notification profile, all of the priorities you want to enable for PFC and all of the output queues you want to pause must be explicitly configured.</p> |
| <b>Options</b>                  | <p><i>code-point-bits</i>—3-bit value in decimal form.</p> <p>The remaining statements are described separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 417</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 317</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 309</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 340</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109</a></li> </ul>                                                                                                                                                                                                                                                                                            |

## code-point (Rewrite Rules)

---

|                                 |                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>code-point [ <i>alias</i> ] [ <i>bit-pattern</i> ];</code>                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service rewrite-rules</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a> ) <a href="#">forwarding-class</a> <a href="#">class-name</a> <a href="#">loss-priority</a> <a href="#">level</a> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                        |
| <b>Description</b>              | Configure a code-point alias or bit set to apply to a forwarding class for a rewrite rule.                                                                                                                                                               |
| <b>Options</b>                  | <i>alias</i> —Name of the alias.<br><br><i>bit-pattern</i> —Value of the code-point bits, in decimal form.                                                                                                                                               |
| <b>Required Privilege Level</b> | <i>interfaces</i> —To view this statement in the configuration.<br><i>interface-control</i> —To add this statement to the configuration.                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Defining CoS Rewrite Rules on page 427</a></li><li>• <a href="#">Understanding CoS Classifiers on page 51</a></li></ul>                                                                              |

## code-point-aliases

---

|                                 |                                                                                                                                                                                                                                    |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>code-point-aliases {<br/>  (<a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a>) {<br/>    <i>alias-name</i> <i>bits</i>;<br/>  }<br/>}</pre>                                                       |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                  |
| <b>Description</b>              | Define an alias for a CoS marker. You can use the alias instead of the bit pattern when you specify the code point during configuration.                                                                                           |
| <b>Options</b>                  | <i>(dscp   dscp-ipv6   ieee-802.1)</i> —Set the type of classifier for which you are creating an alias.<br><br><i>alias-name</i> —Name of the code-point alias.<br><br><i>bits</i> —Value of the code-point bits, in decimal form. |
| <b>Required Privilege Level</b> | <i>interfaces</i> —To view this statement in the configuration.<br><i>interface-control</i> —To add this statement to the configuration.                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Defining CoS Code-Point Aliases on page 404</a></li><li>• <a href="#">Understanding CoS Code-Point Aliases on page 49</a></li></ul>                                            |

## code-points (Application Maps)


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit policy-options <b>application-maps</b> <i>application-map-name</i> <b>application</b> <i>application-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Define one or more code-point aliases or bit sets for an application.                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | <i>aliases</i> —Name of the alias or aliases.<br><br><i>bit-patterns</i> —Value of the code-point bits, in decimal form.                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring an Application Map for DCBX Application Protocol TLV Exchange on page 436</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 389</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 179</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## code-points (CoS)

|                                 |                                                                                                                                                                                                              |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>code-points [ <i>aliases</i> ] [ <i>bit-patterns</i> ];</code>                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service</b> <b>classifiers</b> ( <b>dscp</b>   <b>dscp-ipv6</b>   <b>ieee-802.1</b> ) <i>classifier-name</i> <b>forwarding-class</b> <i>class-name</i> <b>loss-priority</b> <i>level</i> ] |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                            |
| <b>Description</b>              | Configure one or more code-point aliases or bit sets to apply to a forwarding class.                                                                                                                         |
| <b>Options</b>                  | <i>aliases</i> —Name of the alias or aliases.<br><br><i>bit-patterns</i> —Value of the code-point bits, in decimal form.                                                                                     |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Unicast Classifiers on page 273</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 51</a></li> </ul>                 |

## configured-flow-control

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>configured-flow-control {<br/>    rx-buffers (on   off);<br/>    tx-buffers (on   off);<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ether-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Configure Ethernet PAUSE asymmetric flow control on an interface. You can set an interface to generate and send PAUSE messages, and you can set an interface to respond to PAUSE messages sent by the connected peer. You must set both the <b>rx-buffers</b> and the <b>tx-buffers</b> values when you configure asymmetric flow control.</p> <p>Use the <b>flow-control</b> and <b>no-flow-control</b> statements to enable and disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> <div> <b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC) by applying a congestion notification profile to the interface.</div> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> |
| <b>Default</b>                  | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">congestion-notification-profile on page 469</a></li><li>• <a href="#">flow-control on page 486</a></li></ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |



## congestion-notification-profile

|                                                              |                                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                | <pre> congestion-notification-profile <i>profile-name</i> {   input {     ieee-802.1 {       code-point [<i>code-point-bits</i>] {         pfc {           mru <i>mru-value</i>;         }       }     }     cable-length <i>cable-length-value</i>;   }   output {     ieee-802.1 {       code-point [<i>code-point-bits</i>] {         flow-control-queue [<i>queue</i>   <i>list-of-queues</i>];       }     }   } } </pre> |
| <b>Interface Congestion Notification Profile Association</b> | <pre> congestion-notification-profile <i>profile-name</i> { </pre>                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>                                       | <pre> [edit <i>class-of-service</i>], [edit <i>class-of-service interfaces interface-name</i>] </pre>                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>                                   | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>                                           | Configure a congestion notification profile to enable priority-based flow control (PFC) on traffic specified by an IEEE 802.1 code point, and apply the profile to an interface.                                                                                                                                                                                                                                               |



**NOTE:** You must configure PFC for FCoE traffic. Each interface that carries FCoE traffic should be configured for PFC on the FCoE code point (usually 011).

You can attach a maximum of one congestion notification profile to an interface. There is no limit to the total number of congestion notification profiles you can create.



**NOTE:** Configuring or changing PFC on an interface blocks the entire port until the PFC change is completed. After a PFC change is completed, the port is unblocked and traffic resumes. Blocking the port stops ingress and egress traffic, and causes packet loss on all queues on the port until the port is unblocked.

**Options**    *profile-name*—Name of the congestion notification profile.

The remaining statements are explained separately.

**Required Privilege**    interface—To view this statement in the configuration.  
**Level**    interface-control—To add this statement to the configuration.

**Related Documentation**

- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 241](#)
- [Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326](#)
- [Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 317](#)
- [Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic \(FCoE Transit Switch\) on page 309](#)
- [Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications \(FCoE and iSCSI\) on page 340](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)
- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109](#)

## dcbx


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> dcbx {   disable;   interface (interface-name   all) {     disable;     application-map application-map-name;     applications {       no-auto-negotiation;     }     enhanced-transmission-selection {       no-auto-negotiation;       no-recommendation-tlv;       recommendation-tlv {         no-auto-negotiation;       }     }     dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);     priority-flow-control {       no-auto-negotiation;     }   } } </pre> |
| <b>Hierarchy Level</b>          | [edit protocols]                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for EX Series switches.</p> <p><b>mode</b> and <b>recommendation-tlv</b> statements introduced in Junos OS Release 12.2 for the QFX Series.</p>                                                                                                                                                                                                                |
| <b>Description</b>              | Configure DCBX properties.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 625</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 156</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 431</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on EX Series Switches</a></li> <li>• <a href="#">Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</a></li> </ul>               |

## dcbx-version

---

|                                 |                                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);</code>                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">dcbx interface</a> (all   <i>interface-name</i> )]                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                         |
| <b>Description</b>              | <p>Set the DCBX version for the specified interface or interfaces.</p> <p>QFX3500 switches come up in IEEE DCBX mode and then autonegotiate with the connected peer to set the DCBX version.</p> <p>QFabric system Node devices come up using DCBX version 1.01, and then autonegotiate with the connected peer to set the DCBX mode.</p> |
| <b>Default</b>                  | The default DCBX mode is autonegotiation.                                                                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b>auto-negotiate</b>—Automatically negotiate the DCBX version with the connected peer.</p> <p><b>ieee-dcbx</b>—Force the interface to use IEEE DCBX mode, regardless of the peer configuration.</p> <p><b>dcbx-version-1.01</b>—Force the interface to use version 1.01 DCBX mode, regardless of the peer configuration.</p>          |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 625</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 431</a></li><li>• <a href="#">Understanding DCBX on page 170</a></li></ul>                                                                                                         |

## destination-port (Applications)

|                                                                                                                                                                                                                                               |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                 | <code>destination-port <i>port-value</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                        | [edit applications <b>application</b> <i>application-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>                                                                                                                                                                                                                    | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>                                                                                                                                                                                                                            | Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) destination port number, which combines with <b>protocol</b> to identify an application type. The Internet Assigned Numbers Authority (IANA) assigns port numbers. See the IANA <i>Service Name and Transport Protocol Port Number Registry</i> at <a href="http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml">http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml</a> for a list of assigned port numbers. |
| <div>  <b>NOTE:</b> To create an application for iSCSI, use the protocol <code>tcp</code> with the destination port number <code>3260</code>.         </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                                                                                                                                                                                                                                | <i>port-value</i> —Identifier for the port.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b>                                                                                                                                                                                                               | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>                                                                                                                                                                                                                  | <ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 434</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 389</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 179</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul>                  |

## disable (DCBX)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | disable                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">dcbx</a> ]<br><br>[edit protocols <a href="#">dcbx interface</a> <i>interface-name</i> ]                                                                                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 11.3 for EX Series switches.                                                                                                                                                                                                                                               |
| <b>Description</b>              | Disable Data Center Bridging Capability Exchange protocol (DCBX) on one or more 10-Gigabit Ethernet interfaces.                                                                                                                                                                                                                                                                          |
| <b>Default</b>                  | DCBX is enabled by default on all 10-Gigabit or higher Ethernet interfaces.<br><br>DCBX is enabled by default on all 10-Gigabit Ethernet interfaces on EX4500 CEE-enabled switches.                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring DCBX Autonegotiation on page 431</a></li><li>• <i>Disabling DCBX to Disable PFC Autonegotiation on EX Series Switches (CLI Procedure)</i></li><li>• <a href="#">Understanding DCB Features and Requirements on page 156</a></li><li>• <i>Understanding DCB Features and Requirements on EX Series Switches</i></li></ul> |

## drop-probability

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>drop-probability 0 drop-probability <i>high-value</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service drop-profiles</a> <i>profile-name</i> <a href="#">interpolate</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>When configuring WRED, map the packet <b>drop-probability</b> to the fullness of a queue (<b>fill-level</b>). You configure the <b>fill-level</b> and <b>drop-probability</b> statements in related pairs by specifying a low <b>fill-level</b> value at which packets begin to drop (the drop probability is zero until the queue reaches this level of fullness) and a high <b>fill-level</b> value at which packets drop at the highest drop probability. As the queue fills from the low fill level to the high fill level, the rate of packet drop increases in a linear pattern from zero to the high drop probability.</p> |
| <b>Options</b>                  | <p><b>0</b>—Probability that packets will drop at the lowest <b>fill-level</b> value. This is always zero, because until the queue reaches the specified low <b>fill-level</b> value, no packets are scheduled to drop.</p> <p><b>high-value</b>—The maximum probability that packets will drop before queue fullness exceeds the high value of the queue <b>fill-level</b>, expressed as a percentage. If the queue fills beyond the high <b>fill-level</b> value, all packets drop.</p> <p><b>Range:</b> 0 through 100</p>                                                                                                         |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring WRED Drop Profiles on page 279</a></li> <li>• <a href="#">Understanding CoS WRED Drop Profiles on page 149</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                  |

## drop-profile

---

|                                 |                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>drop-profile <i>profile-name</i>;</code>                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service schedulers <i>scheduler-name</i></a> <a href="#">drop-profile-map <i>loss-priority</i></a> (low   medium-high   high) <a href="#">protocol <i>protocol</i></a> ]                                                                           |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                             |
| <b>Description</b>              | Define drop profiles for random early detection (RED). When a packet arrives, RED checks the queue fill level specified in the drop profile. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.   |
| <b>Options</b>                  | <i>profile-name</i> —Name of the drop profile.                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Drop Profile Maps on page 281</a></li><li>• <a href="#">Example: Configuring WRED Drop Profiles on page 279</a></li><li>• <a href="#">Understanding CoS WRED Drop Profiles on page 149</a></li></ul> |

## drop-profile-map

---

|                                 |                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>drop-profile-map <a href="#">loss-priority</a> (low   medium-high   high) <a href="#">protocol <i>protocol</i></a> <a href="#">drop-profile <i>drop-profile-name</i></a>;</code>                                                                                        |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service schedulers <i>scheduler-name</i></a> ]                                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                             |
| <b>Description</b>              | Map a drop profile to a loss priority and protocol for random early detection (RED). When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.   |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Drop Profile Maps on page 281</a></li><li>• <a href="#">Example: Configuring WRED Drop Profiles on page 279</a></li><li>• <a href="#">Understanding CoS WRED Drop Profiles on page 149</a></li></ul> |



## drop-profiles

|                                 |                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>drop-profiles {   profile-name {     interpolate {       fill-level low-value fill-level high-value drop-probability 0 drop-probability high-value;     }   } }</pre>                                                                                                                                                          |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | <p>Define drop profiles for weighted random early detection (WRED).</p> <p>For a packet to be dropped, it must match the drop profile. When a packet arrives, WRED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the WRED algorithm determines whether to drop the arriving packet.</p> |
| <b>Options</b>                  | <p><b>profile-name</b>—Name of the drop profile.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring WRED Drop Profiles on page 279</a></li> <li>• <a href="#">Understanding CoS WRED Drop Profiles on page 149</a></li> </ul>                                                                                                                                 |

## dscp

|                                                           |                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                     | <a href="#">Syntax (Classifier) on page 478</a><br><a href="#">Syntax (Code-Point Alias) on page 478</a><br><a href="#">Syntax (Multidestination Classifier) on page 478</a><br><a href="#">Syntax (Interface Classifier Association) on page 478</a><br><a href="#">Syntax (Rewrite Rule) on page 478</a> |
| <b>Syntax (Classifier)</b>                                | <pre>dscp classifier-name {   import (classifier-name   default);   forwarding-class class-name {     loss-priority level {       code-points [ aliases ] [ bit-patterns ];     }   } }</pre>                                                                                                              |
| <b>Syntax (Code-Point Alias)</b>                          | <pre>dscp alias-name bit-pattern;</pre>                                                                                                                                                                                                                                                                    |
| <b>Syntax (Multidestination Classifier)</b>               | <pre>dscp classifier-name;</pre>                                                                                                                                                                                                                                                                           |
| <b>Syntax (Interface Classifier Association)</b>          | <pre>dscp (classifier-name   default);</pre>                                                                                                                                                                                                                                                               |
| <b>Syntax (Rewrite Rule)</b>                              | <pre>dscp rewrite-name {   import (rewrite-name   default);   forwarding-class class-name {     loss-priority level {       code-point [ aliases ] [ bit-patterns ];     }   } }</pre>                                                                                                                     |
| <b>Hierarchy Level (Classifier)</b>                       | [edit <a href="#">class-of-service classifiers</a> ],                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level (Code-Point Aliases)</b>               | [edit <a href="#">class-of-service code-point-aliases</a> ],                                                                                                                                                                                                                                               |
| <b>Hierarchy Level (Multidestination Classifier)</b>      | [edit <a href="#">class-of-service multi-destination classifiers</a> ],                                                                                                                                                                                                                                    |
| <b>Hierarchy Level (Interface Classifier Association)</b> | [edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number classifiers</a> ],<br>[edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number rewrite-rules</a> ],                                                                                     |
| <b>Hierarchy Level (Rewrite Rule)</b>                     | [edit <a href="#">class-of-service rewrite-rules</a> ]                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                                | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                          |

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>              | Define the Differentiated Services code point (DSCP) mapping that is applied to the packets.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                  | <p><i>classifier-name</i>—Name of the classifier.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Unicast Classifiers on page 273</a></li><li>• <a href="#">Defining CoS Code-Point Aliases on page 404</a></li><li>• <a href="#">Defining CoS Rewrite Rules on page 427</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 429</a></li><li>• <a href="#">Understanding CoS Classifiers on page 51</a></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 154</a></li><li>• <a href="#">Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57</a></li></ul> |

## dscp-ipv6

|                                                     |                                                                                                                                                                                                                                        |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                               | <a href="#">Syntax (Classifier) on page 480</a><br><a href="#">Syntax (Code-Point Alias) on page 480</a><br><a href="#">Syntax (Interface Classifier Association) on page 480</a><br><a href="#">Syntax (Rewrite Rule) on page 480</a> |
| <b>Syntax (Classifier)</b>                          | <pre>dscp-ipv6 classifier-name {   import (classifier-name   default);   forwarding-class class-name {     loss-priority level {       code-points [ aliases ] [ bit-patterns ];     }   } }</pre>                                     |
| <b>Syntax (Code-Point Alias)</b>                    | <pre>dscp-ipv6 alias-name bit-pattern;</pre>                                                                                                                                                                                           |
| <b>Syntax (Interface Classifier Association)</b>    | <pre>dscp-ipv6 (classifier-name   default);</pre>                                                                                                                                                                                      |
| <b>Syntax (Rewrite Rule)</b>                        | <pre>dscp-ipv6 rewrite-name {   import (rewrite-name   default);   forwarding-class class-name {     loss-priority level {       code-point [ aliases ] [ bit-patterns ];     }   } }</pre>                                            |
| <b>Hierarchy (Classifier)</b>                       | [edit <a href="#">class-of-service classifiers</a> ],                                                                                                                                                                                  |
| <b>Hierarchy (Code-Point Alias)</b>                 | [edit <a href="#">class-of-service code-point-aliases</a> ],                                                                                                                                                                           |
| <b>Hierarchy (Interface Classifier Association)</b> | [edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number classifiers</a> ],<br>[edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number rewrite-rules</a> ],                 |
| <b>Hierarchy (Rewrite Rule)</b>                     | [edit <a href="#">class-of-service rewrite-rules</a> ]                                                                                                                                                                                 |
| <b>Release Information</b>                          | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                      |
| <b>Description</b>                                  | Define the Differentiated Services code point (DSCP) IPv6 mapping that is applied to the packets.                                                                                                                                      |



**NOTE:** There is no DSCP IPv6 classifier for multdestination (multicast, broadcast, and destination lookup fail) traffic. Multidestination IPv6 traffic uses the multdestination DSCP classifier.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Unicast Classifiers on page 273</a></li><li>• <a href="#">Defining CoS Code-Point Aliases on page 404</a></li><li>• <a href="#">Defining CoS Rewrite Rules on page 427</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 429</a></li><li>• <a href="#">Understanding CoS Classifiers on page 51</a></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 154</a></li><li>• <a href="#">Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57</a></li></ul> |


---

## dscp-code-point

---

|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>dscp-code-point <i>code-point</i>;</code>                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit class-of-service <a href="#">host-outbound-traffic</a> ]                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                  |
| <b>Description</b>              | Set the value of the DSCP code point in the type of service (ToS) field of the packet generated by the Routing Engine (host).                                                                                                                      |
| <b>Options</b>                  | <b>code-point</b> —Six-bit DSCP code point value.                                                                                                                                                                                                  |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Changing the Host Outbound Traffic Default Queue Mapping on page 415</a></li><li>• <a href="#">Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 47</a></li></ul> |

## egress (Buffer Configuration)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>egress {   percent <i>percent</i>;   <b>buffer-partition</b> (lossless   lossy   multicast) {     percent <i>percent</i>;   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service shared-buffer</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | <p>Configure the global shared buffer pool allocation for egress traffic. The system allocates the shared buffer pool dynamically across its ports as the ports require memory space. Some buffer space is reserved for other buffers such as dedicated buffers (buffers allocated permanently to ports).</p> <p>The percentage you specify is the percentage of available (user-configurable) buffer space allocated to the global shared egress buffer pool. If you allocate less than 100 percent of the available buffer space to the shared buffer pool, the remaining buffer space is added to the dedicated buffer pool. (You cannot directly configure the dedicated buffer pool for each port; dedicated buffers are allocated evenly across all the ports. However, on a port, you can configure the portion of dedicated port buffer space allocated to each queue in the scheduler configuration using the <b>buffer-size</b> option.)</p> |
|                                 | <div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p><b>CAUTION:</b> Changing the buffer configuration is a disruptive event. Traffic stops on <i>all</i> ports until buffer reprogramming is complete.</p> </div> </div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|                                 | <p>You can also partition the shared buffer pool to adjust the egress buffer allocations for different mixes of network traffic using the <b>buffer-partition</b> statement.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                  | The default shared buffer percentage is 100 percent. (All available buffer space is allocated to the shared buffer pool.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | <p><b>percent <i>percent</i></b>—Percentage of available egress buffer space allocated to the shared buffer pool. If the percentage is less than 100 percent, the remaining buffer space is allocated to the dedicated buffer pool.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 365</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |


- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 377](#)
- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 383](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 425](#)
- [Understanding CoS Buffer Configuration on page 131](#)

## enhanced-transmission-selection

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>enhanced-transmission-selection {     no-auto-negotiation;     no-recommendation-tlv;     recommendation-tlv {         no-auto-negotiation;     } }</pre>                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">dcbx</a> <a href="#">interface</a> <i>interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Disable advertising the enhanced transmission selection (ETS) state of the interface to the peer. To disable ETS on the interface, do not enable ETS on the interface in the class-of-service (CoS) configuration.</p> <p>Disabling ETS autonegotiation stops the QFX Series from advertising the ETS Configuration TLV and the ETS Recommendation TLV.</p> <p>Disabling the ETS recommendation TLV stops the QFX Series from advertising the ETS Recommendation TLV, but the ETS Configuration TLV is still advertised.</p> |
| <b>Options</b>                  | <p><b>no-auto-negotiation</b>—Disable automatic negotiation of ETS (Configuration TLV and Recommendation TLV)</p> <p><b>no-recommendation-tlv</b>—Disable automatic negotiation of the ETS Recommendation TLV</p> <p><b>recommendation-tlv</b>—Enable automatic negotiation of ETS Recommendation TLV</p>                                                                                                                                                                                                                       |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 625</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 431</a></li> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 156</a></li> </ul>                                                                                                                                                               |

## ether-type

---


|                                                                                                                                                                  |                                                                                                                                                                                                                                                                                                                                               |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                    | <code>ether-type <i>ether-type</i>;</code>                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>                                                                                                                                           | [edit applications <a href="#">application</a> <i>application-name</i> ]                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>                                                                                                                                       | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                    |
| <b>Description</b>                                                                                                                                               | Two-octet field in an Ethernet frame that defines the protocol encapsulated in the frame payload. See <a href="http://standards.ieee.org/develop/regauth/ethertype/eth.txt">http://standards.ieee.org/develop/regauth/ethertype/eth.txt</a> for a list of Institute of Electrical and Electronics Engineers (IEEE) EtherTypes.                |
| <div> <b>NOTE:</b> To create a FIP application, use the EtherType 0x8914.</div> |                                                                                                                                                                                                                                                                                                                                               |
| <b>Options</b>                                                                                                                                                   | <i>type</i> —Identifier for the EtherType.                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b>                                                                                                                                  | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                       |
| <b>Related Documentation</b>                                                                                                                                     | <ul style="list-style-type: none"><li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 434</a></li><li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 389</a></li><li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 179</a></li></ul> |



## fill-level

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>fill-level <i>low-value</i> fill-level <i>high-value</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service drop-profiles <i>profile-name</i> interpolate</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | When configuring random early detection (RED), map the fullness of a queue to a packet <a href="#">drop-probability</a> value. You configure the <b>fill-level</b> and <b>drop-probability</b> statements in related pairs by specifying a low <b>fill-level</b> value at which packets begin to drop (the drop probability is zero until the queue reaches this level of fullness) and a high <b>fill-level</b> value at which packets drop at the highest drop probability. As the queue fills from the low fill level to the high fill level, the rate of packet drop increases in a linear pattern from zero to the high drop probability. |
| <b>Options</b>                  | <p><b>low-value</b>—Fullness of the queue before packets begin to drop, expressed as a percentage. The low value must be less than the high value.</p> <p><b>Range:</b> 0 through 100</p> <p><b>high-value</b>—Fullness of the queue before it reaches the maximum drop probability. If the queue fills beyond the fill level high value, all packets drop. The high value must be greater than the low value.</p> <p><b>Range:</b> 0 through 100</p>                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring WRED Drop Profiles on page 279</a></li> <li>• <a href="#">Understanding CoS WRED Drop Profiles on page 149</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## flow-control

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | (flow-control   no-flow-control);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ether-options]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Description</b>              | <p>Explicitly enable or disable symmetric Ethernet PAUSE flow control, which regulates the flow of packets from the switch to the remote side of the connection by pausing all traffic flows on a link during periods of network congestion. Symmetric flow control means that Ethernet PAUSE is enabled in both directions. The interface generates and sends Ethernet PAUSE messages when the receive buffers fill to a certain threshold and the interface responds to PAUSE messages received from the connected peer. By default, flow control is disabled.</p> <p>You can configure asymmetric flow control by including the <b>configured-flow-control</b> statement at the [edit interfaces <i>interface-name</i> ether-options hierarchy level. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div> <ul style="list-style-type: none"> <li>• <b>flow-control</b>—Enable flow control; flow control is useful when the remote device is a Gigabit Ethernet switch.</li> <li>• <b>no-flow-control</b>—Disable flow control.</li> </ul> |
| <b>Default</b>                  | Flow control is disabled.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">configured-flow-control on page 468</a></li> <li>• <i>Configuring Gigabit and 10-Gigabit Ethernet Interfaces</i></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 160</a></li> <li>• <i>Junos OS Network Interfaces Library for Routing Devices</i></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |

## flow-control-queue (Output Congestion Notification)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>flow-control-queue [ <i>queue</i>   <i>list-of-queues</i> ];</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service congestion-notification-profile <i>profile-name</i> output ieee-802.1 code-point <i>code-point-bits</i></a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | Specify one or more output queues to pause, to support priority-based flow control (PFC). The specified queues pause when the interface receives a PFC frame with a matching IEEE 802.1p code point.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>                  | <p>Queue 3 (mapped to the fcoe forwarding class) and queue 4 (mapped to the no-loss forwarding class) are programmed as flow control queues to pause. No other output queues are programmed to pause by default.</p> <p>If you configure flow control queues explicitly, only the queues that you specify are programmed to pause. The explicit flow control queue to pause configuration overrides the default setting, so the queues paused in the default configuration are no longer paused by default.</p> <p>For example, if you configure queue 2 as a flow control queue, then queue 2 pauses when congestion occurs, but queues 3 and 4 do not pause because they were not explicitly specified. To enable pause on output queues 2, 3, and 4, you must explicitly configure all three of the queues as flow control queues.</p> <p>The same behavior applies to the IEEE 802.1p code points (priorities) on which PFC is enabled. By default, priorities 3 (011) and 4 (100) are enabled for PFC pause. If you explicitly configure flow control queues to pause, you must also explicitly configure pause for each priority (code point) that you want to pause, because the explicit configuration overrides the default configuration.</p> |
| <b>Options</b>                  | [ <i>queue</i>   <i>list-of-queues</i> ]—The output queue or a list of output queues to pause.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 417</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 317</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 309</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 340</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                               |

- [Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109](#)

## forwarding-class

|                                        |                                                                                                                                                                                                                                                             |
|----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                  | <a href="#">Syntax (Classifier) on page 489</a><br><a href="#">Syntax (Rewrite Rule) on page 489</a><br><a href="#">Syntax (Scheduler Map) on page 489</a><br><a href="#">Syntax (Interface) on page 489</a>                                                |
| <b>Syntax (Classifier)</b>             | <pre>forwarding-class class-name {   loss-priority level {     code-points [ aliases ] [ bit-patterns ];   } }</pre>                                                                                                                                        |
| <b>Syntax (Rewrite Rule)</b>           | <pre>forwarding-class class-name {   loss-priority level {     code-point [ aliases ] [ bit-patterns ];   } }</pre>                                                                                                                                         |
| <b>Syntax (Scheduler Map)</b>          | <pre>forwarding-class class-name {   scheduler scheduler-name; }</pre>                                                                                                                                                                                      |
| <b>Syntax (Interface)</b>              | <pre>forwarding-class class-name;</pre>                                                                                                                                                                                                                     |
| <b>Hierarchy Level (Classifier)</b>    | [edit <a href="#">class-of-service classifiers</a> (dscp   dscp-ipv6   ieee-802.1) <i>classifier-name</i> ],                                                                                                                                                |
| <b>Hierarchy Level (Rewrite Rule)</b>  | [edit <a href="#">class-of-service rewrite-rules</a> ] (dscp   dscp-ipv6   ieee-802.1) <i>rewrite-name</i> ],                                                                                                                                               |
| <b>Hierarchy Level (Scheduler Map)</b> | [edit <a href="#">class-of-service scheduler-maps</a> <i>map-name</i> ],                                                                                                                                                                                    |
| <b>Hierarchy Level (Interface)</b>     | [edit <a href="#">class-of-service interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ]                                                                                                                                   |
| <b>Release Information</b>             | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                           |
| <b>Description</b>                     | Configure forwarding class name and option values (classifier configuration), map rewrite rules to forwarding classes (rewrite rules), map forwarding classes to schedulers (scheduler maps), or map forwarding classes to logical interfaces (interfaces). |
| <b>Options</b>                         | <p><i>class-name</i>—Name of the forwarding class.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                            |
| <b>Required Privilege Level</b>        | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                         |
| <b>Related Documentation</b>           | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Forwarding Classes on page 283</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 289</a></li> </ul>                                                        |

- [Defining CoS Rewrite Rules on page 427](#)
- [Understanding CoS Forwarding Classes on page 68](#)
- [Understanding CoS Rewrite Rules on page 154](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)

## forwarding-class (Fibre Channel Interfaces)

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>forwarding-class <i>lossless-forwarding-class-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>     | [edit <a href="#">class-of-service interfaces</a> <i>fibre-channel-interface-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b> | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>         | <p>Configure a Layer 3 fixed classifier on a Fibre Channel (FC) interface. The fixed classifier places all traffic received from the FC network into the specified forwarding class. The forwarding class must be lossless. (That is, the forwarding class must be either the default <b>fcoe</b> or <b>no-loss</b> forwarding class, or the forwarding class must be configured with the <b>no-loss</b> drop attribute.) If you attempt to specify a lossy forwarding class, the system returns a commit error.</p> <p>FCoE networks typically use priority 3 (IEEE code point 011) for FCoE traffic. The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. If the IEEE 802.1p code point value that the Ethernet network uses for FCoE traffic is different than code point 3, you can rewrite the code point to the value used in your Ethernet (FCoE) network. The lossless forwarding class specified in the fixed classifier uses the <b>rewrite-value</b> statement as the IEEE 802.1p code point (priority) for FCoE traffic on the FCoE network.</p> <p>To rewrite the code point value, include the <b>rewrite-value input ieee code-point code-point-bits</b> statement at the [edit <a href="#">class-of-service interfaces</a> <i>fc-interface-name</i>] hierarchy level.</p> |



**NOTE:** If you are not using the default configuration (priority 3 for FCoE traffic), the lossless forwarding class specified in the FC interface fixed classifier must be mapped to the IEEE 802.1p code point specified in the rewrite value statement.

In order to avoid fate sharing (separate flows that affect each other's throughput), the code point (priority) used for the lossless forwarding class (the code point specified in the rewrite value statement) should be the only code point classified to that forwarding class (at the [edit [class-of-service classifiers](#)] hierarchy level). For example, if the rewrite value uses code point 101 for lossless FCoE forwarding class `fcoe_fc1`, then in the classifier configuration attached to ingress Ethernet interfaces, code point 101 is the only code point that should be classified to the `fcoe_fc1` forwarding class. Now if you also attach a classifier to an interface that maps code point 110 to forwarding class `fcoe_fc1`, then congestion on priority 110 unfairly (and unintentionally) affects the FCoE traffic that uses priority 101. Both priorities 101 and 110 are classified into forwarding class `fcoe_fc1`, so the traffic from both priorities shares the same fate.

**Options** *lossless-forwarding-class-name*—Name of the lossless forwarding class.

|                                 |                                                                                                                                                                                                                                                                                                                                         |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 356</a></li><li>• <a href="#">Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 128</a></li><li>• <a href="#">Understanding CoS Classifiers on page 51</a></li></ul> |

---

## forwarding-class (Host Outbound Traffic)

---

|                                 |                                                                                                                                                                                                                                                    |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | forwarding-class <i>class-name</i> ;                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit class-of-service <a href="#">host-outbound-traffic</a> ]                                                                                                                                                                                     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                  |
| <b>Description</b>              | Define forwarding class name for outbound host traffic (traffic generated by the Routing Engine).                                                                                                                                                  |
| <b>Options</b>                  | <i>class-name</i> —Name of the forwarding class.                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Changing the Host Outbound Traffic Default Queue Mapping on page 415</a></li><li>• <a href="#">Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 47</a></li></ul> |



## forwarding-class-set

---

|                                 |                                                                                                                                                                                                                                                                                                                           |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>forwarding-class-set <i>forwarding-class-set-name</i> {<br/>    output-traffic-control-profile <i>profile-name</i>;<br/>}</code>                                                                                                                                                                                    |
| <b>Hierarchy Level</b>          | [edit <code>class-of-service interfaces <i>interface-name</i></code> ]                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                         |
| <b>Description</b>              | Apply a previously defined forwarding class set to an output traffic control profile.                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <i>forwarding-class-set-name</i> —Name of the forwarding class set.<br><br>The remaining statement is explained separately.                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220</a></li> <li>• <a href="#">Assigning CoS Components to Interfaces on page 429</a></li> <li>• <a href="#">Understanding CoS Forwarding Class Sets (Priority Groups) on page 73</a></li> </ul> |

## forwarding-class-sets

---

|                                 |                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>forwarding-class-sets <i>forwarding-class-set-name</i> {<br/>    class <i>class-name</i>;<br/>}</code>                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <code>class-of-service</code> ]                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Assign forwarding classes to forwarding class sets (priority groups).                                                                                                                                                                                                                                                         |
| <b>Options</b>                  | <i>forwarding-class-set-name</i> —Name of the forwarding class set.<br><br>The remaining statement is explained separately.                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                      |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220</a></li> <li>• <a href="#">Example: Configuring Forwarding Class Sets on page 286</a></li> <li>• <a href="#">Understanding CoS Forwarding Class Sets (Priority Groups) on page 73</a></li> </ul> |

## forwarding-classes

**Syntax**

```
forwarding-classes {
 class {
 class-name {
 queue-num queue-number <no-loss>;
 }
 }
}
```

**Hierarchy Level** [edit [class-of-service](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.  
No-loss option introduced in Junos OS Release 12.3 for the QFX Series.

**Description** Map one or more forwarding classes to a single queue. You can configure up to 12 forwarding classes (8 unicast forwarding classes on queues 0 through 7 and 4 multidestination forwarding classes on queues 8 through 11) and map them to queues. You can map multiple forwarding classes to a single queue using the **class** statement. All forwarding classes mapped to a particular queue must be of the same type, either unicast or multicast. You cannot mix unicast and multicast forwarding classes on the same queue.

You cannot configure weighted random early detection (WRED) packet drop on forwarding classes configured with the no-loss packet drop attribute. Do not associate a drop profile with lossless forwarding classes.



**NOTE:** If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).



**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.



If you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the fcoe and no-loss forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the no-loss option. If you do not specify the no-loss option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if

you explicitly configure the `fcoe` forwarding class and you do not include the `no-loss` option, the `fcoe` forwarding class is lossy, not lossless.

.....

|                              |                                                                                                                                                                                                 |
|------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>               | The statements are explained separately.                                                                                                                                                        |
| <b>Required Privilege</b>    | interfaces—To view this statement in the configuration.                                                                                                                                         |
| <b>Level</b>                 | interface-control—To add this statement to the configuration.                                                                                                                                   |
| <b>Related Documentation</b> | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Forwarding Classes on page 283</a></li><li>• <a href="#">Understanding CoS Forwarding Classes on page 68</a></li></ul> |

## guaranteed-rate

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>guaranteed-rate (rate  percent <i>percentage</i>);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service traffic-control-profiles</a> <i>traffic-control-profile-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Configure a guaranteed minimum rate of transmission for a traffic control profile. The sum of the guaranteed rates of all of the forwarding class sets (priority groups) on a port should not exceed the total port bandwidth. The guaranteed rate also determines the amount of excess (extra) port bandwidth that the priority group (forwarding class set) can share. Extra port bandwidth is allocated among the priority groups on a port in proportion to the guaranteed rate of each priority group.                                                                                      |
|                                 | <p> <b>NOTE:</b> You cannot configure a guaranteed rate for a forwarding class set (priority group) that includes strict-high priority queues. If the traffic control profile is for a forwarding class set that contains strict-high priority queues, do not configure a guaranteed rate.</p>                                                                                                                                                                                                                  |
| <b>Default</b>                  | If you do not specify a guaranteed rate, the guaranteed rate is zero (0) and there is no minimum guaranteed bandwidth.                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|                                 | <p> <b>NOTE:</b> If you do not configure a guaranteed rate for a traffic control profile, the queues that belong to any forwarding class set (priority group) that uses that traffic control profile cannot have a configured transmit rate. The result is that there is no minimum guaranteed bandwidth for those queues and that those queues can be starved during periods of congestion.</p>                                                                                                              |
| <b>Options</b>                  | <p><b>percent <i>percentage</i></b>—Minimum percentage of transmission capacity allocated to the forwarding class set or logical interface.</p> <p><b>Range:</b> 1 through 100 percent</p> <p><b><i>rate</i></b>—Minimum transmission rate allocated to the forwarding class set or logical interface, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation <b>k</b> (1000), <b>m</b> (1,000,000), or <b>g</b> (1,000,000,000).</p> <p><b>Range:</b> 1000 through 10,000,000,000 bps</p> |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                              |

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
  - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
  - [Example: Configuring Minimum Guaranteed Output Bandwidth on page 300](#)
  - [Understanding CoS Traffic Control Profiles on page 98](#)
  - [output-traffic-control-profile on page 516](#)

## host-outbound-traffic

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>host-outbound-traffic {     forwarding-class <i>class-name</i>;     dscp-code-point <i>code-point</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit class-of-service]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | <p>Allow queue selection for traffic generated by the Routing Engine (host). The selected queue must be configured properly. You can also configure specific DSCP code point bits for the type of service (ToS) field of the generated packets. This configuration does not affect transit packets or incoming packets. This is a global configuration that only affects packets originating on the Routing Engine. If you do not configure an output queue for host outbound traffic, the switch uses the default queue mapping.</p> |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Changing the Host Outbound Traffic Default Queue Mapping on page 415</a></li> <li>• <a href="#">Understanding Host Routing Engine Outbound Traffic Queues and Defaults on page 47</a></li> </ul>                                                                                                                                                                                                                                                                                 |

## ieee-802.1

|                                                           |                                                                                                                                                                                                                                                                                                            |
|-----------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                     | <a href="#">Syntax (Classifier) on page 498</a><br><a href="#">Syntax (Code-Point Alias) on page 498</a><br><a href="#">Syntax (Multidestination Classifier) on page 498</a><br><a href="#">Syntax (Interface Classifier Association) on page 498</a><br><a href="#">Syntax (Rewrite Rule) on page 498</a> |
| <b>Syntax (Classifier)</b>                                | <pre> ieee-802.1 classifier-name {   import (classifier-name   default);   forwarding-class class-name {     loss-priority level {       code-points [ aliases ] [ bit-patterns ];     }   } }</pre>                                                                                                       |
| <b>Syntax (Code-Point Alias)</b>                          | <pre> ieee-802.1 alias-name bit-pattern;</pre>                                                                                                                                                                                                                                                             |
| <b>Syntax (Multidestination Classifier)</b>               | <pre> ieee-802.1 classifier-name;</pre>                                                                                                                                                                                                                                                                    |
| <b>Syntax (Interface Classifier Association)</b>          | <pre> ieee-802.1 (classifier-name   default);</pre>                                                                                                                                                                                                                                                        |
| <b>Syntax (Rewrite Rule)</b>                              | <pre> ieee-802.1 rewrite-name {   import (rewrite-name   default);   forwarding-class class-name {     loss-priority level {       code-point [ aliases ] [ bit-patterns ];     }   } }</pre>                                                                                                              |
| <b>Hierarchy Level (Classifier)</b>                       | [edit <a href="#">class-of-service classifiers</a> ],                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level (Code-Point Alias)</b>                 | [edit <a href="#">class-of-service code-point-aliases</a> ],                                                                                                                                                                                                                                               |
| <b>Hierarchy Level (Multidestination Classifier)</b>      | [edit <a href="#">class-of-service multi-destination classifiers</a> ],                                                                                                                                                                                                                                    |
| <b>Hierarchy Level (Interface Classifier Association)</b> | [edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number classifiers</a> ],<br>[edit <a href="#">class-of-service interfaces interface-name unit logical-unit-number rewrite-rules</a> ],                                                                                     |
| <b>Hierarchy Level (Rewrite Rule)</b>                     | [edit <a href="#">class-of-service rewrite-rules</a> ]                                                                                                                                                                                                                                                     |
| <b>Release Information</b>                                | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                          |

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Description</b>              | Configure an IEEE 802.1 classifier, configure an IEEE 802.1 code-point alias, apply a fixed IEEE 802.1 classifier to an interface, or apply an IEEE-802.1 rewrite rule.                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><i>classifier-name</i>—Name of the classifier.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Unicast Classifiers on page 273</a></li><li>• <a href="#">Defining CoS Code-Point Aliases on page 404</a></li><li>• <a href="#">Defining CoS Rewrite Rules on page 427</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 429</a></li><li>• <a href="#">Understanding CoS Classifiers on page 51</a></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 154</a></li></ul> |

## ieee-802.1 (Fibre Channel Interfaces)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ieee-802.1 {   code-point code-point-bits; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service interfaces</a> <i>fibre-channel-interface-name</i> <a href="#">rewrite-value</a> <a href="#">input</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Description</b>              | <p>Configure the IEEE 802.1p code point value to which all traffic received from the Fibre Channel (FC) network on the specified FC interface is rewritten. After the code point value is rewritten, the interface forwards the traffic to the Ethernet (FCoE) network. This works in conjunction with configuring a fixed classifier on the FC interface. The fixed classifier maps all traffic from the FC network into one lossless forwarding class (the lossless forwarding class must be mapped to the code point specified in the rewrite value). Traffic mapped to the lossless forwarding class uses the IEEE 802.1p priority specified by the code point bits in the rewrite value.</p> <p>FCoE networks typically use priority 3 (IEEE code point 011) for FCoE traffic. The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. Rewriting the code point value enables you to change the IEEE 802.1p priority of the FCoE traffic if the Ethernet network uses a different priority than priority 3 (code point 011).</p> <p>The system supports only one IEEE 802.1p code point value per FC interface. You cannot configure more than one IEEE 802.1p rewrite value per FC interface. In addition, you can specify only one rewrite value per local FCoE-FC gateway fabric; all interfaces in the local fabric must use the same rewrite value. Attempting to configure FC interfaces in the same local fabric with different rewrite values generates a commit error. You can specify different rewrite values for interfaces that belong to different local FCoE-FC gateway fabrics.</p> <p>The statement is described separately.</p> |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">forwarding-class (Fibre Channel Interfaces) on page 491</a></li> <li>• <a href="#">Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 356</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 128</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 51</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |



## ieee-802.1 (Input Congestion Notification)

|                                 |                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> ieee-802.1 {   code-point [code-point-bits] {     pfc {       mru mru-value;     }   } } </pre>                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service congestion-notification-profile</a> <i>profile-name</i> <a href="#">input</a> ]                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure an IEEE 802.1 code point and apply priority-based flow control (PFC) to packets with that code point.                                                                                                                                                                                                            |
| <b>Options</b>                  | The statements are described separately.                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 241</a></li> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 417</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 160</a></li> </ul> |

## ieee-802.1 (Output Congestion Notification)

---


|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ieee-802.1 {<br/>    code-point [ code-point-bits ] {<br/>        flow-control-queue [ queue   list-of-queues ];<br/>    }<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service congestion-notification-profile</b> <i>profile-name</i> <b>output</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Description</b>              | Configure an IEEE 802.1 code point and apply priority-based flow control (PFC) to packets with that code point on output queues.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                  | The statements are described separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 417</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326</a></li><li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 317</a></li><li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 309</a></li><li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 340</a></li><li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109</a></li></ul> |

## import

---

|                                 |                                                                                                                                                                                                                                                                                                                                       |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>import (<i>import</i>   default);</code>                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service classifiers</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a> ) <i>classifier-name</i> ],<br>[edit <a href="#">class-of-service rewrite-rules</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a> ) <i>classifier-name</i> ]     |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                     |
| <b>Description</b>              | Specify a default or previously defined classifier.                                                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <p><b><i>import</i></b>—Name of the classifier mapping configured at the [edit <a href="#">class-of-service classifiers</a>] hierarchy level.</p> <p><b>default</b>—Default classifier mapping.</p> <p>The remaining statements are explained separately.</p>                                                                         |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Unicast Classifiers on page 273</a></li> <li>• <a href="#">Defining CoS Rewrite Rules on page 427</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 51</a></li> <li>• <a href="#">Understanding CoS Rewrite Rules on page 154</a></li> </ul> |

## ingress (Buffer Configuration)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>ingress {   <b>buffer-partition</b> (lossless   lossless-headroom   lossy) {     percent <i>percent</i>;   }   percent <i>percent</i>; }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service shared-buffer</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | <p>Configure the global shared buffer pool allocation for ingress traffic. The system allocates the shared buffer pool dynamically across its ports as the ports require memory space. Some buffer space is reserved for buffers such as dedicated buffers (buffers allocated permanently to ports) and headroom buffers (buffers that help prevent packet loss on lossless flows).</p> <p>The percentage you specify is the percentage of available (user-configurable) buffer space allocated to the global shared ingress buffer pool. If you allocate less than 100 percent of the available buffer space to the shared buffer pool, the remaining buffer space is added to the dedicated buffer pool. (You cannot directly configure the dedicated buffer pool for each port; dedicated buffers are allocated evenly across all the ports.)</p> |
|                                 | <div>  <p><b>CAUTION:</b> Changing the buffer configuration is a disruptive event. Traffic stops on <i>all</i> ports until buffer reprogramming is complete.</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|                                 | <p>You can also partition the shared buffer pool to adjust the ingress buffer allocations for different mixes of network traffic using the <b>buffer-partition</b> statement.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Default</b>                  | The default shared buffer percentage is 100 percent. (All available buffer space is allocated to the shared buffer pool.)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Options</b>                  | <p><b>percent <i>percent</i></b>—Percentage of available ingress buffer space allocated to the shared buffer pool. If the percentage is less than 100 percent, the remaining buffer space is allocated to the dedicated buffer pool.</p> <p>The remaining statement is explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 365</a></li> <li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 377</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

- [Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 383](#)
- [Configuring Global Ingress and Egress Shared Buffers on page 425](#)
- [Understanding CoS Buffer Configuration on page 131](#)

## input (Congestion Notification)

|                                 |                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> input {   ieee-802.1 {     code-point [code-point-bits] {       pfc {         mru mru-value;       }     }   }   cable-length cable-length-value; } </pre>                                                                                                                                                           |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service congestion-notification-profile</b> <i>profile-name</i> ]                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure priority-based flow control (PFC) on incoming traffic.                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 241</a></li> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 417</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 160</a></li> </ul> |

## input (Fibre Channel Interfaces)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>input {   ieee-802.1p {     code-point code-point-bits;   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service interfaces</a> <i>fibre-channel-interface-name</i> <a href="#">rewrite-value</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | <p>Configure the IEEE 802.1p code point value to which all traffic received from the Fibre Channel (FC) network on the specified FC interface is rewritten. After the code point value is rewritten, the interface forwards the traffic to the Ethernet (FCoE) network. This works in conjunction with configuring a fixed classifier on the FC interface. The fixed classifier maps all traffic from the FC network into one lossless forwarding class (the lossless forwarding class must be mapped to the code point specified in the rewrite value). Traffic mapped to the lossless forwarding class uses the IEEE 802.1p priority specified by the code point bits in the rewrite value.</p> <p>FCoE networks typically use priority 3 (IEEE code point 011) for FCoE traffic. The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. Rewriting the code point value enables you to change the IEEE 802.1p priority of the FCoE traffic if the Ethernet network uses a different priority than priority 3 (code point 011).</p> <p>The system supports only one IEEE 802.1p code point value per FC interface. You cannot configure more than one IEEE 802.1p rewrite value per FC interface. In addition, you can specify only one rewrite value per local FCoE-FC gateway fabric; all interfaces in the local fabric must use the same rewrite value. Attempting to configure FC interfaces in the same local fabric with different rewrite values generates a commit error. You can specify different rewrite values for interfaces that belong to different local FCoE-FC gateway fabrics.</p> <p>The statements are described separately.</p> |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">forwarding-class (Fibre Channel Interfaces) on page 491</a></li> <li>• <a href="#">Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 356</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 128</a></li> <li>• <a href="#">Understanding CoS Classifiers on page 51</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |

## interface (DCBX)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> interface (<i>interface-name</i>   all) {   disable;   application-map <i>application-map-name</i>;   applications {     no-auto-negotiation;   }   enhanced-transmission-selection {     no-auto-negotiation;     no-recommendation-tlv;     recommendation-tlv {       no-auto-negotiation;     }   }   dcbx-version (auto-negotiate   ieee-dcbx   dcbx-version-1.01);   priority-flow-control {     no-auto-negotiation;   } } </pre>                                                                                                         |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">dcbx</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | <p>Statement introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Statement introduced in Junos OS Release 11.3 for the EX Series switches.</p> <p><b>Mode</b> and <b>recommendation-tlv</b> statements introduced in Junos OS Release 12.2 for the QFX Series.</p>                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Configure DCBX properties on an interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><b><i>interface-name</i></b>—Name of the interface.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | <p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">show dcbx neighbors on page 625</a></li> <li>• <a href="#">Configuring DCBX Autonegotiation on page 431</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on page 156</a></li> <li>• <a href="#">Understanding DCB Features and Requirements on EX Series Switches</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## interfaces (Class of Service)

```
Syntax interfaces {
 interface-name {
 congestion-notification-profile profile-name {
 }
 forwarding-class lossless-forwarding-class-name;
 forwarding-class-set forwarding-class-set-name {
 output-traffic-control-profile profile-name;
 }
 rewrite-value {
 input {
 ieee-802.1 {
 code-point code-point-bits;
 }
 }
 }
 }
 unit logical-unit-number {
 classifiers {
 (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);
 }
 forwarding-class class-name;
 rewrite-rules {
 (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);
 }
 }
}
```

**Hierarchy Level** [edit [class-of-service](#)]

**Release Information** Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description** Configure interface-specific CoS properties for incoming packets.

**Options** *interface-name*—Name of the interface.

The statements are explained separately.

**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring Unicast Classifiers on page 273](#)
- [Example: Configuring Forwarding Classes on page 283](#)
- [Example: Configuring Forwarding Class Sets on page 286](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
- [Assigning CoS Components to Interfaces on page 429](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
- [Defining CoS Rewrite Rules on page 427](#)



- *Interfaces Overview*

## interpolate

---

|                                 |                                                                                                                                                                                                     |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> interpolate {   fill-level low-value fill-level high-value;   drop-probability 0 drop-probability high-value; } </pre>                                                                        |
| <b>Hierarchy Level</b>          | [edit class-of-service <b>drop-profiles</b> <i>profile-name</i> ]                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                   |
| <b>Description</b>              | <p>Specify values for interpolating the relationship between queue fill level and drop probability.</p> <p>The statements are explained separately.</p>                                             |
| <b>Required Privilege Level</b> | <p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring WRED Drop Profiles on page 279</a></li> <li>• <a href="#">Understanding CoS WRED Drop Profiles on page 149</a></li> </ul> |

## loss-priority (Classifiers)

---

|                                 |                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>loss-priority level {<br/>    code-points [ aliases ] [ bit-patterns ];<br/>}</code>                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service classifiers</a> ( <a href="#">dscp</a>   <a href="#">dscp-ipv6</a>   <a href="#">ieee-802.1</a> ) <i>classifier-name</i><br><a href="#">forwarding-class</a> <i>class-name</i> ]                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Configure packet loss priority value for a specific set of code-point aliases and bit patterns.                                                                                                                                                                                                                             |
| <b>Options</b>                  | <p><i>level</i>—Can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>low</b>—Packet has low loss priority.</li><li>• <b>medium-high</b>—Packet has medium-high loss priority.</li><li>• <b>high</b>—Packet has high loss priority.</li></ul> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                         |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Unicast Classifiers on page 273</a></li><li>• <a href="#">Understanding CoS Classifiers on page 51</a></li></ul>                                                                                                                                   |

## loss-priority (Drop Profiles)

|                                 |                                                                                                                                                                                                                                                                                                                                          |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>loss-priority <i>level</i> <i>protocol</i> <i>protocol</i> <b>drop-profile</b> <i>profile-name</i>;</code>                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service</b> <b>schedulers</b> <i>scheduler-name</i> <b>drop-profile-map</b> ]                                                                                                                                                                                                                                          |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                        |
| <b>Description</b>              | Configure packet loss priority value for a drop profile mapped to a system drop profile.                                                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><b><i>level</i></b>—Can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>low</b>—Packet has low loss priority.</li> <li>• <b>medium-high</b>—Packet has medium-high loss priority.</li> <li>• <b>high</b>—Packet has high loss priority.</li> </ul> <p>The remaining statements are explained separately.</p> |
| <b>Required Privilege Level</b> | <p><b>interfaces</b>—To view this statement in the configuration.</p> <p><b>interface-control</b>—To add this statement to the configuration.</p>                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring Drop Profile Maps on page 281</a></li> <li>• <a href="#">Example: Configuring WRED Drop Profiles on page 279</a></li> <li>• <a href="#">Understanding CoS WRED Drop Profiles on page 149</a></li> </ul>                                                        |

## loss-priority (Rewrite Rules)

---

|                                 |                                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>loss-priority <i>level</i> {<br/>    <code>code-point</code> (<i>alias</i>   <i>bit-pattern</i>);<br/>}</code>                                                                                                                                                                                                        |
| <b>Hierarchy Level</b>          | [edit <code>class-of-service rewrite-rules</code> ( <code>dscp</code>   <code>dscp-ipv6</code>   <code>ieee-802.1</code> ) <i>rewrite-name</i><br><code>forwarding-class</code> <i>class-name</i> ]                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                           |
| <b>Description</b>              | Specify a loss priority to which to apply a rewrite rule. The rewrite rule sets the code-point aliases and bit patterns for a specific forwarding class and loss priority. Packets that match the forwarding class and loss priority are rewritten with the rewrite code-point alias or bit pattern.                        |
| <b>Options</b>                  | <p><i>level</i>—Can be one of the following:</p> <ul style="list-style-type: none"><li>• <b>low</b>—Packet has low loss priority.</li><li>• <b>medium-high</b>—Packet has medium-high loss priority.</li><li>• <b>high</b>—Packet has high loss priority.</li></ul> <p>The remaining statement is explained separately.</p> |
| <b>Required Privilege Level</b> | <p><code>interfaces</code>—To view this statement in the configuration.</p> <p><code>interface-control</code>—To add this statement to the configuration.</p>                                                                                                                                                               |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Defining CoS Rewrite Rules on page 427</a></li><li>• <a href="#">Understanding CoS Rewrite Rules on page 154</a></li></ul>                                                                                                                                              |


---

## multi-destination

---

|                                 |                                                                                                                                                                                                                                                                                                             |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>multi-destination {<br/>  classifiers {<br/>    (dscp   ieee-802.1) classifier-name;<br/>  }<br/>}</pre>                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                                                                                                    |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                           |
| <b>Description</b>              | Define a multicast CoS behavior aggregate (BA) classifier.                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | The statements are explained separately.                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                    |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Multidestination (Multicast, Broadcast, DLF) Classifiers on page 277</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 429</a></li><li>• <a href="#">Understanding CoS Classifiers on page 51</a></li></ul> |

## mrp

|                                                                                                                                                                                                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                        | <code>mrp <i>mrp-value</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                               | [edit <a href="#">class-of-service congestion-notification-profile <i>profile-name</i> input ieee-802.1 code-point <i>code-point-bits</i> pfc</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                           | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                   | Configure the maximum receive unit (MRU) of the interface in bytes (incoming packet sizes must be less than or equal to the MRU, or the packets are dropped). The system uses the MRU and the cable length to calculate the amount of buffer headroom reserved to support priority-based flow control (PFC). The lower the MRU and the shorter the cable length, the less headroom buffer space is required for PFC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <div>  <p><b>NOTE:</b> You can also set a maximum transmission unit (MTU) value (the largest packet size the interface sends) for interfaces by including the <code>mtu</code> statement at the [edit <a href="#">interfaces <i>interface-name</i></a>] hierarchy level.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                       | For priority 3 traffic, the default MRU value is 2500 bytes.<br><br>For priority 4 traffic, the default MRU value is 9612 bytes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                       | <b><i>mrp-value</i></b> —Value of the maximum packet receive unit size in bytes (generally from 1500 to 9216 bytes, but there is no configuration restriction).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                      | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                         | <ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 417</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 317</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 309</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 340</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 160</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109</a></li> </ul> |

## output (Congestion Notification)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>output {   ieee-802.1 {     code-point [code-point-bits] {       flow-control-queue [queue   list-of-queues];     }   } }</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service congestion-notification-profile</a> <i>profile-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>              | Configure priority-based flow control (PFC) on output queues.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Options</b>                  | The remaining statements are explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 417</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 317</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 309</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 340</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109</a></li> </ul> |

## output-traffic-control-profile

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>output-traffic-control-profile <i>profile-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service interfaces</a> <i>interface-name</i> <a href="#">forwarding-class-set</a> <i>forwarding-class-set-name</i> ]                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                      |
| <b>Description</b>              | Apply an output traffic scheduling and shaping profile to a forwarding class set (priority group).                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <i>profile-name</i> —Name of the traffic-control profile to apply to the specified forwarding class set.                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <code>interfaces</code> —To view this statement in the configuration.<br><code>interface-control</code> —To add this statement to the configuration.                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220</a></li><li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 298</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 429</a></li><li>• <a href="#">Understanding CoS Traffic Control Profiles on page 98</a></li></ul> |



## pfc (Input Congestion Notification)

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>pfc {<br/>    <b>mru</b> <i>mru-value</i>;<br/>}</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service</b> <b>congestion-notification-profile</b> <i>profile-name</i> <b>input</b> <b>ieee-802.1</b> <b>code-point</b> <i>code-point-bits</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>              | Enable and configure ingress interface priority-based flow control (PFC).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Options</b>                  | The remaining statement is explained separately.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 417</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE IEEE 802.1p Priorities on Different FCoE Transit Switch Interfaces on page 326</a></li> <li>• <a href="#">Example: Configuring Two or More Lossless FCoE Priorities on the Same FCoE Transit Switch Interface on page 317</a></li> <li>• <a href="#">Example: Configuring Lossless FCoE Traffic When the Converged Ethernet Network Does Not Use IEEE 802.1p Priority 3 for FCoE Traffic (FCoE Transit Switch) on page 309</a></li> <li>• <a href="#">Example: Configuring Lossless IEEE 802.1p Priorities on Ethernet Interfaces for Multiple Applications (FCoE and iSCSI) on page 340</a></li> <li>• <a href="#">Understanding CoS Flow Control (Ethernet PAUSE and PFC) on page 160</a></li> <li>• <a href="#">Understanding CoS IEEE 802.1p Priorities for Lossless Traffic Flows on page 109</a></li> </ul> |

## policy-options

```
Syntax policy-options
 application-maps application-map-name {
 application application-name {
 code-points [aliases] [bit-patterns];
 }
 }
 policy-statement policy-name {
 term term-name {
 from {
 family family-name;
 match-conditions;
 policy subroutine-policy-name;
 prefix-list prefix-list-name;
 prefix-list-filter prefix-list-name match-type <actions>;
 route-filter destination-prefix match-type <actions>;
 source-address-filter source-prefix match-type <actions>;
 }
 to {
 match-conditions;
 policy subroutine-policy-name;
 }
 then actions;
 }
 }
 }
```

**Hierarchy Level** [edit]

**Release Information** Statement introduced in Junos OS Release 12.1 for the QFX Series.  
Statement introduced in Junos OS Release 12.1 for the EX Series.


**Description** Configure options such as application maps for DCBX application protocol exchange and policy statements.

**Required Privilege Level** storage—To view this statement in the configuration.  
storage-control—To add this statement to the configuration.

**Related Documentation**

- [Defining an Application for DCBX Application Protocol TLV Exchange on page 434](#)
- [Example: Configuring DCBX Application Protocol TLV Exchange on page 389](#)
- [Example: Configuring DCBX to Support an iSCSI Application](#)
- [Understanding DCBX Application Protocol TLV Exchange on page 179](#)
- [Understanding DCBX Application Protocol TLV Exchange on EX Series Switches](#)

## priority (Schedulers)


|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <code>priority <i>priority</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                         | [edit <a href="#">class-of-service schedulers</a> <i>scheduler-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                     | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Specify the packet-scheduling drop priority value.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <p><b><i>priority</i></b>—It can be one of the following:</p> <ul style="list-style-type: none"> <li>• <b>low</b>—Scheduler has low priority.</li> <li>• <b>strict-high</b>—Scheduler has strict high priority. You can configure only one queue as a strict-high priority queue. Strict-high priority allocates the scheduled bandwidth to the queue before any other queue receives bandwidth. Other queues receive the bandwidth that remains after the strict-high queue has been serviced.</li> </ul> |
| <div>  <p><b>NOTE:</b> We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 289</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 86</a></li> </ul>                                                                                                                                                                                                   |

## priority-flow-control

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>priority-flow-control {<br/>    no-auto-negotiation;<br/>}</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Hierarchy Level</b>          | [edit protocols <a href="#">dcbx interface</a> (all   <i>interface-name</i> )]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>Statement introduced in Junos OS Release 11.3 for EX Series switches.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Description</b>              | Disable autonegotiation of priority-based flow control (PFC) on one or more Ethernet interfaces. Autonegotiation enables PFC on an interface only if the switch and the peer device connected to the switch both support PFC and have the same PFC configuration. Disabling autonegotiation on an interface forces the interface to use the PFC state (enabled or disabled) that is configured on the switch by the configuration and assignment of the congestion notification profile.                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b>no-auto-negotiation</b> —Disable automatic negotiation of PFC.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 625</a></li><li>• <a href="#">Configuring CoS PFC (Congestion Notification Profiles) on page 417</a></li><li>• <a href="#">Configuring Priority-Based Flow Control for an EX Series Switch (CLI Procedure)</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 431</a></li><li>• <a href="#">Example: Configuring CoS PFC for FCoE Traffic on page 241</a></li><li>• <a href="#">Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches</a></li><li>• <a href="#">Understanding Priority-Based Flow Control</a></li><li>• <a href="#">Understanding DCB Features and Requirements on page 156</a></li></ul> |

## protocol (Applications)

|                                                                                                                                                                                                                           |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                             | <code>protocol (tcp   udp);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Hierarchy Level</b>                                                                                                                                                                                                    | [edit applications <b>application</b> <i>application-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>                                                                                                                                                                                                | Statement introduced in Junos OS Release 12.1 for EX Series switches.<br>Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Description</b>                                                                                                                                                                                                        | Networking protocol type, which combines with <b>destination-port</b> to identify an application type.                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <div>  <b>NOTE:</b> To create an application for iSCSI, use the protocol <b>tcp</b> with the destination port number <b>3260</b>. </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Options</b>                                                                                                                                                                                                            | <b>tcp</b> —Transmission Control Protocol<br><br><b>udp</b> —User Datagram Protocol                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Required Privilege Level</b>                                                                                                                                                                                           | <b>interface</b> —To view this statement in the configuration.<br><b>interface-control</b> —To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Related Documentation</b>                                                                                                                                                                                              | <ul style="list-style-type: none"> <li>• <a href="#">Defining an Application for DCBX Application Protocol TLV Exchange on page 434</a></li> <li>• <a href="#">Example: Configuring DCBX Application Protocol TLV Exchange on page 389</a></li> <li>• <a href="#">Example: Configuring DCBX to Support an iSCSI Application</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on page 179</a></li> <li>• <a href="#">Understanding DCBX Application Protocol TLV Exchange on EX Series Switches</a></li> </ul> |

## protocol (Drop Profile Map)

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>protocol <i>protocol</i> <b>drop-profile</b> <i>profile-name</i>;</code>                                                                                                                                                                                                                                                                                                  |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service</b> <b>schedulers</b> <i>scheduler-name</i> <b>drop-profile-map</b> <b>loss-priority</b> (low   medium-high   high)]                                                                                                                                                                                                                                  |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                               |
| <b>Description</b>              | Configure the protocol type for the specified drop profile.                                                                                                                                                                                                                                                                                                                     |
| <b>Options</b>                  | <p><b><i>protocol</i></b>—Type of protocol. The protocol can be:</p> <ul style="list-style-type: none"><li>• <b>any</b>—Accept any protocol type.</li></ul> <p>The remaining statement is explained separately.</p>                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                             |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220</a></li><li>• <a href="#">Example: Configuring Drop Profile Maps on page 281</a></li><li>• <a href="#">Example: Configuring WRED Drop Profiles on page 279</a></li><li>• <a href="#">Understanding CoS WRED Drop Profiles on page 149</a></li></ul> |

## queue-num

|                            |                                                                                                                                             |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>queue-num <i>queue-number</i> &lt;no-loss&gt;;</code>                                                                                 |
| <b>Hierarchy Level</b>     | [edit <a href="#">class-of-service forwarding-classes class</a> <i>class-name</i> ]                                                         |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>No-loss option introduced in Junos OS Release 12.3 for the QFX Series. |
| <b>Description</b>         | Map a forwarding class to an output queue number. Optionally, configure the forwarding class as a lossless forwarding class.                |

You can map some or all of the eight unicast forwarding classes to a unicast queue (0 through 7) or some or all of the four multdestination (multicast, broadcast, destination lookup fail) forwarding classes to the same multdestination queue (8 through 11), providing that you do not map one forwarding class to more than one queue. The queue to which you map a forwarding class determines if the forwarding class is a unicast or multdestination forwarding class.

You cannot configure weighted random early detection (WRED) packet drop on forwarding classes configured with the no-loss packet drop attribute. Do not associate a drop profile with lossless forwarding classes.



**NOTE:** If you map more than one forwarding class to a queue, all of the forwarding classes mapped to the queue must have the same packet drop attribute (all of the forwarding classes must be lossy, or all of the forwarding classes mapped to a queue must be lossless).



**NOTE:** If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless fcoe and no-loss forwarding classes. If you explicitly configure lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

If you are using Junos OS Release 12.3 or later, the default configuration is the same as the default configuration for Junos OS Release 12.2, and the default behavior is the same (the fcoe and no-loss forwarding classes receive lossless treatment). However, if you explicitly configure lossless forwarding classes, you can configure up to six lossless forwarding classes by specifying the no-loss option. If you do not specify the no-loss option in an explicit forwarding class configuration, the forwarding class is lossy. For example, if you explicitly configure the fcoe forwarding class and you do not include the no-loss option, the fcoe forwarding class is lossy, not lossless.

|                                 |                                                                                                                                                                                                                                                 |
|---------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Options</b>                  | <b><i>queue-number</i></b> —Number of the CoS unicast queue (0 through 7) or the CoS multidestination queue (8 through 11).<br><br><b><i>no-loss</i></b> —Optional packet drop attribute keyword to configure the forwarding class as lossless. |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring Forwarding Classes on page 283</a></li><li>• <a href="#">Understanding CoS Forwarding Classes on page 68</a></li></ul>                                                 |

---

## recommendation-tlv

---

|                                 |                                                                                                                                                                                                                                                                                                                                  |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>recommendation-tlv {<br/>    no-auto-negotiation;<br/>}</pre>                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level</b>          | [edit protocols <b>dcbx</b> <b>interface</b> <i>interface-name</i> <b>enhanced-transmission-selection</b> ]                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.2 for the QFX Series.                                                                                                                                                                                                                                                                |
| <b>Description</b>              | Enable DCBX to send the ETS Recommendation TLV (also known as the Information TLV) on egress. This feature is valid only if the interface DCBX mode is IEEE DCBX. If the interface DCBX mode is DCBX version 1.01, this statement has no effect. (DCBX version 1.01 does not advertise separate TLVs for individual attributes.) |
| <b>Default</b>                  | DCBX-enabled interfaces send the ETS recommendation TLV unless it is disabled.                                                                                                                                                                                                                                                   |
| <b>Options</b>                  | <b><i>no-auto-negotiation</i></b> —Disable sending of the ETS recommendation TLV.                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | routing—To view this statement in the configuration.<br>routing-control—To add this statement to the configuration.                                                                                                                                                                                                              |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">show dcbx neighbors on page 625</a></li><li>• <a href="#">Configuring DCBX Autonegotiation on page 431</a></li></ul>                                                                                                                                                         |



## rewrite-rules

|                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                  |
|------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>List of Syntax</b>                                            | <a href="#">Syntax (Rewrite Rule Configuration) on page 525</a><br><a href="#">Syntax (Rewrite Rule Association with Interface) on page 525</a>                                                                                                                                                                                                                                                  |
| <b>Syntax (Rewrite Rule Configuration)</b>                       | <pre>rewrite-rules {   (dscp   dscp-ipv6   ieee-802.1   exp) rewrite-name {     import (rewrite-name   default);     forwarding-class class-name {       loss-priority priority code-point (alias   bits);     }   } }</pre>                                                                                                                                                                     |
| <b>Syntax (Rewrite Rule Association with Interface)</b>          | <pre>rewrite-rules {   (dscp   dscp-ipv6   ieee-802.1   exp) rewrite-name; }</pre>                                                                                                                                                                                                                                                                                                               |
| <b>Hierarchy Level (Rewrite Rule Configuration)</b>              | [edit <a href="#">class-of-service</a> ],                                                                                                                                                                                                                                                                                                                                                        |
| <b>Hierarchy Level (Rewrite Rule Association with Interface)</b> | [edit <a href="#">class-of-service interfaces</a> <i>interface-name</i> <a href="#">unit</a> <i>logical-unit-number</i> ]                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                                       | Statement introduced in Junos OS Release 11.1 for the QFX Series.<br>EXP statement introduced in Junos OS Release 12.3 for the QFX Series.                                                                                                                                                                                                                                                       |
| <b>Description</b>                                               | <p>Configure rewrite rules that map traffic to code points when traffic exits the system, and apply the rewrite rules to a specific interface.</p> <p>MPLS EXP rewrite rules can only be bound to logical interfaces, not to physical interfaces. You can configure as many EXP rewrite rules as you want, but you can use only 16 EXP rewrite rules on switch interfaces at any given time.</p> |
| <b>Options</b>                                                   | The statements are explained separately.                                                                                                                                                                                                                                                                                                                                                         |
| <b>Required Privilege Level</b>                                  | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>                                     | <ul style="list-style-type: none"> <li>• <a href="#">Defining CoS Rewrite Rules on page 427</a></li> <li>• <a href="#">Configuring Rewrite Rules for MPLS EXP Classifiers</a></li> <li>• <a href="#">Understanding CoS Rewrite Rules on page 154</a></li> <li>• <a href="#">Understanding CoS MPLS EXP Classifiers and Rewrite Rules</a></li> </ul>                                              |

## rewrite-value (Fibre Channel Interfaces)

```
Syntax rewrite-value {
 input {
 ieee-802.1p {
 code-point code-point-bits;
 }
 }
 }
```

**Hierarchy Level** [edit [class-of-service interfaces](#) *fibre-channel-interface-name*]

**Description** Configure the IEEE 802.1p code point value (priority) for all traffic received from the Fibre Channel (FC) network on the specified FC interface. Instead of using the default priority 3 (011) for FCoE traffic, the priority is rewritten to the specified priority before being forwarded. This works in conjunction with configuring a fixed classifier on the FC interface. The fixed classifier maps all traffic from the FC network into one lossless forwarding class (the lossless forwarding class must be mapped to the code point specified in the rewrite value). Traffic mapped to the lossless forwarding class uses the IEEE 802.1p priority specified by the code point bits in the rewrite value.

FCoE traffic typically uses priority 3 (IEEE code point 011). The QFX Series default configuration uses IEEE 802.1p priority 3 for FCoE traffic. Rewriting the code point value enables you to change the IEEE 802.1p priority of the FCoE traffic if the Ethernet network uses a different priority than priority 3 (code point 011) for FCoE traffic.

The system supports only one IEEE 802.1p code point value per FC interface, so you cannot configure more than one IEEE 802.1p rewrite value per FC interface. In addition, you can specify only one rewrite value per local FCoE-FC gateway fabric; all interfaces in the local fabric must use the same rewrite value. Attempting to configure FC interfaces in the same local fabric with different rewrite values generates a commit error. You can specify different rewrite values for interfaces that belong to different local FCoE-FC gateway fabrics.



**NOTE:** In order to avoid fate sharing (separate flows that affect each others' throughput), the code point used for the rewrite value should be the only code point used for the lossless FCoE forwarding class (the forwarding class used for the fixed classifier on the Fibre Channel interface). When you configure classifiers for ingress Ethernet interfaces, map only the rewrite value code point to the forwarding class.

For example, if the rewrite value uses code point 101 for lossless FCoE forwarding class `fcoe_fc1`, then in the classifier configuration attached to ingress Ethernet interfaces, code point 101 is the only code point that should be classified to the `fcoe_fc1` forwarding class. Now if you also attach a classifier to an interface that maps code point 110 to forwarding class `fcoe_fc1`, then congestion on priority 110 unfairly (and unintentionally) affects the FCoE traffic that uses priority 101. Both priorities 101 and 110 are classified into

forwarding class `fcoe_fc1`, so the traffic from both priorities shares the same fate.


.....  
The remaining statements are described separately.

**Required Privilege Level**    interfaces—To view this statement in the configuration.  
                                 interface-control—To add this statement to the configuration.

**Related Documentation**

- [forwarding-class \(Fibre Channel Interfaces\) on page 491](#)
- [Example: Configuring IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 356](#)
- [Understanding CoS IEEE 802.1p Priority Remapping on an FCoE-FC Gateway on page 128](#)
- [Understanding CoS Classifiers on page 51](#)

## rx-buffers

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | rx-buffers (on   off);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | [edit interfaces <i>interface-name</i> ether-options <a href="#">configured-flow-control</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Release Information</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Description</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <p>Enable or disable an interface to generate and send Ethernet PAUSE messages. If you enable the receive buffers to generate and send PAUSE messages, when the receive buffers reach a certain level of fullness, the interface sends a PAUSE message to the connected peer. If the connected peer is properly configured, it stops transmitting frames to the interface on the entire link. When the interface receive buffer empties below a certain threshold, the interface sends a message to the connected peer to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>rx-buffers</b> statement with the <b>tx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p> |
| <div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Default</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | on   off—Enable or disable an interface to generate and send Ethernet PAUSE messages.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Required Privilege Level</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Related Documentation</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <ul style="list-style-type: none"> <li>• <a href="#">flow-control on page 486</a></li> <li>• <a href="#">tx-buffers on page 541</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 420](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 421](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)

## scheduler

---

|                                 |                                                                                                                                                                                                                                                                                                          |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>scheduler <i>scheduler-name</i>;</code>                                                                                                                                                                                                                                                            |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service scheduler-maps <i>map-name</i> forwarding-class <i>class-name</i></a> ]                                                                                                                                                                                               |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                        |
| <b>Description</b>              | Map a scheduler to a forwarding class using a scheduler map.                                                                                                                                                                                                                                             |
| <b>Options</b>                  | <i>scheduler-name</i> —Name of the scheduler to map to the forwarding class.                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                 |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 289</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 86</a></li> </ul> |

## scheduler-map

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <code>scheduler-map <i>map-name</i>;</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service traffic-control-profiles <i>traffic-control-profile-name</i></a> ]                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Associate a scheduler map with a traffic control profile.                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <i>map-name</i> —Name of the scheduler map.                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220</a></li> <li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 298</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 289</a></li> <li>• <a href="#">Understanding CoS Traffic Control Profiles on page 98</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 86</a></li> </ul> |

## scheduler-maps



---

|                                 |                                                                                                                                                                                                                                                                                                      |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>scheduler-maps {<br/>  map-name {<br/>    forwarding-class class-name scheduler scheduler-name;<br/>  }<br/>}</pre>                                                                                                                                                                             |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                                                                                             |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                    |
| <b>Description</b>              | Specify a scheduler map name to map a scheduler configuration to a forwarding class.                                                                                                                                                                                                                 |
| <b>Options</b>                  | <p><i>map-name</i>—Name of the scheduler map.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                          |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                  |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220</a></li><li>• <a href="#">Example: Configuring Queue Schedulers on page 289</a></li><li>• <a href="#">Understanding CoS Output Queue Schedulers on page 86</a></li></ul> |

## schedulers

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre> schedulers {   scheduler-name {     buffer-size (percent <i>percentage</i>   remainder);     drop-profile-map loss-priority (low   medium-high   high) protocol <i>protocol</i> drop-profile       drop-profile-name;     explicit-congestion-notification;     priority <i>priority</i>;     shaping-rate (<i>rate</i>   percent <i>percentage</i>);     transmit-rate (percent <i>percentage</i>);   } } </pre>                    |
| <b>Hierarchy Level</b>          | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Description</b>              | Specify scheduler name and parameter values such minimum bandwidth ( <b>transmit-rate</b> ), maximum bandwidth ( <b>shaping-rate</b> ), and priority ( <b>priority</b> ).                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>scheduler-name</b> —Name of the scheduler.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                             |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220</a></li> <li>• <a href="#">Example: Configuring Queue Schedulers on page 289</a></li> <li>• <a href="#">Example: Configuring Drop Profile Maps on page 281</a></li> <li>• <a href="#">Example: Configuring ECN</a></li> <li>• <a href="#">Understanding CoS Output Queue Schedulers on page 86</a></li> </ul> |

## shaping-rate

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>shaping-rate (rate   percent <i>percentage</i>);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Hierarchy Level</b>     | [edit <b>class-of-service</b> <b>schedulers</b> <i>scheduler-name</i> ],<br>[edit <b>class-of-service</b> <b>traffic-control-profiles</b> <i>profile-name</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Description</b>         | <p>Configure the shaping rate. The shaping rate throttles the rate of packet transmission by setting a maximum bandwidth (rate in bits per second) or a maximum percentage of bandwidth for a queue or a forwarding class set. You specify the maximum bandwidth for a queue by using a scheduler map to associate a forwarding class (queue) with a scheduler that has a configured shaping rate. You specify the maximum bandwidth for a forwarding class set by setting the shaping rate for a traffic control profile, and then applying the traffic control profile and a forwarding class set to an interface.</p> <p>We recommend that you configure the shaping rate as an absolute maximum usage and not as additional usage beyond the configured transmit rate (the minimum guaranteed bandwidth for a queue) or the configured guaranteed rate (the minimum guaranteed bandwidth for a forwarding class set).</p> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> When you set the maximum bandwidth (shaping-rate value) for a queue or for a priority group at 100 Kbps or less, the traffic shaping behavior is accurate only within +/- 20 percent of the configured shaping-rate value.</p> </div> <div style="border: 1px solid #ccc; padding: 10px; margin-top: 10px;"> <p> <b>NOTE:</b> We recommend that you always apply a shaping rate to strict-high priority queues to prevent them from starving other queues. If you do not apply a shaping rate to limit the amount of bandwidth a strict-high priority queue can use, then the strict-high priority queue can use all of the available port bandwidth and starve other queues on the port.</p> </div> |
| <b>Default</b>             | If you do not configure a shaping rate, the default shaping rate is 100 percent (all of the available bandwidth), which is the equivalent of no rate shaping.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Options</b>             | <p><b>percent <i>percentage</i></b>—Shaping rate as a percentage of the available interface bandwidth.<br/><b>Range:</b> 1 through 100 percent</p> <p><b><i>rate</i></b>—Peak (maximum) rate, in bits per second (bps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).<br/><b>Range:</b> 1000 through 10,000,000,000 bps</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |



**Required Privilege** interfaces—To view this statement in the configuration.  
**Level** interface-control—To add this statement to the configuration.

- Related Documentation**
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
  - [Example: Configuring Queue Schedulers on page 289](#)
  - [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
  - [Understanding CoS Output Queue Schedulers on page 86](#)
  - [Understanding CoS Traffic Control Profiles on page 98](#)

## shared-buffer

```
Syntax shared-buffer {
 egress {
 buffer-partition (lossless | lossy | multicast) {
 percent percent
 }
 percent percent;
 }
 ingress {
 percent percent;
 buffer-partition (lossless | lossless-headroom | lossy) {
 percent percent
 }
 }
 }
```

Hierarchy Level [edit [class-of-service](#)]

Release Information Statement introduced in Junos OS Release 12.3 for the QFX Series.

**Description** Configure the global shared buffer pool allocation to ports. Shared buffers are a pool of buffer space that the system can allocate dynamically across all of its ports as memory space is needed. Some buffer space is reserved for dedicated buffers (buffers allocated permanently to ports), headroom buffers (buffers that help prevent packet loss on lossless flows), and other buffers.

Configure the way the system uses the available (user-configurable) buffer space by setting the **shared-buffer** percentage for the ingress buffer pool and for the egress buffer pool.

The percentage you specify is the percentage of available buffer space allocated to the global shared ingress buffer pool or to the global shared egress buffer pool. If you allocate less than 100 percent of the available buffer space to the shared buffer pool, the remaining buffer space is added to the dedicated buffer pool. (You cannot directly configure the dedicated buffer pool for each port; dedicated buffers are allocated evenly across all the ports.)



**CAUTION:** Changing the buffer configuration is a disruptive event. Traffic stops on *all* ports until the buffer reprogramming is complete.

You can also partition the ingress shared buffer pool and the egress shared buffer pool to adjust the buffer allocations for different mixes of network traffic (best-effort, lossless, multicast) using the **buffer-partition** statement.

**Options** The statements are explained separately.

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Required Privilege Level</b> | interfaces—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 365</a></li><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 377</a></li><li>• <a href="#">Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 383</a></li><li>• <a href="#">Configuring Global Ingress and Egress Shared Buffers on page 425</a></li><li>• <a href="#">Understanding CoS Buffer Configuration on page 131</a></li></ul> |

## traceoptions (Class of Service)

|                            |                                                                                                                                                                                                                                    |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <pre> traceoptions {     file <i>filename</i> &lt;size <i>size</i>&gt; &lt;files <i>number</i>&gt;     &lt;world-readable   no-world-readable&gt;;     flag <i>flag</i> &lt;<i>flag-modifier</i>&gt;;     no-remote-trace } </pre> |
| <b>Hierarchy Level</b>     | [edit <a href="#">class-of-service</a> ]                                                                                                                                                                                           |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                  |
| <b>Description</b>         | Set class-of-service (CoS) tracing options.                                                                                                                                                                                        |



**NOTE:** The `traceoptions` statement is not supported on the QFabric system.

**Default** Traceoptions is disabled.

**Options** **file *filename***—Name of the file to receive the tracing operation output. Enclose the name in quotation marks. Traceoption output files are located in the `/var/log/` directory.

**files *number***—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file.0***. The traceoption output continues in a second trace file named ***trace-file.1***. When ***trace-file.1*** reaches its maximum size, output continues in a third file named ***trace-file.2***, and so on. When the maximum number of trace files is reached, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the size option.

**Range:** 2 through 1000 files

**Default:** 1 trace file

***flag***—Tracing operation to perform. To specify more than one tracing operation, include multiple ***flag*** statements:

- **all**—Trace all operations.
- **asynch**—Trace asynchronous configuration processing.
- **chassis-scheduler**—Trace chassis stream scheduler processing.
- **cos-adjustment**—Trace CoS rate adjustments.
- **dynamic**—Trace dynamic CoS functions.
- **hardware-database**—Trace the chassis hardware database related processing.
- **init**—Trace initialization events.

- **performance-monitor**—Trace performance monitor counters.
- **process**—Trace configuration processing.
- **restart**—Trace restart processing.
- **route-socket**—Trace route-socket events.
- **show**—Trace show command servicing.
- **snmp**—Trace SNMP-related processing.
- **util**—Trace utilities.

The following are the global tracing options:

- **all**—Perform all tracing operations
- **parse**—Trace parser processing.

**no-remote-trace**—(Optional) Disable remote tracing.

**no-world-readable**—(Optional) Prevent any user from reading the log file.

**size size**—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named **trace-file** reaches its maximum size, it is renamed **trace-file.0**. Incoming tracefile data is logged in the now empty **trace-file**. When **trace-file** again reaches its maximum size, **trace-file.0** is renamed **trace-file.1** and **trace-file** is renamed **trace-file.0**. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

**Syntax:** **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

**Range:** 10 KB through the maximum file size of 4 GB (maximum is lower if 4 GB is not supported on your system)

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.




|                           |                                                             |
|---------------------------|-------------------------------------------------------------|
| <b>Required Privilege</b> | routing—To view this statement in the configuration.        |
| <b>Level</b>              | routing-control—To add this statement to the configuration. |

## traffic-control-profiles

---

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | <pre>traffic-control-profiles <i>profile-name</i> {<br/>    <b>guaranteed-rate</b> (<i>rate</i>  percent <i>percentage</i>);<br/>    <b>scheduler-map</b> <i>map-name</i>;<br/>    <b>shaping-rate</b> (<i>rate</i>  percent <i>percentage</i>);<br/>}</pre>                                                                                                                                                                                                                                                                                                                 |
| <b>Hierarchy Level</b>          | [edit <b>class-of-service</b> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>Description</b>              | Configure traffic shaping and scheduling profiles for forwarding class sets (priority groups) to implement enhanced transmission selection (ETS) or for logical interfaces.                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>Options</b>                  | <p><b>profile-name</b>—Name of the traffic-control profile. This name is also used to specify an output traffic control profile.</p> <p>The remaining statements are explained separately.</p>                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | <p>interfaces—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"><li>• <a href="#">Example: Configuring CoS Hierarchical Port Scheduling (ETS) on page 220</a></li><li>• <a href="#">Example: Configuring Traffic Control Profiles (Priority Group Scheduling) on page 298</a></li><li>• <a href="#">Example: Configuring Forwarding Class Sets on page 286</a></li><li>• <a href="#">Assigning CoS Components to Interfaces on page 429</a></li><li>• <a href="#">output-traffic-control-profile on page 516</a></li><li>• <a href="#">Understanding CoS Traffic Control Profiles on page 98</a></li></ul> |

## transmit-rate

|                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>              | <code>transmit-rate (rate   percent <i>percentage</i>);</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>Hierarchy Level</b>     | [edit <code>class-of-service schedulers scheduler-name</code> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>Release Information</b> | Statement introduced in Junos OS Release 11.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>         | Specify the minimum transmission rate or percentage for a queue (forwarding class) scheduler. The transmit rate also determines the amount of excess (extra) priority group bandwidth that the queue can share. Extra priority group bandwidth is allocated among the queues in the priority group in proportion to the transmit rate of each queue.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|                            | <p> <b>NOTE:</b> The <code>transmit-rate</code> setting works only if you also configure the <code>guaranteed-rate</code> in the traffic control profile that is attached to the forwarding class set to which the queue belongs. If you do not configure the guaranteed rate, the minimum guaranteed rate for individual queues that you set using the <code>transmit-rate</code> statement does not work. The sum of all queue transmit rates in a forwarding class set should not exceed the traffic control profile guaranteed rate.</p> <p> <b>NOTE:</b> You cannot configure a transmit rate for strict-high priority queues. Queues (forwarding classes) with a configured transmit rate cannot be included in a forwarding class set that has strict-high priority queues.</p> <p> <b>NOTE:</b> For transmit rates below 1 Gbps, we recommend that you configure the transmit rate as a percentage instead of as a fixed rate. This is because the system converts fixed rates into percentages and may round small fixed rates to a lower percentage. For example, a fixed rate of 350 Mbps is rounded down to 3 percent instead of 3.5 percent.</p> |
| <b>Default</b>             | If you do not include this statement, the default scheduler transmission rate and buffer size percentages for queues 0 through 11 are:                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

| Queue Number        | Default Minimum Guaranteed Bandwidth |
|---------------------|--------------------------------------|
| 0 (best-effort)     | 5 %                                  |
| 1                   | 0                                    |
| 2                   | 0                                    |
| 3 (fcoe)            | 35 %                                 |
| 4 (no-loss)         | 35 %                                 |
| 5                   | 0                                    |
| 6                   | 0                                    |
| 7 (network control) | 5 %                                  |
| 8 (mcast)           | 20 %                                 |
| 9                   | 0                                    |
| 10                  | 0                                    |
| 11                  | 0                                    |

Configure schedulers if you want to change the minimum guaranteed bandwidth and other queue characteristics.

**Options** **rate**—Minimum transmission rate for the queue, in bps. You can specify a value in bits-per-second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

**Range:** 1000 through 10,000,000,000 bps

**percent** **percentage**—Minimum percentage of transmission capacity allocated to the queue. A percentage of zero means that there is no minimum bandwidth guarantee for the queue.

**Range:** 0 through 100 percent


**Required Privilege Level** interfaces—To view this statement in the configuration.  
interface-control—To add this statement to the configuration.

**Related Documentation**

- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)



## tx-buffers

|                                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Syntax</b>                   | tx-buffers (on   off);                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Hierarchy Level</b>          | [edit interfaces <i>interface-name</i> ether-options <a href="#">configured-flow-control</a> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>Release Information</b>      | Statement introduced in Junos OS Release 12.1 for the QFX Series.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Description</b>              | <p>Enable or disable an interface to respond to received Ethernet PAUSE messages. If you enable the transmit buffers to respond to PAUSE messages, when the interface receives a PAUSE message from the connected peer, the interface stops transmitting frames on the entire link. When the receive buffer on the connected peer empties below a certain threshold, the peer interface sends a message to the paused interface to resume sending frames.</p> <p>Ethernet PAUSE prevents buffers from overflowing and dropping packets during periods of network congestion. If the other devices in the network are also configured to support PAUSE, PAUSE supports lossless operation. Use the <b>tx-buffers</b> statement with the <b>rx-buffers</b> statement to configure asymmetric Ethernet PAUSE on an interface. (Use the <b>flow-control</b> statement to enable symmetric PAUSE and the <b>no-flow-control</b> statement to disable symmetric PAUSE on an interface. Symmetric flow control and asymmetric flow control are mutually exclusive features. If you attempt to configure both, the switch returns a commit error.)</p> |
|                                 | <div>  <p><b>NOTE:</b> Ethernet PAUSE temporarily stops transmitting all traffic on a link when the buffers fill to a certain threshold. To temporarily pause traffic on individual “lanes” of traffic (each lane contains the traffic associated with a particular IEEE 802.1p code point, so there can be eight lanes of traffic on a link), use priority-based flow control (PFC).</p> <p>Ethernet PAUSE and PFC are mutually exclusive features, so you cannot configure both of them on the same interface. If you attempt to configure both Ethernet PAUSE and PFC on an interface, the switch returns a commit error.</p> </div>                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Default</b>                  | Flow control is disabled. You must explicitly configure Ethernet PAUSE flow control on interfaces.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Options</b>                  | on   off—Enable or disable an interface to respond to an Ethernet PAUSE message.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>Required Privilege Level</b> | interface—To view this statement in the configuration.<br>interface-control—To add this statement to the configuration.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>Related Documentation</b>    | <ul style="list-style-type: none"> <li>• <a href="#">flow-control on page 486</a></li> <li>• <a href="#">rx-buffers on page 528</a></li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

- [Enabling and Disabling CoS Symmetric Ethernet PAUSE Flow Control on page 420](#)
- [Configuring CoS Asymmetric Ethernet PAUSE Flow Control on page 421](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)

---

## unit

---

**Syntax**     `unit logical-unit-number {  
                  classifiers {  
                    (dscp | dscp-ipv6 | ieee-802.1 | exp) (classifier-name | default);  
                  }  
                  forwarding-class class-name;  
                  rewrite-rules {  
                    (dscp | dscp-ipv6 | ieee-802.1) (classifier-name | default);  
                  }  
                }`

**Hierarchy Level**     [edit [class-of-service interfaces](#) *interface-name*]

**Release Information**     Statement introduced in Junos OS Release 11.1 for the QFX Series.

**Description**     Configure a logical interface on the physical device. You must configure a logical interface to use the physical device.

**Options**     *logical-unit-number*—Number of the logical unit.

**Range:** 0 through 16,385

The remaining statements are explained separately.

**Required Privilege Level**     *interfaces*—To view this statement in the configuration.  
                                      *interface-control*—To add this statement to the configuration.

**Related Documentation**     • [Assigning CoS Components to Interfaces on page 429](#)

## PART 3

# Administration

- [Routine Monitoring on page 545](#)
- [Operational Commands on page 553](#)



## CHAPTER 8

# Routine Monitoring

- [Monitoring CoS Classifiers on page 545](#)
- [Monitoring CoS Forwarding Classes on page 546](#)
- [Monitoring Interfaces That Have CoS Components on page 547](#)
- [Monitoring CoS Rewrite Rules on page 548](#)
- [Monitoring CoS Scheduler Maps on page 549](#)
- [Monitoring CoS Value Aliases on page 550](#)

### Monitoring CoS Classifiers

**Purpose** Display the mapping of incoming CoS values to forwarding class and loss priority for each classifier.

**Action** To monitor CoS classifiers in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier
```

To monitor a particular classifier in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier name classifier-name
```

To monitor a particular type of classifier in the CLI, enter the CLI command:

```
user@switch> show class-of-service classifier type classifier-type
```

**Meaning** [Table 111 on page 545](#) summarizes key output fields for CoS classifiers.

**Table 111: Summary of Key CoS Classifier Output Fields**

| Field           | Values                                                                                                                                                                                                                                                                |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Classifier      | Name of a classifier.                                                                                                                                                                                                                                                 |
| Code point type | Type of classifier: <ul style="list-style-type: none"><li>• <b>dscp</b>—All classifiers of the DSCP type.</li><li>• <b>ieee-802.1</b>—All classifiers of the IEEE 802.1 type.</li><li>• <b>ieee-mcast</b>—All classifiers of the IEEE 802.1 multicast type.</li></ul> |

Table 111: Summary of Key CoS Classifier Output Fields (*continued*)

| Field            | Values                                                                                                                                                                                              |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Index            | Internal index of the classifier.                                                                                                                                                                   |
| Code point       | DSCP or IEEE 802.1 code point value of the incoming packets, in bits. These values are used for classification.                                                                                     |
| Forwarding Class | Name of the forwarding class that the classifier assigns to an incoming packet. This class affects the forwarding and scheduling policies that are applied to the packet as it transits the switch. |
| Loss Priority    | Loss priority value that the classifier assigns to the incoming packet based on its code point value.                                                                                               |

- Related Documentation**
- [Defining CoS Unicast BA Classifiers \(DSCP, DSCP IPv6, IEEE 802.1p\) on page 404](#)
  - [Defining CoS Multidestination \(Multicast, Broadcast, DLF\) BA Classifiers on page 406](#)

## Monitoring CoS Forwarding Classes

**Purpose** Use the monitoring functionality to view the current assignment of CoS forwarding classes to queue numbers on the system.

**Action** To monitor CoS forwarding classes in the CLI, enter the following CLI command:

```
user@switch> show class-of-service forwarding-class
```

**Meaning** [Table 112 on page 546](#) summarizes key output fields for CoS forwarding classes.

Table 112: Summary of Key CoS Forwarding Class Output Fields

| Field            | Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Forwarding Class | <p>Names of forwarding classes assigned to queue numbers. By default, the following unicast forwarding classes are assigned to queues 0, 3, 4, and 7, respectively:</p> <ul style="list-style-type: none"> <li>• <b>best-effort</b>—Provides no special CoS handling of packets. Loss priority is typically not carried in a CoS value.</li> <li>• <b>fcoe</b>—Provides guaranteed delivery for Fibre Channel over Ethernet (FCoE) traffic.</li> <li>• <b>no-loss</b>—Provides guaranteed delivery for TCP lossless traffic</li> <li>• <b>network-control</b>—Packets can be delayed but not dropped.</li> </ul> <p>By default, the following multidestination forwarding class is assigned to queue 8:</p> <ul style="list-style-type: none"> <li>• <b>mcast</b>—Provides no special CoS handling of packets.</li> </ul> |

Table 112: Summary of Key CoS Forwarding Class Output Fields (*continued*)

| Field   | Values                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Queue   | Queue number corresponding to the forwarding class name.<br><br>By default, four queues (0, 3, 4, and 7) are assigned to unicast forwarding classes and one queue (8) is assigned to a multidestination forwarding class.                                                                                                                                                                                                                                                                                                                                                                |
| No-Loss | Packet drop attribute associated with each forwarding class: <ul style="list-style-type: none"> <li>Disabled—The forwarding class is configured for lossy transport (packets might drop during periods of congestion)</li> <li>Enabled—The forwarding class is configured for lossless transport</li> </ul> <p><b>NOTE:</b> To achieve lossless transport, you must ensure that priority-based flow control (PFC) and DCBX are properly configured on the lossless priority (IEEE 802.1p code point), and that sufficient port bandwidth is reserved for the lossless traffic flows.</p> |

- Related Documentation**
- [Defining CoS Forwarding Classes on page 408](#)
  - [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)

## Monitoring Interfaces That Have CoS Components

**Purpose** Use the monitoring functionality to display details about the physical and logical interfaces and the CoS components assigned to them.

**Action** To monitor interfaces that have CoS components in the CLI, enter the command:

```
user@switch> show class-of-service interface
```

To monitor a specific interface in the CLI, enter the command:

```
user@switch> show class-of-service interface interface-name
```

**Meaning** [Table 113 on page 547](#) summarizes key output fields for CoS interfaces.

Table 113: Summary of Key CoS Interfaces Output Fields

| Field              | Values                                                              |
|--------------------|---------------------------------------------------------------------|
| Physical interface | Name of a physical interface to which CoS components are assigned.  |
| Index              | Index of this interface or the internal index of a specific object. |
| Queues supported   | Number of queues you can configure on the interface.                |
| Queues in use      | Number of queues currently configured.                              |
| Scheduler map      | Name of the scheduler map associated with this interface.           |

Table 113: Summary of Key CoS Interfaces Output Fields (*continued*)

| Field                         | Values                                                                                                                                                                                                                                        |
|-------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Congestion-notification       | Status of congestion notification (enabled or disabled).                                                                                                                                                                                      |
| Rewrite Input IEEE Code-point | (Fibre Channel NP_Port interfaces only) IEEE 802.1p code point (priority) the interface assigns to incoming Fibre Channel (FC) traffic when the interface encapsulates the FC traffic in Ethernet before forwarding it onto the FCoE network. |
| Logical Interface             | Name of a logical interface on the physical interface to which CoS components are assigned.                                                                                                                                                   |
| Object                        | Category of an object—for example, <b>classifier</b> , <b>scheduler-map</b> , or <b>rewrite</b> .                                                                                                                                             |
| Name                          | Name of the object—for example, <b>ba-classifier</b> .                                                                                                                                                                                        |
| Type                          | Type of the object—for example, <b>ieee8021p</b> for a classifier.                                                                                                                                                                            |

**Related Documentation** • [Assigning CoS Components to Interfaces on page 429](#)

## Monitoring CoS Rewrite Rules

- Purpose** Use the monitoring functionality to display information about CoS value rewrite rules, which are based on the forwarding class and loss priority.
- Action** To monitor CoS rewrite rules in the CLI, enter the CLI command:
- ```
user@switch> show class-of-service rewrite-rule
```
- To monitor a particular rewrite rule in the CLI, enter the CLI command:
- ```
user@switch> show class-of-service rewrite-rule name rewrite-rule-name
```
- To monitor a particular type of rewrite rule (for example, DSCP, DSCP IPv6, or IEEE-802.1) in the CLI, enter the CLI command:
- ```
user@switch> show class-of-service rewrite-rule type rewrite-rule-type
```
- Meaning** [Table 114 on page 548](#) summarizes key output fields for CoS rewrite rules.

Table 114: Summary of Key CoS Rewrite Rule Output Fields

Field	Values
Rewrite rule	Name of the rewrite rule.

Table 114: Summary of Key CoS Rewrite Rule Output Fields (*continued*)

Field	Values
Code point type	Rewrite rule type: <ul style="list-style-type: none"> • dscp—For IPv4 DiffServ traffic. • dscp-ipv6—For IPv6 Diffserv traffic. • ieee-802.1—For Layer 2 traffic.
Index	Internal index for the rewrite rule.
Forwarding class	Name of the forwarding class that is used to determine CoS values for rewriting in combination with loss priority. Rewrite rules are applied to CoS values in outgoing packets based on forwarding class and loss priority setting.
Loss priority	Level of loss priority that is used to determine CoS values for rewriting in combination with forwarding class.
Code point	Rewrite code point value.

Related Documentation • [Defining CoS Rewrite Rules on page 427](#)

Monitoring CoS Scheduler Maps

Purpose Use the monitoring functionality to display assignments of CoS forwarding classes to schedulers.

Action To monitor CoS scheduler maps in the CLI, enter the CLI command:

```
user@switch> show class-of-service scheduler-map
```

To monitor a specific scheduler map in the CLI, enter the CLI command:

```
user@switch> show class-of-service scheduler-map scheduler-map-name
```

Meaning [Table 115 on page 549](#) summarizes key output fields for CoS scheduler maps.

Table 115: Summary of Key CoS Scheduler Maps Output Fields

Field	Values
Scheduler map	Name of the scheduler map.
Index	Index of a specific object—scheduler maps, schedulers, or drop profiles.
Scheduler	Name of the scheduler.
Forwarding class	Names of the forwarding classes to which the scheduler is assigned.

Table 115: Summary of Key CoS Scheduler Maps Output Fields (*continued*)

Field	Values
Transmit rate	Configured transmit rate of the scheduler as a percentage of the total interface bandwidth.
Priority	<p>Scheduling priority of a queue:</p> <ul style="list-style-type: none"> • strict-high or high—Packets in this queue are transmitted first. Only one queue can be configured as strict-high or high. • low—Packets in this queue are transmitted after packets in the strict-high queue.
Drop Profiles	Name and index of a drop profile that is assigned to a specific loss priority and protocol pair.
Loss Priority	Drop profile associated with each packet loss priority. You can configure different drop profiles for low , medium-high , and high loss priority traffic.
Protocol	Transport protocol of the drop profile for the particular priority.
Name	Name of the drop profile.

Related Documentation • [Defining CoS Queue Schedulers on page 410](#)

Monitoring CoS Value Aliases

Purpose Use the monitoring functionality to display information about the CoS value aliases that the system is currently using to represent DSCP and IEEE 802.1p code point bits.

Action To monitor CoS value aliases in the CLI, enter the CLI command:

```
user@switch> show class-of-service code-point-aliases
```

To monitor a specific type of code-point alias (for example, DSCP or IEEE 802.1) in the CLI, enter the CLI command:

```
user@switch> show class-of-service code-point-aliases ieee-802.1
```

Meaning [Table 116 on page 550](#) summarizes key output fields for CoS value aliases.

Table 116: Summary of Key CoS Value Alias Output Fields

Field	Values
Code point type	<p>Type of the CoS value:</p> <ul style="list-style-type: none"> • dscp—Examines Layer 3 packet headers for IP packet classification. • ieee-802.1—Examines Layer 2 packet headers for packet classification.
Alias	Name given to a set of bits—for example, af11 is a name for bits 001010 .

Table 116: Summary of Key CoS Value Alias Output Fields (*continued*)

Field	Values
Bit pattern	Set of bits associated with the alias.

Related Documentation • [Defining CoS Code-Point Aliases on page 404](#)

CHAPTER 9

Operational Commands

- `show class-of-service`
- `show class-of-service classifier`
- `show class-of-service code-point-aliases`
- `show class-of-service congestion-notification`
- `show class-of-service drop-profile`
- `show class-of-service forwarding-class`
- `show class-of-service forwarding-class-set`
- `show class-of-service forwarding-table`
- `show class-of-service forwarding-table classifier`
- `show class-of-service forwarding-table classifier mapping`
- `show class-of-service forwarding-table drop-profile`
- `show class-of-service forwarding-table rewrite-rule`
- `show class-of-service forwarding-table rewrite-rule mapping`
- `show class-of-service forwarding-table scheduler-map`
- `show class-of-service interface`
- `show class-of-service multi-destination`
- `show class-of-service rewrite-rule`
- `show class-of-service scheduler-map`
- `show class-of-service shared-buffer`
- `show class-of-service traffic-control-profile`
- `show dcbx`
- `show dcbx neighbors`
- `show interfaces queue`
- `show pfe next-hop`
- `show pfe route`
- `show pfe terse`
- `show pfe version`

show class-of-service

Syntax	show class-of-service
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the class-of-service (CoS) information.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Monitoring CoS Value Aliases on page 550 • Monitoring CoS Classifiers on page 545 • Monitoring CoS Forwarding Classes on page 546 • Monitoring Interfaces That Have CoS Components on page 547 • Monitoring CoS Scheduler Maps on page 549 • Monitoring CoS Rewrite Rules on page 548
List of Sample Output	show class-of- service on page 555
Output Fields	Table 117 on page 554 lists the output fields for the show class-of-service command. Output fields are listed in the approximate order in which they appear.

Table 117: show class-of-service Output Fields

Field Name	Field Description	Level of Output
Forwarding class	The forwarding class configuration: <ul style="list-style-type: none"> • Forwarding class—Name of the forwarding class. • ID—Forwarding class ID. • Queue—Queue number. 	All levels
Code point type	The type of code-point alias: <ul style="list-style-type: none"> • dscp—Aliases for DiffServ code point (DSCP) values. • ieee-802.1—Aliases for IEEE 802.1p values. 	All levels
Alias	Names given to CoS values.	All levels
Bit pattern	Set of bits associated with an alias.	All levels
Classifier	Name of the classifier.	All levels
Code point	Code-point values.	All levels
Loss priority	Loss priority assigned to specific CoS values and aliases of the classifier.	All levels
Rewrite rule	Name of the rewrite rule if one has been configured.	All levels

Table 117: show class-of-service Output Fields (*continued*)

Field Name	Field Description	Level of Output
Drop profile	Name of the drop profile.	All levels
Type	Type of drop profile. QFX Series supports only the discrete type of drop-profile.	All levels
Fill level	Percentage of queue buffer fullness in a drop profile at which packets begin to drop during periods of congestion.	All levels
Scheduler map	Name of the scheduler map.	All levels
Scheduler	Name of the scheduler.	All levels
Transmit rate	Transmission rate of the scheduler.	All levels
Buffer size	Delay buffer size in the queue.	All levels
Drop profiles	Drop profiles configured for the specified scheduler.	All levels
Protocol	Transport protocol corresponding to the drop profile.	All levels
Name	Name of the drop profile.	All levels
Queues supported	Number of queues that can be configured on the interface.	All levels
Queues in use	Number of queues currently configured.	All levels
Physical interface	Name of the physical interface.	All levels
Scheduler map	Name of the scheduler map.	All levels
Congestion-notification	Enabled if a congestion notification profile is applied to the interface; disabled if no congestion notification profile is applied to the interface.	All levels
Forwarding class set	Name of the forwarding class set (priority group).	
Index	Internal index of an object.	All levels

Sample Output

show class-of- service

```

user@switch> show class-of-service
Forwarding class      ID      Queue
best-effort          0        0
fcoe                  1        3
no-loss               2        4
network-control      3        7
mcast                 8        8

Code point type: dscp

```

```

Alias          Bit pattern
af11           001010
af12           001100
...           ...

Code point type: ieee-802.1
Alias          Bit pattern
af11           100
...           ...

Classifier: dscp-default, Code point type: dscp, Index: 7
Code point    Forwarding class    Loss priority
000000        best-effort         low
000001        best-effort         low
...           ...                 ...

Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
Code point    Forwarding class    Loss priority
000           best-effort         low
001           best-effort         low
010           best-effort         low
011           fcoe                low
100           no-loss             low
101           best-effort         low
110           network-control     low
111           network-control     low

Drop profile:<default-drop-profile>, Type: discrete, Index: 1
Fill level
100

Scheduler map: <default>, Index: 2

Scheduler: <default-be>, Forwarding class: best-effort, Index: 21
Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer
Limit: none,
Priority: low
Excess Priority: low
drop-profile-map-set-type: mark
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low            any        1      <default-drop-profile>
  Medium high    any        1      <default-drop-profile>
  High           any        1      <default-drop-profile>

Scheduler: <default-fcoe>, Forwarding class: fcoe, Index: 50
Transmit rate: 35 percent, Rate Limit: none, Buffer size: 35 percent, Buffer
Limit: none,
Priority: low
Excess Priority: low
drop-profile-map-set-type: mark
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low            any        1      <default-drop-profile>
  Medium high    any        1      <default-drop-profile>
  High           any        1      <default-drop-profile>

Scheduler: <default-noloss>, Forwarding class: no-loss, Index: 51
Transmit rate: 35 percent, Rate Limit: none, Buffer size: 35 percent, Buffer
Limit: none,
Priority: low

```



```

Excess Priority: low
drop-profile-map-set-type: mark
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           any       1      <default-drop-profile>
  Medium high   any       1      <default-drop-profile>
  High          any       1      <default-drop-profile>

Scheduler: <default-nc>, Forwarding class: network-control, Index: 23
  Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent, Buffer
Limit: none,
  Priority: low
  Excess Priority: low
  drop-profile-map-set-type: mark
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low           any       1      <default-drop-profile>
    Medium high   any       1      <default-drop-profile>
    High          any       1      <default-drop-profile>

Scheduler: <default-mcast>, Forwarding class: mcast, Index: 49
  Transmit rate: 20 percent, Rate Limit: none, Buffer size: 20 percent, Buffer
Limit: none,
  Priority: low
  Excess Priority: low
  drop-profile-map-set-type: mark
  Drop profiles:
    Loss priority  Protocol  Index  Name
    Low           any       1      <default-drop-profile>
    Medium high   any       1      <default-drop-profile>
    High          any       1      <default-drop-profile>

Physical interface: xe-0/0/0, Index: 129
Queues supported: 12, Queues in use: 12
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled

Physical interface: xe-0/0/1, Index: 130
Queues supported: 12, Queues in use: 12
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled

...           ...           ...

Forwarding class set: lan-fcset, Type: normal-type, Forwarding class set index:
7
  Forwarding class                Index
  best-effort                     0

```

show class-of-service classifier

Syntax	<pre>show class-of-service classifier <name <i>name</i>> <type dscp type dscp-ipv6 type exp type ieee-802.1 type inet-precedence></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	For each class-of-service (CoS) classifier, display the mapping of code point value to forwarding class and loss priority.
Options	<p>none—Display all classifiers.</p> <p>name <i>name</i>—(Optional) Display named classifier.</p> <p>type dscp—(Optional) Display all classifiers of the Differentiated Services code point (DSCP) type.</p> <p>type dscp-ipv6—(Optional) Display all classifiers of the DSCP for IPv6 type.</p> <p>type exp—(Optional) Display all classifiers of the MPLS experimental (EXP) type.</p> <p>type ieee-802.1—(Optional) Display all classifiers of the ieee-802.1 type.</p> <p>type inet-precedence—(Optional) Display all classifiers of the inet-precedence type.</p>
Required Privilege Level	view
List of Sample Output	<p>show class-of-service classifier type ieee-802.1 on page 559</p> <p>show class-of-service classifier type ieee-802.1 (QFX Series) on page 559</p>
Output Fields	Table 118 on page 558 describes the output fields for the show class-of-service classifier command. Output fields are listed in the approximate order in which they appear.

Table 118: show class-of-service classifier Output Fields

Field Name	Field Description
Classifier	Name of the classifier.
Code point type	Type of the classifier: exp (not on EX Series switch), dscp , dscp-ipv6 (not on EX Series switch), ieee-802.1 , or inet-precedence .
Index	Internal index of the classifier.
Code point	Code point value used for classification
Forwarding class	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.

Table 118: show class-of-service classifier Output Fields (*continued*)

Field Name	Field Description
Loss priority	Loss priority value used for classification. For most platforms, the value is high or low . For some platforms, the value is high , medium-high , medium-low , or low .

Sample Output

show class-of-service classifier type ieee-802.1

```

user@host> show class-of-service classifier type ieee-802.1
Classifier: ieee802.1-default, Code point type: ieee-802.1, Index: 3
Code Point      Forwarding Class      Loss priority
000             best-effort           low
001             best-effort           high
010             expedited-forwarding  low
011             expedited-forwarding  high
100             assured-forwarding    low
101             assured-forwarding    medium-high
110             network-control       low
111             network-control       high

Classifier: users-ieee802.1, Code point type: ieee-802.1
Code point      Forwarding class      Loss priority
100             expedited-forwarding  low

```

show class-of-service classifier type ieee-802.1 (QFX Series)

```

user@switch> show class-of-service classifier type ieee-802.1
Classifier: ieee8021p-default, Code point type: ieee-802.1, Index: 11
Code point      Forwarding class      Loss priority
000             best-effort           low
001             best-effort           low
010             best-effort           low
011             fcoe                  low
100             no-loss               low
101             best-effort           low
110             network-control       low
111             network-control       low

Classifier: ieee-mcast, Code point type: ieee-802.1, Index: 46
Code point      Forwarding class      Loss priority
000             mcast                 low
001             mcast                 low
010             mcast                 low
011             mcast                 low
100             mcast                 low
101             mcast                 low
110             mcast                 low
111             mcast                 low

```

show class-of-service code-point-aliases

Syntax	<code>show class-of-service code-point-aliases</code> <code><dscp dscp-ipv6 exp ieee-802.1 inet-precedence></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the mapping of class-of-service (CoS) code point aliases to corresponding bit patterns.
Options	<p>none—Display code point aliases of all code point types.</p> <p>dscp—(Optional) Display Differentiated Services code point (DSCP) aliases.</p> <p>dscp-ipv6—(Optional) Display IPv6 DSCP aliases.</p> <p>exp—(Optional) Display MPLS EXP code point aliases.</p> <p>ieee-802.1—(Optional) Display IEEE-802.1 code point aliases.</p> <p>inet-precedence—(Optional) Display IPv4 precedence code point aliases.</p>
Required Privilege Level	view
List of Sample Output	show class-of-service code-point-aliases exp on page 561
Output Fields	Table 119 on page 560 describes the output fields for the show class-of-service code-point-aliases command. Output fields are listed in the approximate order in which they appear.

Table 119: show class-of-service code-point-aliases Output Fields

Field Name	Field Description
Code point type	Type of the code points displayed: dscp , dscp-ipv6 (not on EX Series switch), exp (not on EX Series switch or the QFX Series), ieee-802.1 , or inet-precedence (not on the QFX Series).
Alias	Alias for a bit pattern.
Bit pattern	Bit pattern for which the alias is displayed.

Sample Output

show class-of-service code-point-aliases exp

```
user@host> show class-of-service code-point-aliases exp
Code point type: exp
Alias      Bit pattern
af11       100
af12       101
be         000
be1        001
cs6        110
cs7        111
ef         010
ef1        011
nc1        110
nc2        111
```

show class-of-service congestion-notification

Syntax	show class-of-service congestion-notification
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display whether priority-based flow control (PFC) is enabled for each IEEE 802.1p code point.
Options	none —Display the PFC state for all IEEE 802.1p code points.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show class-of-service interface on page 587 • Example: Configuring CoS PFC for FCoE Traffic on page 241 • Troubleshooting Dropped FCoE Traffic on page 701
Output Fields	Table 120 on page 562 describes the output fields for the show class-of-service congestion-notification command. Output fields are listed in the approximate order in which they appear.

Table 120: show class-of-service congestion-notification Output Fields

Field Name	Field Description
Type	Type of interfaces on which congestion notification is applied. Congestion notification is applied on input interfaces.
Index	Index of this congestion notification profile.
Name	Name of the congestion notification profile.
Cable Length	Length of the attached physical cable in meters. The default value is 100 meters.
Priority	IEEE 802.1p code point.
PFC	State of PFC for the corresponding code point, either enabled or disabled .
MRU	<p>Maximum receive unit of the interface in bytes. (Incoming traffic that exceeds the MRU size of an interface is dropped.) The default values are:</p> <ul style="list-style-type: none"> • 2500 bytes for priority 3 traffic • 9216 bytes for priority 4 traffic <p>NOTE: If you configure flow control on a priority that is not one of the default flow control priorities, the default MRU value is 2500 bytes. For example, if you configure flow control on priority 5 and you do not configure an MRU value, the default MRU value is 2500 bytes.</p>

Table 120: show class-of-service congestion-notification Output Fields (*continued*)

Field Name	Field Description
Flow-Control-Queues	Output queue mapping to IEEE 802.1p code points (priorities). Explicit output queue to priority mapping overwrites the default configuration, and only explicitly mapped queues are displayed in the output. Flow control is only enabled on a queue when you enable PFC on the corresponding priority in the input stanza of the congestion notification profile.

Sample Output

show class-of-service congestion-notification

```
user@switch> show class-of-service congestion-notification
```

```
Name: fcoe_p3_cnp, Index: 12037
```

```
Type: Input
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2500
100	Enabled	9216
101	Disabled	
110	Disabled	
111	Disabled	

```
Type: Output
```

Priority	Flow-Control-Queues
000	
	0
001	
	1
010	
	2
011	
	3
100	
	4
101	
	5
110	
	6
111	
	7

```
Name: fcoe_p3_p5_cnp, Index: 46484
```

```
Type: Input
```

```
Cable Length: 100 m
```

Priority	PFC	MRU
000	Disabled	
001	Disabled	
010	Disabled	
011	Enabled	2240
100	Disabled	
101	Enabled	2240
110	Disabled	
111	Disabled	

```
Type: Output
```

Priority	Flow-Control-Queues
----------	---------------------

011	3
101	5

show class-of-service drop-profile

Syntax	<code>show class-of-service drop-profile</code> <code><profile-name <i>profile-name</i>></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display data points for each class-of-service (CoS) random early detection (RED) drop profile.
Options	none —Display all drop profiles. profile-name <i>profile-name</i> —(Optional) Display the specified profile only.
Required Privilege Level	view
List of Sample Output	show class-of-service drop-profile on page 566 show class-of-service drop-profile (EX4200 Switch) on page 566 show class-of-service drop-profile (EX8200 Switch) on page 566
Output Fields	Table 121 on page 565 describes the output fields for the show class-of-service drop-profile command. Output fields are listed in the approximate order in which they appear.

Table 121: show class-of-service drop-profile Output Fields

Field Name	Field Description
Drop profile	Name of a drop profile.
Type	Type of drop profile: <ul style="list-style-type: none"> • discrete (default) • interpolated (EX8200 switches only)
Index	Internal index of this drop profile.
Fill Level	Percentage fullness of a queue.
Drop probability	Drop probability at this fill level.

Sample Output

show class-of-service drop-profile

```

user@host> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
    100         100
Drop profile: user-drop-profile, Type: interpolated, Index: 2989
  Fill level    Drop probability
     0           0
     1           1
     2           2
     4           4
     5           5
     6           6
     8           8
    10          10
    12          15
    14          20
    15          23
... 64 entries total
    90          96
    92          96
    94          97
    95          98
    96          98
    98          99
    99          99
   100         100

```

show class-of-service drop-profile (EX4200 Switch)

```

user@switch> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level
    100
Drop profile: dp1, Type: discrete, Index: 40496
  Fill level
    10

```

show class-of-service drop-profile (EX8200 Switch)

```

user@switch> show class-of-service drop-profile
Drop profile: <default-drop-profile>, Type: discrete, Index: 1
  Fill level    Drop probability
    100         100
Drop profile: dp1, Type: interpolated, Index: 40496
  Fill level    Drop probability
     0           0
     1          80
     2          90
     4          90
     5          90
     6          90
     8          90
    10          90
    12          91
    14          91
    15          91
    16          91

```

18	91
20	91
22	92
24	92
25	92
26	92
28	92
30	92
32	93
34	93
35	93
36	93
38	93
40	93
42	94
44	94
45	94
46	94
48	94
49	94
51	95
52	95
54	95
55	95
56	95
58	95
60	95
62	96
64	96
65	96
66	96
68	96
70	96
72	97
74	97
75	97
76	97
78	97
80	97
82	98
84	98
85	98
86	98
88	98
90	98
92	99
94	99
95	99
96	99
98	99
99	99
100	100
Drop profile: dp2, Type: discrete, Index: 40499	
Fill level	Drop probability
10	5
50	50

show class-of-service forwarding-class

Syntax	show class-of-service forwarding-class
Release Information	Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display information about forwarding classes, including the mapping of forwarding classes to queue numbers.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • <i>Example: Configuring CoS on EX Series Switches</i> • <i>Monitoring CoS Forwarding Classes</i> • <i>Defining CoS Forwarding Classes (CLI Procedure)</i> • <i>Configuring CoS Traffic Classification for Ingress Queuing on Oversubscribed Ports on EX8200 Line Cards (CLI Procedure)</i>
List of Sample Output	show class-of-service forwarding-class on page 569 show class-of-service forwarding-class (EX8200 Switch) on page 569 show class-of-service forwarding-class (QFX Series) on page 569
Output Fields	Table 122 on page 568 describes the output fields for the show class-of-service forwarding-class command. Output fields are listed in the approximate order in which they appear.

Table 122: show class-of-service forwarding-class Output Fields

Field Name	Field Description
Forwarding class	Name of the forwarding class.
ID	Forwarding class identifier.
Queue	CoS queue mapped to the forwarding class.
Policing priority	Not supported on EX Series switches or the QFX Series and can be ignored.
Fabric priority	(EX8200 switches only) Fabric priority for the forwarding class, either high or low . Determines the priority of packets entering the switch fabric.

Table 122: show class-of-service forwarding-class Output Fields (*continued*)

Field Name	Field Description
No-Loss	<p>(QFX Series only) Packet loss attribute to differentiate lossless forwarding classes from lossy forwarding classes:</p> <ul style="list-style-type: none"> Disabled—Lossless transport is not configured on the forwarding class (packet drop attribute is drop). Enabled—Lossless transport is configured on the forwarding class (packet drop attribute is no-loss).

Sample Output

show class-of-service forwarding-class

```

user@switch> show class-of-service forwarding-class
Forwarding class      ID      Queue Policing priority
best-effort           0        0      normal
expedited-forwarding  1        5      normal
assured-forwarding    2        1      normal
network-control       3        7      normal

```

Sample Output

show class-of-service forwarding-class (EX8200 Switch)

```

user@switch> show class-of-service forwarding-class
Forwarding class      ID      Queue Fabric priority
best-effort           0        0      low
expedited-forwarding  1        5      low
assured-forwarding    2        1      low
network-control       3        7      low
mcast-be              4        2      low
mcast-ef              5        4      low
mcast-af              6        6      low

```

Sample Output

show class-of-service forwarding-class (QFX Series)

```

user@switch> show class-of-service forwarding-class
Forwarding class      ID      Queue Policing priority No-Loss
best-effort           0        0      normal      Disabled
fcoe                  1        3      normal      Enabled
no-loss               2        4      normal      Enabled
network-control       3        7      normal      Disabled
mcast                 8        8      normal      Disabled

```

show class-of-service forwarding-class-set

Syntax	show class-of-service forwarding-class-set <forwarding-class-set-name>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the forwarding classes associated with each forwarding class set.
Options	<p>none—Display all forwarding class sets.</p> <p>forwarding-class-set-name—(Optional) Display the forwarding classes associated with the specified forwarding class set.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Understanding CoS Fabric Forwarding Class Sets on page 186 • Troubleshooting Dropped FCoE Traffic on page 701
Output Fields	Table 123 on page 570 describes the output fields for the show class-of-service forwarding-class-set command. Output fields are listed in the approximate order in which they appear.

Table 123: show class-of-service forwarding-class-set Output Fields

Field Name	Field Description
Forwarding class set	Name of the forwarding class set.
Type	Internal Junos OS type.
Forwarding class set index	Index of this forwarding class set.
Forwarding class	Name of a forwarding class.
Index	Index of this forwarding class.

Sample Output

show class-of-service forwarding-class-set

```

user@switch> show class-of-service forwarding-class-set
Forwarding class set: san_fcset, Type: normal-type, Forwarding class set index:
37839
  Forwarding class      Index
  fcoe                  1

Forwarding class set: lan_fcset, Type: normal-type, Forwarding class set index:
37840
  Forwarding class      Index

```

best-effort 0

Forwarding class set: multicast_fcset, Type: normal-type, Forwarding class set
index: 37841

Forwarding class	Index
mcast	8

show class-of-service forwarding-table

List of Syntax	Syntax on page 572 Syntax (TX Matrix and TX Matrix Plus Router) on page 572
Syntax	show class-of-service forwarding-table
Syntax (TX Matrix and TX Matrix Plus Router)	show class-of-service forwarding-table <lcc number> <sfc number>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the entire class-of-service (CoS) configuration as it exists in the forwarding table. Executing this command is equivalent to executing all show class-of-service forwarding-table commands in succession.
Options	<p>lcc number—(TX Matrix and TX Matrix Plus router only) (Optional) On a TX Matrix router, display the forwarding table configuration for a specific T640 router (or line-card chassis) configured in a routing matrix. On a TX Matrix Plus router, display the forwarding table configuration for a specific router (or line-card chassis) configured in the routing matrix.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none"> • 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix. • 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix. • 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix. • 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix. <p>sfc number—(TX Matrix Plus routers only) (Optional) Display the forwarding table configuration for the TX Matrix Plus router. Replace <i>number</i> with 0.</p>
Required Privilege Level	view
List of Sample Output	show class-of-service forwarding-table on page 573 show class-of-service forwarding-table lcc (TX Matrix Plus Router) on page 574
Output Fields	See the output field descriptions for show class-of-service forwarding-table commands: <ul style="list-style-type: none"> • show class-of-service forwarding-table classifier • show class-of-service forwarding-table classifier mapping • show class-of-service forwarding-table drop-profile

- *show class-of-service forwarding-table fabric scheduler-map*
- *show class-of-service forwarding-table loss-priority-map*
- *show class-of-service forwarding-table loss-priority-map mapping*
- *show class-of-service forwarding-table rewrite-rule*
- *show class-of-service forwarding-table rewrite-rule mapping*
- *show class-of-service forwarding-table scheduler-map*

Sample Output

show class-of-service forwarding-table

```

user@host> show class-of-service forwarding-table
Classifier table index: 9, # entries: 8, Table type: EXP
Entry #   Code point   Forwarding-class #   PLP
0         000           0                   0
1         001           0                   1
2         010           1                   0
3         011           1                   1
4         100           2                   0
5         101           2                   1
6         110           3                   0
7         111           3                   1

Interface      Index      Table Index/
              Q num      Table type
sp-0/0/0.1001  66         11      IPv4 precedence
sp-0/0/0.2001  67         11      IPv4 precedence
sp-0/0/0.16383 68         11      IPv4 precedence
fe-0/0/0.0     69         11      IPv4 precedence

Interface: sp-0/0/0 (Index: 129, Map index: 2, Map type: FINAL,
Num of queues: 2):
  Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
    Tx rate: 0 Kb (95%), Buffer size: 95 percent
  Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
  Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
    Tx rate: 0 Kb (5%), Buffer size: 5 percent
  Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/0 (Index: 137, Map index: 2, Map type: FINAL,
Num of queues: 2):
  Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
    Tx rate: 0 Kb (95%), Buffer size: 95 percent
  Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
  Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
    Tx rate: 0 Kb (5%), Buffer size: 5 percent
  Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

Interface: fe-0/0/1 (Index: 138, Map index: 2, Map type: FINAL,
Num of queues: 2):
  Entry 0 (Scheduler index: 16, Forwarding-class #: 0):
    Tx rate: 0 Kb (95%), Buffer size: 95 percent
  Priority low

```

```

    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1
Entry 1 (Scheduler index: 18, Forwarding-class #: 3):
    Tx rate: 0 Kb (5%), Buffer size: 5 percent
Priority low
    PLP high: 1, PLP low: 1, PLP medium-high: 1, PLP medium-low: 1

...

RED drop profile index: 1, # entries: 1
      Drop
Entry  Fullness(%)  Probability(%)
   0         100         100

```

show class-of-service forwarding-table lcc (TX Matrix Plus Router)

```

user@host> show class-of-service forwarding-table lcc 0
lcc0-re0:

```

```

-----
Classifier table index: 9, # entries: 64, Table type: IPv6 DSCP

```

Entry #	Code point	Forwarding-class #	PLP
0	000000	0	0
1	000001	0	0
2	000010	0	0
3	000011	0	0
4	000100	0	0
5	000101	0	0
6	000110	0	0
7	000111	0	0
8	001000	0	0
9	001001	0	0
10	001010	0	0
11	001011	0	0
12	001100	0	0
13	001101	0	0
14	001110	0	0
15	001111	0	0
16	010000	0	0
17	010001	0	0
18	010010	0	0
19	010011	0	0
20	010100	0	0
21	010101	0	0
22	010110	0	0
23	010111	0	0
24	011000	0	0
25	011001	0	0
26	011010	0	0
27	011011	0	0
28	011100	0	0
29	011101	0	0
30	011110	0	0
31	011111	0	0
32	100000	0	0
33	100001	0	0
34	100010	0	0
35	100011	0	0
36	100100	0	0
37	100101	0	0
38	100110	0	0
39	100111	0	0

40	101000	0	0
41	101001	0	0
42	101010	0	0
43	101011	0	0
44	101100	0	0
45	101101	0	0
46	101110	0	0
...			

show class-of-service forwarding-table classifier

Syntax	show class-of-service forwarding-table classifier
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the mapping of code point value to queue number and loss priority for each classifier as it exists in the forwarding table.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show class-of-service forwarding-table classifier on page 576
Output Fields	Table 124 on page 576 describes the output fields for the show class-of-service forwarding-table classifier command. Output fields are listed in the approximate order in which they appear.

Table 124: show class-of-service forwarding-table classifier Output Fields

Field Name	Field Description
Classifier table index	Index of the classifier table.
entries	Total number of entries.
Table type	Type of code points in the table: DSCP , EXP (not on the QFX Series), IEEE 802.1 , IPv4 precedence (not on the QFX Series), or IPv6 DSCP .
Entry #	Entry number.
Code point	Code point value used for classification.
Forwarding-class #	Forwarding class to which the code point is assigned.
PLP	Packet loss priority value set by classification. For most platforms, the value can be 0 or 1 . For some platforms, the value is 0 , 1 , 2 , or 3 . The value 0 represents low PLP. The value 1 represents high PLP. The value 2 represents medium-low PLP. The value 3 represents medium-high PLP.

Sample Output

show class-of-service forwarding-table classifier

```

user@host> show class-of-service forwarding-table classifier
Classifier table index: 62436, # entries: 64, Table type: DSCP

Entry #   Code point   Forwarding-class #   PLP

```

0	000000	0	0
1	000001	0	0
2	000010	0	0
3	000011	0	0
4	000100	0	0
5	000101	0	0
6	000110	0	0
7	000111	0	0
8	001000	0	0
9	001001	0	0
10	001010	1	1
11	001011	0	0
...			
60	111100	0	0
61	111101	0	0
62	111110	0	0
63	111111	0	0

show class-of-service forwarding-table classifier mapping

Syntax	show class-of-service forwarding-table classifier mapping
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	For each logical interface, display either the table index of the classifier for a given code point type or the queue number (if it is a fixed classification) in the forwarding table.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show class-of-service forwarding-table classifier mapping on page 578
Output Fields	Table 125 on page 578 describes the output fields for the show class-of-service forwarding-table classifier mapping command. Output fields are listed in the approximate order in which they appear.

Table 125: show class-of-service forwarding-table classifier mapping Output Fields

Field Name	Field Description
Table index/ Q num	If the table type is Fixed , the number of the queue to which the interface is mapped. For all other types, this value is the classifier index number.
Interface	Name of the logical interface. This field can also show the physical interface (QFX Series).
Index	Logical interface index.
Table type	Type of code points in the table: DSCP , EXP (not on the QFX Series), Fixed , IEEE 802.1 , IPv4 precedence (not on the QFX Series), or IPv6 DSCP .

Sample Output

show class-of-service forwarding-table classifier mapping

```

user@host> show class-of-service forwarding-table classifier mapping
Table index/
Interface      Index  Q num  Table type
so-5/0/0.0     10    62436  DSCP
so-0/1/0.0     11    62436  DSCP
so-0/2/0.0     12      1  Fixed
so-0/2/1.0     13    62436  DSCP
so-0/2/1.0     13    62437  IEEE 802.1
so-0/2/2.0     14    62436  DSCP
so-0/2/2.0     14    62438  IPv4 precedence

```


show class-of-service forwarding-table drop-profile

Syntax	show class-of-service forwarding-table drop-profile
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the data points of all random early detection (RED) drop profiles as they exist in the forwarding table.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show class-of-service forwarding-table drop-profile on page 580
Output Fields	Table 126 on page 580 describes the output fields for the show class-of-service forwarding-table drop-profile command. Output fields are listed in the approximate order in which they appear.

Table 126: show class-of-service forwarding-table drop-profile Output Fields

Field Name	Field Description
RED drop profile index	Index of this drop profile.
# entries	Number of entries in a particular RED drop profile index.
Entry	Drop profile entry number.
Fullness(%)	Percentage fullness of a queue.
Drop probability(%)	Drop probability at this fill level.

Sample Output

show class-of-service forwarding-table drop-profile

```

user@host> show class-of-service forwarding-table drop-profile
RED drop profile index: 4, # entries: 1
      Drop
Entry   Fullness(%)  Probability(%)
  0         100           100

RED drop profile index: 8742, # entries: 3
      Drop
Entry   Fullness(%)  Probability(%)
  0         10           10
  1         20           20
  2         30           30

```


RED drop profile index: 24627, # entries: 64

Entry	Fullness(%)	Drop Probability(%)
0	0	0
1	1	1
2	2	2
3	4	4
...		
61	98	99
62	99	99
63	100	100

RED drop profile index: 25393, # entries: 64

Entry	Fullness(%)	Drop Probability(%)
0	0	0
1	1	1
2	2	2
3	4	4
...		
61	98	98
62	99	99
63	100	100

show class-of-service forwarding-table rewrite-rule

Syntax	show class-of-service forwarding-table rewrite-rule
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display mapping of queue number and loss priority to code point value for each rewrite rule as it exists in the forwarding table.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show class-of-service forwarding-table rewrite-rule on page 582
Output Fields	Table 127 on page 582 describes the output fields for the show class-of-service forwarding-table rewrite-rule command. Output fields are listed in the approximate order in which they appear.

Table 127: show class-of-service forwarding-table rewrite-rule Output Fields

Field Name	Field Description
Rewrite table index	Index for this rewrite rule.
# entries	Number of entries in this rewrite rule.
Table type	Type of table: DSCP , EXP (not on the QFX Series), EXP-PUSH-3 (not on the QFX Series), EXP-SWAP-PUSH-2 , (J Series routers only), IEEE 802.1 , IPv4 precedence (not on the QFX Series), IPv6 DSCP , or Fixed .
Q#	Queue number to which this entry is assigned.
Low bits	Code point value for low-priority loss profile.
State	State of this code point: enabled , rewritten , or disabled .
High bits	Code point value for high-priority loss profile.

Sample Output

show class-of-service forwarding-table rewrite-rule

```

user@host> show class-of-service forwarding-table rewrite-rule
Rewrite table index: 3753, # entries: 4, Table type: DSCP
Q#      Low bits  State      High bits  State
0       000111  Enabled    001010    Enabled
2       000000  Disabled   001100    Enabled

```

1	101110	Enabled	110111	Enabled
3	110000	Enabled	111000	Enabled

show class-of-service forwarding-table rewrite-rule mapping

Syntax	show class-of-service forwarding-table rewrite-rule mapping
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	For each logical interface, display the table identifier of the rewrite rule map for each code point type.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show class-of-service forwarding-table rewrite-rule mapping on page 584
Output Fields	Table 128 on page 584 describes the output fields for the show class-of-service forwarding-table rewrite-rule mapping command. Output fields are listed in the approximate order in which they appear.

Table 128: show class-of-service forwarding-table rewrite-rule mapping Output Fields

Field Name	Field Description
Interface	Name of the logical interface. This field can also show the physical interface (QFX Series).
Index	Logical interface index.
Table index	Rewrite table index.
Type	Type of classifier: DSCP , EXP (not on the QFX Series), EXP-PUSH-3 (not on the QFX Series), EXP-SWAP-PUSH-2 (not on the QFX Series), Frame-Relay DE (J Series routers only), IEEE 802.1 , IPv4 precedence (not on the QFX Series), IPv6 DSCP , or Fixed .

Sample Output

show class-of-service forwarding-table rewrite-rule mapping

```

user@host> show class-of-service forwarding-table rewrite-rule mapping
Interface      Index  Table index  Type
so-5/0/0.0     10     3753        DSCP
so-0/1/0.0     11     3753        DSCP
so-0/2/0.0     12     3753        DSCP
so-0/2/1.0     13     3753        DSCP
so-0/2/2.0     14     3753        DSCP
so-0/2/3.0     15     3753        DSCP

```

show class-of-service forwarding-table scheduler-map

Syntax	show class-of-service forwarding-table scheduler-map
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	For each physical interface, display the scheduler map information as it exists in the forwarding table.
Options	This command has no options.
Required Privilege Level	view
List of Sample Output	show class-of-service forwarding-table scheduler-map on page 586
Output Fields	Table 129 on page 585 describes the output fields for the show class-of-service forwarding-table scheduler-map command. Output fields are listed in the approximate order in which they appear.

Table 129: show class-of-service forwarding-table scheduler-map Output Fields

Field Name	Field Description
Interface	Name of the physical interface.
Index	Physical interface index.
Map index	Scheduler map index.
Num of queues	Number of queues defined in this scheduler map.
Entry	Number of this entry in the scheduler map.
Scheduler index	Scheduler policy index.
Forwarding-class #	Forwarding class number to which this entry is applied.
Tx rate	Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword remainder , which indicates that the scheduler receives the remaining bandwidth of the interface.
Max buffer delay	Amount of transmit delay (in milliseconds) or buffer size of the queue. This amount is a percentage of the total interface buffer allocation or the keyword remainder , which indicates that the buffer is sized according to what remains after other scheduler buffer allocations.
Priority	<ul style="list-style-type: none"> high—Queue priority is high. low—Queue priority is low.
PLP high	Drop profile index for a high packet loss priority profile.

Table 129: show class-of-service forwarding-table scheduler-map Output Fields (*continued*)

Field Name	Field Description
PLP low	Drop profile index for a low packet loss priority profile.
PLP medium-high	Drop profile index for a medium-high packet loss priority profile.
PLP medium-low	Drop profile index for a medium-low packet loss priority profile.
TCP PLP high	Drop profile index for a high TCP packet loss priority profile.
TCP PLP low	Drop profile index for a low TCP packet loss priority profile.
Policy is exact	If this line appears in the output, exact rate limiting is enabled. Otherwise, no rate limiting is enabled.

Sample Output

show class-of-service forwarding-table scheduler-map

```

user@host> show class-of-service forwarding-table scheduler-map
Interface: so-5/0/0 (Index: 9, Map index: 17638, Num of queues: 2):
  Entry 0 (Scheduler index: 6090, Forwarding-class #: 0):
    Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
    Priority low
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742
    Policy is exact
  Entry 1 (Scheduler index: 38372, Forwarding-class #: 1):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
    Priority high
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742

Interface: at-6/1/0 (Index: 10, Map index: 17638, Num of queues: 2):
  Entry 0 (Scheduler index: 6090, Forwarding-class #: 0):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (30%), Max buffer delay: 39 bytes (0%)
    Priority high
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742
  Entry 1 (Scheduler index: 38372, Forwarding-class #: 1):
    Traffic chunk: Max = 0 bytes, Min = 0 bytes
    Tx rate: 0 Kb (40%), Max buffer delay: 68 bytes (0%)
    Priority low
    PLP high: 25393, PLP low: 24627, TCP PLP high: 25393, TCP PLP low: 8742

```

show class-of-service interface

Syntax	<pre>show class-of-service interface <comprehensive detail> <interface-name></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Forwarding class map information added in Junos OS Release 9.4.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>Command introduced in Junos OS Release 12.1 for the PTX Series Packet Transport Routers.</p> <p>Command introduced in Junos OS Release 12.2 for the ACX Series Universal Access routers.</p> <p>Options detail and comprehensive introduced in Junos OS Release 11.4.</p>
Description	Display the logical and physical interface associations for the classifier, rewrite rules, and scheduler map objects.
Options	<p>none—Display CoS associations for all physical and logical interfaces.</p> <p>comprehensive—(M Series, MX Series, and T Series routers) (Optional) Display comprehensive quality-of-service (QoS) information about all physical and logical interfaces.</p> <p>detail—(M Series, MX Series, and T Series routers) (Optional) Display QoS and CoS information based on the interface.</p> <p>If the interface <i>interface-name</i> is a physical interface, the output includes:</p> <ul style="list-style-type: none"> • Brief QoS information about the physical interface • Brief QoS information about the logical interface • CoS information about the physical interface • Brief information about filters or policers of the logical interface • Brief CoS information about the logical interface <p>If the interface <i>interface-name</i> is a logical interface, the output includes:</p> <ul style="list-style-type: none"> • Brief QoS information about the logical interface • Information about filters or policers for the logical interface • CoS information about the logical interface <p>interface-name—(Optional) Display class-of-service (CoS) associations for the specified interface.</p>
Required Privilege Level	view
List of Sample Output	show class-of-service interface (Physical) on page 598

[show class-of-service interface \(Logical\) on page 599](#)
[show class-of-service interface \(Gigabit Ethernet\) on page 599](#)
[show class-of-service interface \(PPPoE Interface\) on page 599](#)
[show class-of-service interface \(T4000 Routers with Type 5 FPCs\) on page 599](#)
[show class-of-service interface detail on page 600](#)
[show class-of-service interface comprehensive on page 600](#)
[show class-of-service interface \(ACX Series Routers\) on page 610](#)

Output Fields [Table 130 on page 588](#) describes the output fields for the **show class-of-service interface** command. Output fields are listed in the approximate order in which they appear.

Table 130: show class-of-service interface Output Fields

Field Name	Field Description
Physical interface	Name of a physical interface.
Index	Index of this interface or the internal index of this object.
Dedicated Queues	Status of dedicated queues configured on an interface. Supported only on Trio MPC/MIC interfaces on MX Series routers.
Queues supported	Number of queues you can configure on the interface.
Queues in use	Number of queues currently configured.
Total non-default queues created	Number of queues created in addition to the default queues. Supported only on Trio MPC/MIC interfaces on MX Series routers.
Rewrite Input IEEE Code-point	(QFX Series only) IEEE 802.1p code point (priority) rewrite value. Incoming traffic from the Fibre Channel (FC) SAN is classified into the forwarding class specified in the native FC interface (NP_Port) fixed classifier and uses the priority specified as the IEEE 802.1p rewrite value.
Shaping rate	Maximum transmission rate on the physical interface. You can configure the shaping rate on the physical interface, or on the logical interface, but not on both. Therefore, the Shaping rate field is displayed for either the physical interface or the logical interface.
Scheduler map	Name of the output scheduler map associated with this interface.
Scheduler map forwarding class sets	(QFX Series only) Name of the fabric forwarding class set scheduler map associated with a QFabric system Interconnect device interface.
Input shaping rate	For Gigabit Ethernet IQ2 PICs, maximum transmission rate on the input interface.
Input scheduler map	For Gigabit Ethernet IQ2 PICs, name of the input scheduler map associated with this interface.
Chassis scheduler map	Name of the scheduler map associated with the packet forwarding component queues.
Rewrite	Name and type of the rewrite rules associated with this interface.
Classifier	Name and type of classifiers associated with this interface.

Table 130: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Forwarding-class-map	Name of the forwarding map associated with this interface.
Congestion-notification	(QFX Series only) Congestion notification state, enabled or disabled .
Logical interface	Name of a logical interface.
Object	Category of an object: Classifier , Fragmentation-map (for LSQ interfaces only), Scheduler-map , Rewrite , or Translation Table (for IQE PICs only).
Name	Name of an object.
Type	Type of an object: dscp , dscp-ipv6 , exp , ieee-802.1 , ip , or inet-precedence .
Link-level type	Encapsulation on the physical interface.
MTU	MTU size on the physical interface.
Speed	Speed at which the interface is running.
Loopback	Whether loopback is enabled and the type of loopback.
Source filtering	Whether source filtering is enabled or disabled.
Flow control	Whether flow control is enabled or disabled.
Auto-negotiation	(Gigabit Ethernet interfaces) Whether autonegotiation is enabled or disabled.
Remote-fault	(Gigabit Ethernet interfaces) Remote fault status. <ul style="list-style-type: none"> • Online—Autonegotiation is manually configured as online. • Offline—Autonegotiation is manually configured as offline.

Table 130: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Device flags	<p>The Device flags field provides information about the physical device and displays one or more of the following values:</p> <ul style="list-style-type: none"> • Down—Device has been administratively disabled. • Hear-Own-Xmit—Device receives its own transmissions. • Link-Layer-Down—The link-layer protocol has failed to connect with the remote endpoint. • Loopback—Device is in physical loopback. • Loop-Detected—The link layer has received frames that it sent, thereby detecting a physical loopback. • No-Carrier—On media that support carrier recognition, no carrier is currently detected. • No-Multicast—Device does not support multicast traffic. • Present—Device is physically present and recognized. • Promiscuous—Device is in promiscuous mode and recognizes frames addressed to all physical addresses on the media. • Quench—Transmission on the device is quenched because the output buffer is overflowing. • Recv-All-Multicasts—Device is in multicast promiscuous mode and therefore provides no multicast filtering. • Running—Device is active and enabled.
Interface flags	<p>The Interface flags field provides information about the physical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> • Admin-Test—Interface is in test mode and some sanity checking, such as loop detection, is disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Hardware-Down—Interface is nonfunctional or incorrectly connected. • Link-Layer-Down—Interface keepalives have indicated that the link is incomplete. • No-Multicast—Interface does not support multicast traffic. • No-receive No-transmit—Passive monitor mode is configured on the interface. • Point-To-Point—Interface is point-to-point. • Pop all MPLS labels from packets of depth—MPLS labels are removed as packets arrive on an interface that has the pop-all-labels statement configured. The depth value can be one of the following: <ul style="list-style-type: none"> • 1—Takes effect for incoming packets with one label only. • 2—Takes effect for incoming packets with two labels only. • [1 2]—Takes effect for incoming packets with either one or two labels. • Promiscuous—Interface is in promiscuous mode and recognizes frames addressed to all physical addresses. • Recv-All-Multicasts—Interface is in multicast promiscuous mode and provides no multicast filtering. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational.

Table 130: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Flags	<p>The Logical interface flags field provides information about the logical interface and displays one or more of the following values:</p> <ul style="list-style-type: none"> • ACFC Encapsulation—Address control field Compression (ACFC) encapsulation is enabled (negotiated successfully with a peer). • Device-down—Device has been administratively disabled. • Disabled—Interface is administratively disabled. • Down—A hardware failure has occurred. • Clear-DF-Bit—GRE tunnel or IPsec tunnel is configured to clear the Don't Fragment (DF) bit. • Hardware-Down—Interface protocol initialization failed to complete successfully. • PFC—Protocol field compression is enabled for the PPP session. • Point-To-Point—Interface is point-to-point. • SNMP-Traps—SNMP trap notifications are enabled. • Up—Interface is enabled and operational.
Encapsulation	Encapsulation on the logical interface.
Admin	Administrative state of the interface (Up or Down)
Link	Status of physical link (Up or Down).
Proto	Protocol configured on the interface.
Input Filter	Names of any firewall filters to be evaluated when packets are received on the interface, including any filters attached through activation of dynamic service.
Output Filter	Names of any firewall filters to be evaluated when packets are transmitted on the interface, including any filters attached through activation of dynamic service.
Link flags	<p>Provides information about the physical link and displays one or more of the following values:</p> <ul style="list-style-type: none"> • ACFC—Address control field compression is configured. The Point-to-Point Protocol (PPP) session negotiates the ACFC option. • Give-Up—Link protocol does not continue connection attempts after repeated failures. • Loose-LCP—PPP does not use the Link Control Protocol (LCP) to indicate whether the link protocol is operational. • Loose-LMI—Frame Relay does not use the Local Management Interface (LMI) to indicate whether the link protocol is operational. • Loose-NCP—PPP does not use the Network Control Protocol (NCP) to indicate whether the device is operational. • Keepalives—Link protocol keepalives are enabled. • No-Keepalives—Link protocol keepalives are disabled. • PFC—Protocol field compression is configured. The PPP session negotiates the PFC option.
Hold-times	Current interface hold-time up and hold-time down, in milliseconds.
CoS queues	Number of CoS queues configured.

Table 130: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Last flapped	Date, time, and how long ago the interface went from down to up. The format is Last flapped: year-month-day hour:minute:second:timezone (hour:minute:second ago) . For example, Last flapped: 2002-04-26 10:52:40 PDT (04:33:20 ago) .
Statistics last cleared	Number and rate of bytes and packets received and transmitted on the physical interface. <ul style="list-style-type: none"> • Input bytes—Number of bytes received on the interface. • Output bytes—Number of bytes transmitted on the interface. • Input packets—Number of packets received on the interface. • Output packets—Number of packets transmitted on the interface.
IPv6 transit statistics	Number of IPv6 transit bytes and packets received and transmitted on the logical interface if IPv6 statistics tracking is enabled.
Input errors	Input errors on the interface. The labels are explained in the following list: <ul style="list-style-type: none"> • Errors—Sum of the incoming frame aborts and FCS errors. • Drops—Number of packets dropped by the input queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. • Framing errors—Number of packets received with an invalid frame checksum (FCS). • Runts—Number of frames received that are smaller than the runt threshold. • Giants—Number of frames received that are larger than the giant threshold. • Bucket Drops—Drops resulting from the traffic load exceeding the interface transmit or receive leaky bucket configuration. • Policed discards—Number of frames that the incoming packet match code discarded because they were not recognized or not of interest. Usually, this field reports protocols that Junos OS does not handle. • L3 incompletes—Number of incoming packets discarded because they failed Layer 3 (usually IPv4) sanity checks of the header. For example, a frame with less than 20 bytes of available IP header is discarded. Layer 3 incomplete errors can be ignored by configuring the ignore-l3-incompletes statement. • L2 channel errors—Number of times the software did not find a valid logical interface for an incoming frame. • L2 mismatch timeouts—Number of malformed or short packets that caused the incoming packet handler to discard the frame as unreadable. • HS link CRC errors—Number of errors on the high-speed links between the ASICs responsible for handling the router interfaces. • HS link FIFO overflows—Number of FIFO overflows on the high-speed links between the ASICs responsible for handling the router interfaces.

Table 130: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Output errors	<p>Output errors on the interface. The labels are explained in the following list:</p> <ul style="list-style-type: none"> • Carrier transitions—Number of times the interface has gone from down to up. This number does not normally increment quickly, increasing only when the cable is unplugged, the far-end system is powered down and up, or another problem occurs. If the number of carrier transitions increments quickly (perhaps once every 10 seconds), the cable, the far-end system, or the PIC is malfunctioning. • Errors—Sum of the outgoing frame aborts and FCS errors. • Drops—Number of packets dropped by the output queue of the I/O Manager ASIC. If the interface is saturated, this number increments once for every packet that is dropped by the ASIC's RED mechanism. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Drops field does not always use the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p> <ul style="list-style-type: none"> • Aged packets—Number of packets that remained in shared packet SDRAM so long that the system automatically purged them. The value in this field should never increment. If it does, it is most likely a software bug or possibly malfunctioning hardware. • HS link FIFO underflows—Number of FIFO underflows on the high-speed links between the ASICs responsible for handling the router interfaces. • MTU errors—Number of packets whose size exceeds the MTU of the interface.
Egress queues	Total number of egress queues supported on the specified interface.
Queue counters	<p>CoS queue number and its associated user-configured forwarding class name.</p> <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), the Dropped packets field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
SONET alarms SONET defects	<p>(SONET) SONET media-specific alarms and defects that prevent the interface from passing packets. When a defect persists for a certain period, it is promoted to an alarm. Based on the router configuration, an alarm can ring the red or yellow alarm bell on the router or light the red or yellow alarm LED on the craft interface. See these fields for possible alarms and defects: SONET PHY, SONET section, SONET line, and SONET path.</p>
SONET PHY	<p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET PHY field has the following subfields:</p> <ul style="list-style-type: none"> • PLL Lock—Phase-locked loop • PHY Light—Loss of optical signal

Table 130: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
SONET section	<p>Counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET section field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B1—Bit interleaved parity for SONET section overhead • SEF—Severely errored framing • LOS—Loss of signal • LOF—Loss of frame • ES-S—Errored seconds (section) • SES-S—Severely errored seconds (section) • SEFS-S—Severely errored framing seconds (section)
SONET line	<p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET line field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B2—Bit interleaved parity for SONET line overhead • REI-L—Remote error indication (near-end line) • RDI-L—Remote defect indication (near-end line) • AIS-L—Alarm indication signal (near-end line) • BERR-SF—Bit error rate fault (signal failure) • BERR-SD—Bit error rate defect (signal degradation) • ES-L—Errored seconds (near-end line) • SES-L—Severely errored seconds (near-end line) • UAS-L—Unavailable seconds (near-end line) • ES-LFE—Errored seconds (far-end line) • SES-LFE—Severely errored seconds (far-end line) • UAS-LFE—Unavailable seconds (far-end line)

Table 130: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
SONET path	<p>Active alarms and defects, plus counts of specific SONET errors with detailed information.</p> <ul style="list-style-type: none"> • Seconds—Number of seconds the defect has been active. • Count—Number of times that the defect has gone from inactive to active. • State—State of the error. A state other than OK indicates a problem. <p>The SONET path field has the following subfields:</p> <ul style="list-style-type: none"> • BIP-B3—Bit interleaved parity for SONET section overhead • REI-P—Remote error indication • LOP-P—Loss of pointer (path) • AIS-P—Path alarm indication signal • RDI-P—Path remote defect indication • UNEQ-P—Path unequipped • PLM-P—Path payload (signal) label mismatch • ES-P—Errored seconds (near-end STS path) • SES-P—Severely errored seconds (near-end STS path) • UAS-P—Unavailable seconds (near-end STS path) • ES-PFE—Errored seconds (far-end STS path) • SES-PFE—Severely errored seconds (far-end STS path) • UAS-PFE—Unavailable seconds (far-end STS path)
Received SONET overhead Transmitted SONET overhead	<p>Values of the received and transmitted SONET overhead:</p> <ul style="list-style-type: none"> • C2—Signal label. Allocated to identify the construction and content of the STS-level SPE and for PDI-P. • F1—Section user channel byte. This byte is set aside for the purposes of users. • K1 and K2—These bytes are allocated for APS signaling for the protection of the multiplex section. • J0—Section trace. This byte is defined for STS-1 number 1 of an STS-<i>N</i> signal. Used to transmit a 1-byte fixed-length string or a 16-byte message so that a receiving terminal in a section can verify its continued connection to the intended transmitter. • S1—Synchronization status. The S1 byte is located in the first STS-1 number of an STS-<i>N</i> signal. • Z3 and Z4—Allocated for future use.
Received path trace Transmitted path trace	<p>SONET/SDH interfaces allow path trace bytes to be sent inband across the SONET/SDH link. Juniper Networks and other router manufacturers use these bytes to help diagnose misconfigurations and network errors by setting the transmitted path trace message so that it contains the system hostname and name of the physical interface. The received path trace value is the message received from the router at the other end of the fiber. The transmitted path trace value is the message that this router transmits.</p>
HDLC configuration	<p>Information about the HDLC configuration.</p> <ul style="list-style-type: none"> • Policing bucket—Configured state of the receiving policer. • Shaping bucket—Configured state of the transmitting shaper. • Giant threshold—Giant threshold programmed into the hardware. • Runt threshold—Runt threshold programmed into the hardware.

Table 130: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Packet Forwarding Engine configuration	Information about the configuration of the Packet Forwarding Engine: <ul style="list-style-type: none"> • Destination slot—FPC slot number. • PLP byte—Packet Level Protocol byte.
CoS information	Information about the CoS queue for the physical interface. <ul style="list-style-type: none"> • CoS transmit queue—Queue number and its associated user-configured forwarding class name. • Bandwidth %—Percentage of bandwidth allocated to the queue. • Bandwidth bps—Bandwidth allocated to the queue (in bps). • Buffer %—Percentage of buffer space allocated to the queue. • Buffer usec—Amount of buffer space allocated to the queue, in microseconds. This value is nonzero only if the buffer size is configured in terms of time. • Priority—Queue priority: low or high. • Limit—Displayed if rate limiting is configured for the queue. Possible values are none and exact. If exact is configured, the queue transmits only up to the configured bandwidth, even if excess bandwidth is available. If none is configured, the queue transmits beyond the configured bandwidth if bandwidth is available.
Forwarding classes	Total number of forwarding classes supported on the specified interface.
Egress queues	Total number of egress queues supported on the specified interface.
Queue	Queue number.
Forwarding classes	Forwarding class name.
Queued Packets	Number of packets queued to this queue.
Queued Bytes	Number of bytes queued to this queue. The byte counts vary by PIC type.
Transmitted Packets	Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the Packet Forwarding Engine Chassis Queues field) shows the prefragmentation values.
Transmitted Bytes	Number of bytes transmitted by this queue. The byte counts vary by PIC type.
Tail-dropped packets	Number of packets dropped because of tail drop.

Table 130: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
RED-dropped packets	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> Low, non-TCP—Number of low-loss priority non-TCP packets dropped because of RED. Low, TCP—Number of low-loss priority TCP packets dropped because of RED. High, non-TCP—Number of high-loss priority non-TCP packets dropped because of RED. High, TCP—Number of high-loss priority TCP packets dropped because of RED. (MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> Low—Number of low-loss priority packets dropped because of RED. Medium-low—Number of medium-low loss priority packets dropped because of RED. Medium-high—Number of medium-high loss priority packets dropped because of RED. High—Number of high-loss priority packets dropped because of RED. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
RED-dropped bytes	<p>Number of bytes dropped because of RED. The byte counts vary by PIC type.</p> <ul style="list-style-type: none"> (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> Low, non-TCP—Number of low-loss priority non-TCP bytes dropped because of RED. Low, TCP—Number of low-loss priority TCP bytes dropped because of RED. High, non-TCP—Number of high-loss priority non-TCP bytes dropped because of RED. High, TCP—Number of high-loss priority TCP bytes dropped because of RED. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
Transmit rate	Configured transmit rate of the scheduler. The rate is a percentage of the total interface bandwidth.
Rate Limit	<p>Rate limiting configuration of the queue. Possible values are :</p> <ul style="list-style-type: none"> None—No rate limit. exact—Queue transmits at the configured rate.
Buffer size	Delay buffer size in the queue.
Priority	Scheduling priority configured as low or high .
Excess Priority	Priority of the excess bandwidth traffic on a scheduler: low , medium-low , medium-high , high , or none .

Table 130: show class-of-service interface Output Fields (*continued*)

Field Name	Field Description
Drop profiles	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> • Loss priority—Packet loss priority for drop profile assignment. • Protocol—Transport protocol for drop profile assignment. • Index—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • Name—Name of the drop profile. • Type—Type of the drop profile: discrete or interpolated. • Fill Level—Percentage fullness of a queue. • Drop probability—Drop probability at this fill level.
Excess Priority	Priority of the excess bandwidth traffic on a scheduler.
Drop profiles	<p>Display the assignment of drop profiles.</p> <ul style="list-style-type: none"> • Loss priority—Packet loss priority for drop profile assignment. • Protocol—Transport protocol for drop profile assignment. • Index—Index of the indicated object. Objects that have indexes in this output include schedulers and drop profiles. • Name—Name of the drop profile. • Type—Type of the drop profile: discrete or interpolated. • Fill Level—Percentage fullness of a queue. • Drop probability—Drop probability at this fill level.
Adjustment information	<p>Display the assignment of shaping-rate adjustments on a scheduler node or queue.</p> <ul style="list-style-type: none"> • Adjusting application—Application that is performing the shaping-rate adjustment. <ul style="list-style-type: none"> • The adjusting application can appear as anclp LS-0, which is the Junos OS Access Node Control Profile process (anclpd) that performs shaping-rate adjustments on schedule nodes. • The adjusting application can also appear as pppoe, which adjusts the shaping-rate and overhead-accounting class-of-service attributes on dynamic subscriber interfaces in a broadband access network based on access line parameters in Point-to-Point Protocol over Ethernet (PPPoE) Tags [TR-101]. This feature is supported on MPC/MIC interfaces on MX Series routers. The shaping rate is based on the actual-data-rate-downstream attribute. The overhead accounting value is based on the access-loop-encapsulation attribute and specifies whether the access loop uses Ethernet (frame mode) or ATM (cell mode). • Adjustment type—Type of adjustment: absolute or delta. • Configured shaping rate—Shaping rate configured for the scheduler node or queue. • Adjustment value—Value of adjusted shaping rate. • Adjustment target—Level of shaping-rate adjustment performed: node or queue. • Adjustment overhead-accounting mode—Configured shaping mode: frame or cell.

Sample Output

show class-of-service interface (Physical)

```

user@host> show class-of-service interface so-0/2/3
Physical interface: so-0/2/3, Index: 135
Queues supported: 8, Queues in use: 4

```

```

Total non-default queues created: 4
Scheduler map: <default>, Index: 2032638653

Logical interface: fe-0/0/1.0, Index: 68, Dedicated Queues: no
Shaping rate: 32000

```

Object	Name	Type	Index
Scheduler-map	<default>		27
Rewrite	exp-default	exp	21
Classifier	exp-default	exp	5
Classifier	ipprec-compatibility	ip	8
Forwarding-class-map	exp-default	exp	5

show class-of-service interface (Logical)

```

user@host> show class-of-service interface so-0/2/3.0
Logical interface: so-0/2/3.0, Index: 68, Dedicated Queues: no
Shaping rate: 32000

```

Object	Name	Type	Index
Scheduler-map	<default>		27
Rewrite	exp-default	exp	21
Classifier	exp-default	exp	5
Classifier	ipprec-compatibility	ip	8
Forwarding-class-map	exp-default	exp	5

show class-of-service interface (Gigabit Ethernet)

```

user@host> show class-of-service interface ge-6/2/0
Physical interface: ge-6/2/0, Index: 175
Queues supported: 4, Queues in use: 4
Scheduler map: <default>, Index: 2
Input scheduler map: <default>, Index: 3
Chassis scheduler map: <default-chassis>, Index: 4

```

show class-of-service interface (PPPoE Interface)

```

user@host> show class-of-service interface pp0.1
Logical interface: pp0.1, Index: 85

```

Object	Name	Type	Index
Traffic-control-profile	tcp-pppoe.o.pp0.1	Output	2726446535
Classifier	ipprec-compatibility	ip	13

```

Adjusting application: PPPoE
Adjustment type: absolute
Adjustment value: 5000000
Adjustment overhead-accounting mode: cell
Adjustment target: node

```

show class-of-service interface (T4000 Routers with Type 5 FPCs)

```

user@host> show class-of-service interface xe-4/0/0
Physical interface: xe-4/0/0, Index: 153
Queues supported: 8, Queues in use: 4
Shaping rate: 5000000000 bps
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled

Logical interface: xe-4/0/0.0, Index: 77

```

Index	Object	Name	Type
13	Classifier	ipprec-compatibility	ip

show class-of-service interface detail

```
user@host> show class-of-service interface ge-0/3/0 detail
```

```
Physical interface: ge-0/3/0, Enabled, Physical link is Up
Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled, Auto-negotiation: Enabled,
Remote fault: Online
```

```
Device flags : Present Running
Interface flags: SNMP-Traps Internal: 0x4000
```

```
Physical interface: ge-0/3/0, Index: 138
Queues supported: 4, Queues in use: 5
Shaping rate: 50000 bps
Scheduler map: interface-scheduler-map, Index: 58414
Input shaping rate: 10000 bps
878674 Input scheduler map: scheduler-map, Index: 15103
Chassis scheduler map: <default-chassis>, Index: 4
Congestion-notification: Disabled
```

```
Logical interface ge-0/3/0.0
```

```
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
```

```
inet
```

```
mpls
```

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/3/0.0	up	up	inet		

Interface	Admin	Link	Proto	Input Policer	Output Policer
ge-0/3/0.0	up	up	inet		

```
Logical interface: ge-0/3/0.0, Index: 68
```

Object	Name	Type	Index
Rewrite	exp-default	exp (mpls-any)	33
Classifier	exp-default	exp	10
Classifier	ipprec-compatibility	ip	13

```
Logical interface ge-0/3/0.1
```

```
Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.2 ] Encapsulation: ENET2
```

```
inet
```

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/3/0.1	up	up	inet		

Interface	Admin	Link	Proto	Input Policer	Output Policer
ge-0/3/0.1	up	up	inet		

```
Logical interface: ge-0/3/0.1, Index: 69
```

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

show class-of-service interface comprehensive

```
user@host> show class-of-service interface ge-0/3/0 comprehensive
```

```
Physical interface: ge-0/3/0, Enabled, Physical link is Up
```

```
Interface index: 138, SNMP ifIndex: 601, Generation: 141
```

```
Link-level type: Ethernet, MTU: 1518, Speed: 1000mbps, BPDU Error: None,
MAC-REWRITE Error: None, Loopback: Disabled, Source filtering: Disabled, Flow
control: Enabled,
```

```
Auto-negotiation: Enabled, Remote fault: Online
```

```
Device flags : Present Running
```

```
Interface flags: SNMP-Traps Internal: 0x4000
```

```

CoS queues      : 4 supported, 4 maximum usable queues
Schedulers     : 256
Hold-times      : Up 0 ms, Down 0 ms
Current address: 00:14:f6:f4:b4:5d, Hardware address: 00:14:f6:f4:b4:5d
Last flapped    : 2010-09-07 06:35:22 PDT (15:14:42 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps
IPv6 total statistics:
Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Ingress traffic statistics at Packet Forwarding Engine:
Input bytes : 0 0 bps
Input packets: 0 0 pps
Drop bytes : 0 0 bps
Drop packets: 0 0 pps
Label-switched interface (LSI) traffic statistics:
Input bytes : 0 0 bps
Input packets: 0 0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Policed discards: 0, L3
incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0,
Resource errors: 0
Output errors:
Carrier transitions: 5, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0,
FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0
Ingress queues: 4 supported, 5 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 af3                0                0                0
1 af2                0                0                0
2 ef2                0                0                0
3 ef1                0                0                0

Egress queues: 4 supported, 5 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

0 af3                0                0                0
1 af2                0                0                0
2 ef2                0                0                0
3 ef1                0                0                0

Active alarms : None
Active defects : None
MAC statistics:
Total octets      Receive      Transmit
Total packets     0            0
Unicast packets   0            0
Broadcast packets 0            0
Multicast packets 0            0

```

```

CRC/Align errors                0                0
FIFO errors                     0                0
MAC control frames              0                0
MAC pause frames                0                0
Oversized frames                0
Jabber frames                   0
Fragment frames                 0
VLAN tagged frames              0
Code violations                  0
Filter statistics:
  Input packet count            0
  Input packet rejects          0
  Input DA rejects              0
  Input SA rejects              0
  Output packet count           0
  Output packet pad count       0
  Output packet error count     0
  CAM destination filters: 0, CAM source filters: 0
Autonegotiation information:
  Negotiation status: Complete
  Link partner:
    Link mode: Full-duplex, Flow control: Symmetric/Asymmetric, Remote fault:
OK
  Local resolution:
    Flow control: Symmetric, Remote fault: Link OK
Packet Forwarding Engine configuration:
  Destination slot: 0
CoS information:
  Direction : Output
  CoS transmit queue            Bandwidth          Buffer Priority
Limit                           %          bps      %          usec          high
  2 ef2                         39          19500    0          120
none
  Direction : Input
  CoS transmit queue            Bandwidth          Buffer Priority
Limit                           %          bps      %          usec          low
  0 af3                         30          3000     45          0
none

Physical interface: ge-0/3/0, Enabled, Physical link is Up
Interface index: 138, SNMP ifIndex: 601
Forwarding classes: 16 supported, 5 in use
Ingress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
  Transmitted:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :          0          0 pps
    RED-dropped bytes   :          0          0 bps
Queue: 1, Forwarding classes: af2
  Queued:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
  Transmitted:
    Packets      :          0          0 pps

```

```

Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef2
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: ef1
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Forwarding classes: 16 supported, 5 in use
Egress queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: af2
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: ef2
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RL-dropped packets : 0 0 pps
RL-dropped bytes : 0 0 bps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 pps

```

```

    RED-dropped bytes      :                0          0 bps
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets : Not Available
    RL-dropped packets    :                0          0 pps
    RL-dropped bytes      :                0          0 bps
    RED-dropped packets   :                0          0 pps
    RED-dropped bytes     :                0          0 bps

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 5 in use
Queue: 0, Forwarding classes: af3
  Queued:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets :                0          0 pps
    RED-dropped packets   : Not Available
    RED-dropped bytes     : Not Available
Queue: 1, Forwarding classes: af2
  Queued:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets :                0          0 pps
    RED-dropped packets   : Not Available
    RED-dropped bytes     : Not Available
Queue: 2, Forwarding classes: ef2
  Queued:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
  Transmitted:
    Packets                :                0          0 pps
    Bytes                  :                0          0 bps
    Tail-dropped packets :                0          0 pps
    RED-dropped packets   : Not Available
    RED-dropped bytes     : Not Available
Queue: 3, Forwarding classes: ef1
  Queued:
    Packets                :            108546          0 pps
    Bytes                  :       12754752       376 bps
  Transmitted:
    Packets                :            108546          0 pps
    Bytes                  :       12754752       376 bps
    Tail-dropped packets :                0          0 pps
    RED-dropped packets   : Not Available
    RED-dropped bytes     : Not Available

Physical interface: ge-0/3/0, Index: 138
Queues supported: 4, Queues in use: 5
Shaping rate: 50000 bps

```


Scheduler map: interface-scheduler-map, Index: 58414

Scheduler: ef2, Forwarding class: ef2, Index: 39155

Transmit rate: 39 percent, Rate Limit: none, Buffer size: 120 us, Buffer

Limit: none, Priority: high

Excess Priority: unspecified

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	1	< default-drop-profile>
Medium low	any	1	< default-drop-profile>
Medium high	any	1	< default-drop-profile>
High	any	1	< default-drop-profile>

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Input shaping rate: 10000 bps

Input scheduler map: scheduler-map

Scheduler map: scheduler-map, Index: 15103

Scheduler: af3, Forwarding class: af3, Index: 35058

Transmit rate: 30 percent, Rate Limit: none, Buffer size: 45 percent, Buffer

Limit: none, Priority: low

Excess Priority: unspecified

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	40582	green
Medium low	any	1	< default-drop-profile>
Medium high	any	1	< default-drop-profile>
High	any	18928	yellow

Drop profile: green, Type: discrete, Index: 40582

Fill level	Drop probability
50	0
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: yellow, Type: discrete, Index: 18928

Fill level	Drop probability
50	0
100	100

Chassis scheduler map: < default-drop-profile>

Scheduler map: < default-drop-profile>, Index: 4

Scheduler: < default-drop-profile>, Forwarding class: af3, Index: 25

Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer

Limit: none, Priority: low

Excess Priority: low

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	1	< default-drop-profile>
Medium low	any	1	< default-drop-profile>
Medium high	any	1	< default-drop-profile>
High	any	1	< default-drop-profile>

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Scheduler: < default-drop-profile>, Forwarding class: af2, Index: 25
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none, Priority: low
 Excess Priority: low
 Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	1	< default-drop-profile>
Medium low	any	1	< default-drop-profile>
Medium high	any	1	< default-drop-profile>
High	any	1	< default-drop-profile>

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Scheduler: < default-drop-profile>, Forwarding class: ef2, Index: 25
 Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer Limit: none, Priority: low
 Excess Priority: low
 Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	1	< default-drop-profile>
Medium low	any	1	< default-drop-profile>
Medium high	any	1	< default-drop-profile>
High	any	1	< default-drop-profile>

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

Fill level	Drop probability
100	100

Drop profile: < default-drop-profile>, Type: discrete, Index: 1

```

Fill level      Drop probability
    100          100

Scheduler: < default-drop-profile>, Forwarding class: ef1, Index: 25
  Transmit rate: 25 percent, Rate Limit: none, Buffer size: 25 percent, Buffer
Limit: none, Priority: low
  Excess Priority: low
  Drop profiles:
    Loss priority  Protocol    Index    Name
    Low           any         1        < default-drop-profile>
    Medium low    any         1        < default-drop-profile>
    Medium high   any         1        < default-drop-profile>
    High          any         1        < default-drop-profile>
Drop profile: , Type: discrete, Index: 1
  Fill level      Drop probability
    100          100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level      Drop probability
    100          100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level      Drop probability
    100          100
Drop profile: < default-drop-profile>, Type: discrete, Index: 1
  Fill level      Drop probability
    100          100
  Congestion-notification: Disabled
Forwarding class
priority Policing priority
af3      normal
af2      normal
ef2      normal
ef1      normal
af1      normal
ID        Queue  Restricted queue  Fabric
0         0       0                 low
1         1       1                 low
2         2       2                 high
3         3       3                 high
4         4       0                 low

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152) (Generation 159)
  Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1 ] Encapsulation: ENET2
  Traffic statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Local statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Transit statistics:
    Input bytes : 0
    Output bytes : 0
    Input packets: 0
    Output packets: 0
  Protocol inet, MTU: 1500, Generation: 172, Route table: 0
    Flags: Sendbcst-pkt-to-re
    Input Filters: filter-in-ge-0/3/0.0-i,
    Policar: Input: p1-ge-0/3/0.0-inet-i
  Protocol mpls, MTU: 1488, Maximum labels: 3, Generation: 173, Route table: 0

```

Flags: Is-Primary
Output Filters: exp-filter,,,,,

Logical interface ge-0/3/0.0 (Index 68) (SNMP ifIndex 152)
Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.1] Encapsulation: ENET2
Input packets : 0
Output packets: 0

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/3/0.0	up	up	inet	filter-in-ge-0/3/0.0-i	
			mpls		exp-filter
Interface	Admin	Link	Proto	Input Policer	Output Policer
ge-0/3/0.0	up	up	inet	p1-ge-0/3/0.0-inet-i	
			mpls		

Filter: filter-in-ge-0/3/0.0-i

Counters:

Name	Bytes	Packets
count-filter-in-ge-0/3/0.0-i	0	0

Filter: exp-filter

Counters:

Name	Bytes	Packets
count-exp-seven-match	0	0
count-exp-zero-match	0	0

Policers:

Name	Packets
p1-ge-0/3/0.0-inet-i	0

Logical interface: ge-0/3/0.0, Index: 68

Object	Name	Type	Index
Rewrite	exp-default	exp (mpls-any)	33

Rewrite rule: exp-default, Code point type: exp, Index: 33

Forwarding class	Loss priority	Code point
af3	low	000
af3	high	001
af2	low	010
af2	high	011
ef2	low	100
ef2	high	101
ef1	low	110
ef1	high	111

Object	Name	Type	Index
Classifier	exp-default	exp	10

Classifier: exp-default, Code point type: exp, Index: 10

Code point	Forwarding class	Loss priority
000	af3	low
001	af3	high
010	af2	low
011	af2	high
100	ef2	low
101	ef2	high
110	ef1	low
111	ef1	high

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13

Code point	Forwarding class	Loss priority		
000	af3	low		
001	af3	high		
010	af3	low		
011	af3	high		
100	af3	low		
101	af3	high		
110	ef1	low		
111	ef1	high		
Forwarding class	ID	Queue	Restricted queue	Fabric
priority Policing priority				
af3	0	0	0	low
normal				
af2	1	1	1	low
normal				
ef2	2	2	2	high
normal				
ef1	3	3	3	high
normal				
af1	4	4	0	low
normal				

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154) (Generation 160)

Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.2] Encapsulation: ENET2

Traffic statistics:

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Local statistics:

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0

Transit statistics:

Input bytes : 0 0 bps
Output bytes : 0 0 bps
Input packets: 0 0 pps
Output packets: 0 0 pps

Protocol inet, MTU: 1500, Generation: 174, Route table: 0

Flags: Sendbcst-pkt-to-re

Logical interface ge-0/3/0.1 (Index 69) (SNMP ifIndex 154)

Flags: SNMP-Traps 0x4000 VLAN-Tag [0x8100.2] Encapsulation: ENET2

Input packets : 0

Output packets: 0

Interface	Admin	Link	Proto	Input Filter	Output Filter
ge-0/3/0.1	up	up	mpls		
Interface	Admin	Link	Proto	Input Policier	Output Policier
ge-0/3/0.1	up	up	mpls		

Logical interface: ge-0/3/0.1, Index: 69

Object	Name	Type	Index
--------	------	------	-------

Classifier		ipprec-compatibility		ip		13
Classifier: ipprec-compatibility, Code point type: inet-precedence, Index: 13						
Code point	Forwarding class		Loss priority			
000	af3		low			
001	af3		high			
010	af3		low			
011	af3		high			
100	af3		low			
101	af3		high			
110	ef1		low			
111	ef1		high			
Forwarding class		ID	Queue	Restricted queue	Fabric	
priority	Policing priority					
af3		0	0	0	low	
	normal					
af2		1	1	1	low	
	normal					
ef2		2	2	2	high	
	normal					
ef1		3	3	3	high	
	normal					
af1		4	4	0	low	
	normal					

show class-of-service interface (ACX Series Routers)

```
user@host-g11# show class-of-service interface
Physical interface: at-0/0/0, Index: 130
Queues supported: 4, Queues in use: 4
Scheduler map: <default>, Index: 2
Congestion-notification: Disabled
```

```
Logical interface: at-0/0/0.0, Index: 69
```

```
Logical interface: at-0/0/0.32767, Index: 70
```

```
Physical interface: at-0/0/1, Index: 133
```

```
Queues supported: 4, Queues in use: 4
```

```
Scheduler map: <default>, Index: 2
```

```
Congestion-notification: Disabled
```

```
Logical interface: at-0/0/1.0, Index: 71
```

```
Logical interface: at-0/0/1.32767, Index: 72
```

```
Physical interface: ge-0/1/0, Index: 146
```

```
Queues supported: 8, Queues in use: 5
```

```
Scheduler map: <default>, Index: 2
```

```
Congestion-notification: Disabled
```

Object	Name	Type	Index
Rewrite	dscp-default	dscp	31
Classifier	d1	dscp	11331
Classifier	ci	ieee8021p	583

```
Logical interface: ge-0/1/0.0, Index: 73
```

Object	Name	Type	Index
Rewrite	custom-exp	exp (mpls-any)	46413

Logical interface: ge-0/1/0.1, Index: 74

Logical interface: ge-0/1/0.32767, Index: 75

Physical interface: ge-0/1/1, Index: 147

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Logical interface: ge-0/1/1.0, Index: 76

Physical interface: ge-0/1/2, Index: 148

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Rewrite	ri	ieee8021p (outer)	35392
Classifier	ci	ieee8021p	583

Physical interface: ge-0/1/3, Index: 149

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Logical interface: ge-0/1/3.0, Index: 77

Object	Name	Type	Index
Rewrite	custom-exp2	exp (mpls-any)	53581

Physical interface: ge-0/1/4, Index: 150

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Physical interface: ge-0/1/5, Index: 151

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Physical interface: ge-0/1/6, Index: 152

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Physical interface: ge-0/1/7, Index: 153

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	d1	dscp	11331

Physical interface: ge-0/2/0, Index: 154

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Physical interface: ge-0/2/1, Index: 155

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Logical interface: ge-0/2/1.0, Index: 78

Logical interface: ge-0/2/1.32767, Index: 79

Physical interface: xe-0/3/0, Index: 156

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Logical interface: xe-0/3/0.0, Index: 80

Physical interface: xe-0/3/1, Index: 157

Queues supported: 8, Queues in use: 5

Scheduler map: <default>, Index: 2

Congestion-notification: Disabled

Object	Name	Type	Index
Classifier	ipprec-compatibility	ip	13

Logical interface: xe-0/3/1.0, Index: 81

[edit]

user@host-g11#

show class-of-service multi-destination

Syntax	show class-of-service multi-destination
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	For each class-of-service (CoS) multideestination classifier, display the classifier type.
Options	none —Display all multideestination classifiers.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show class-of-service classifier on page 558 • Understanding Applying CoS Classifiers and Rewrite Rules to Interfaces on page 57 • Understanding CoS Fabric Forwarding Class Sets on page 186
Output Fields	Table 131 on page 613 describes the output fields for the show class-of-service multi-destination command. Output fields are listed in the approximate order in which they appear.

Table 131: show class-of-service multi-destination Output Fields

Field Name	Field Description
Family ethernet	Family to which the classifier belongs.
Classifier Name	Name of the classifier.
Classifier Type	Type of the classifier: dscp or ieee-802.1 .
Classifier Index	Internal index of the classifier.

Sample Output

show class-of-service multi-destination

```
user@switch> show class-of-service multi-destination
```

```
Family ethernet:
Classifier Name      Classifier Type      Classifier Index
ba-mcast-classifier  ieee-802.1         62376
```

show class-of-service rewrite-rule

Syntax	show class-of-service rewrite-rule <name <i>name</i> > <type <i>type</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the mapping of forwarding classes and loss priority to code point values.
Options	<p>none—Display all rewrite rules.</p> <p>name <i>name</i>—(Optional) Display the specified rewrite rule.</p> <p>type <i>type</i>—(Optional) Display the rewrite rule of the specified type. The rewrite rule type can be one of the following:</p> <ul style="list-style-type: none"> dscp—For IPv4 traffic. dscp-ipv6—For IPv6 traffic. exp—For MPLS traffic. frame-relay-de—(J Series routers only) For Frame Relay traffic. ieee-802.1—For Layer 2 traffic. inet-precedence—For IPv4 traffic.
Required Privilege Level	view
List of Sample Output	show class-of-service rewrite-rule type dscp on page 615 show class-of-service rewrite-rule type dscp (QFX Series) on page 615
Output Fields	Table 132 on page 614 describes the output fields for the show class-of-service rewrite-rule command. Output fields are listed in the approximate order in which they appear.

Table 132: show class-of-service rewrite-rule Output Fields

Field Name	Field Description
Rewrite rule	Name of the rewrite rule.
Code point type	Type of rewrite rule: dscp , dscp-ipv6 , exp , frame-relay-de , or inet-precedence .
Forwarding class	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router or switch.
Index	Internal index for this particular rewrite rule.
Loss priority	Loss priority for rewriting.

Table 132: show class-of-service rewrite-rule Output Fields (*continued*)

Field Name	Field Description
Code point	Code point value to rewrite.

Sample Output

show class-of-service rewrite-rule type dscp

```

user@host> show class-of-service rewrite-rule type dscp
Rewrite rule: dscp-default, Code point type: dscp
  Forwarding class      Loss priority      Code point
  gold                  high               000000
  silver                low                110000
  silver                high               111000
  bronze                low                001010
  bronze                high               001100
  lead                  high               101110

Rewrite rule: abc-dscp-rewrite, Code point type: dscp, Index: 3245
  Forwarding class      Loss priority      Code point
  gold                  low                000111
  gold                  high               001010
  silver                low                110000
  silver                high               111000
  bronze                high               001100
  lead                  low                101110
  lead                  high               110111

```

Sample Output

show class-of-service rewrite-rule type dscp (QFX Series)

```

user@host> show class-of-service rewrite-rule type dscp
Rewrite rule: dscp-default, Code point type: dscp, Index: 31
  Forwarding class      Loss priority      Code point
  best-effort            low                000000
  best-effort            high               000000
  fcoe                   low                101110
  fcoe                   high               101110
  no-loss                low                001010
  no-loss                high               001100
  network-control        low                110000
  network-control        high               111000

```

show class-of-service scheduler-map

Syntax	<code>show class-of-service scheduler-map</code> <code><name></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display the mapping of schedulers to forwarding classes and a summary of scheduler parameters for each entry.
Options	none —Display all scheduler maps. name —(Optional) Display a summary of scheduler parameters for each forwarding class to which the named scheduler is assigned.
Required Privilege Level	view
List of Sample Output	show class-of-service scheduler-map on page 617
Output Fields	Table 133 on page 616 describes the output fields for the show class-of-service scheduler-map command. Output fields are listed in the approximate order in which they appear.

Table 133: show class-of-service scheduler-map Output Fields

Field Name	Field Description
Scheduler map	Name of the scheduler map.
Index	Index of the indicated object. Objects having indexes in this output include scheduler maps, schedulers, and drop profiles.
Scheduler	Name of the scheduler.
Forwarding class	Classification of a packet affecting the forwarding, scheduling, and marking policies applied as the packet transits the router.
Transmit rate	Configured transmit rate of the scheduler (in bps). The rate is a percentage of the total interface bandwidth, or the keyword remainder , which indicates that the scheduler receives the remaining bandwidth of the interface.
Rate Limit	Rate limiting configuration of the queue. Possible values are none , meaning no rate limiting, and exact , meaning the queue only transmits at the configured rate.
Maximum buffer delay	Amount of transmit delay (in milliseconds) or the buffer size of the queue. The buffer size is shown as a percentage of the total interface buffer allocation, or by the keyword remainder to indicate that the buffer is sized according to what remains after other scheduler buffer allocations.
Priority	Scheduling priority: low or high .

Table 133: show class-of-service scheduler-map Output Fields (*continued*)

Field Name	Field Description
Excess priority	Priority of excess bandwidth: low , medium-low , medium-high , high , or none .
Explicit Congestion Notification	(QFX Series only) Explicit congestion notification (ECN) state: <ul style="list-style-type: none"> Disable—ECN is disabled on the specified scheduler Enable—ECN is enabled on the specified scheduler ECN is disabled by default.
Adjust minimum	Minimum shaping rate for an adjusted queue, in bps.
Adjust percent	Bandwidth adjustment applied to a queue, in percent.
Drop profiles	Table displaying the assignment of drop profiles by name and index to a given loss priority and protocol pair.
Loss priority	Packet loss priority for drop profile assignment.
Protocol	Transport protocol for drop profile assignment.
Name	Name of the drop profile.

Sample Output

show class-of-service scheduler-map

```

user@host> show class-of-service scheduler-map
Scheduler map: dd-scheduler-map, Index: 84

Scheduler: aa-scheduler, Index: 8721, Forwarding class: aa-forwarding-class
Transmit rate: 30 percent, Rate Limit: none, Maximum buffer delay: 39 ms,
Priority: high
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           non-TCP   8724   aa-drop-profile
  Low           TCP       9874   bb-drop-profile
  High          non-TCP   8833   cc-drop-profile
  High          TCP       8484   dd-drop-profile

Scheduler: bb-scheduler, Forwarding class: aa-forwarding-class
Transmit rate: 40 percent, Rate limit: none, Maximum buffer delay: 68 ms,
Priority: high
Drop profiles:
  Loss priority  Protocol  Index  Name
  Low           non-TCP   8724   aa-drop-profile
  Low           TCP       9874   bb-drop-profile
  High          non-TCP   8833   cc-drop-profile
  High          TCP       8484   dd-drop-profile

```

show class-of-service shared-buffer

Syntax	show class-of-service shared-buffer <egress ingress>
Release Information	Command introduced in Junos OS Release 12.3 for the QFX Series.
Description	Display the shared buffer allocation and partitioning configuration.
Options	<p>none—Display ingress and egress shared buffer settings.</p> <p>egress—(Optional) Display the egress shared buffer settings.</p> <p>ingress—(Optional) Display the ingress shared buffer settings.</p>
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Best-Effort Unicast Traffic on page 365 • Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Multicast Traffic on page 377 • Example: Recommended Configuration of the Shared Buffer Pool for Networks with Mostly Lossless Traffic on page 383 • Configuring Global Ingress and Egress Shared Buffers on page 425 • Understanding CoS Buffer Configuration on page 131
List of Sample Output	show class-of-service shared-buffer on page 619
Output Fields	Table 134 on page 618 describes the output fields for the show class-of-service shared-buffer command. Output fields are listed in the approximate order in which they appear.

Table 134: show class-of-service shared-buffer Output Fields

Field Name	Field Description
Ingress	Ingress shared buffer configuration.
Total Buffer	Total buffer space available to the ports in KB. This is the combined dedicated buffer pool and shared buffer pool.
Dedicated Buffer	Buffer space allocated to the dedicated buffer pool in KB.
Shared Buffer	Buffer space allocated to the shared buffer pool in KB.
Lossless	Buffer space allocated to the lossless traffic buffer pool in KB.

Table 134: show class-of-service shared-buffer Output Fields (*continued*)

Field Name	Field Description
Lossless Headroom	Buffer space allocated to the lossless headroom traffic buffer pool to support priority-based flow control (PFC) and Ethernet PAUSE in KB. (Ingress ports only.)
Lossy	Buffer space allocated to the lossy (best-effort) traffic buffer pool in KB.
Lossless Headroom Utilization	Utilization of the ingress lossless headroom buffer pool. (These fields can help you to determine how much headroom buffer space you need to reserve to support PFC and Ethernet PAUSE for lossless flows.)
Node Device	Index number that identifies the switch. On a QFX3500 switch, this field always has a value of zero (0).
Total	Size of the lossless headroom ingress buffer pool in KB.
Used	Amount in KB of lossless headroom ingress buffer used.
Free	Amount in KB of lossless headroom ingress buffer free (unused).
Egress	Egress shared buffer configuration.
Multicast	Buffer space allocated to the multicast traffic buffer pool in KB. (Egress ports only.)

Sample Output

show class-of-service shared-buffer

```
user@switch> show class-of-service shared-buffer
```

```
Ingress:
```

```
Total Buffer      : 9360.00 KB
Dedicated Buffer   : 2158.00 KB
Shared Buffer      : 7202.00 KB
  Lossless        : 648.18 KB
  Lossless Headroom : 3240.90 KB
  Lossy           : 3312.92 KB
```

```
Lossless Headroom Utilization:
```

```
Node Device      Total      Used      Free
0                3240.90 KB  0.00 KB  3240.90 KB
```

```
Egress:
```

```
Total Buffer      : 9360.00 KB
Dedicated Buffer   : 2704.00 KB
Shared Buffer      : 6656.00 KB
  Lossless        : 3328.00 KB
  Multicast       : 1264.64 KB
  Lossy           : 2063.36 KB
```

show class-of-service traffic-control-profile

Syntax	<code>show class-of-service traffic-control-profile</code> <code><profile-name></code>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 12.2 for ACX Series Routers.
Description	For Gigabit Ethernet IQ PICs, Channelized IQ PICs, EQ DPCs, and Trio MPC/MIC interfaces only, display traffic shaping and scheduling profiles. (ACX Series routers) For ATM IMA pseudowire interfaces, display traffic shaping and scheduling profiles.
Options	none —Display all profiles. profile-name —(Optional) Display information about a single profile.
Required Privilege Level	view
List of Sample Output	show class-of-service traffic-control-profile on page 622 show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC) on page 622 show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces) on page 622
Output Fields	Table 135 on page 620 describes the output fields for the show class-of-service traffic-control-profile command. Output fields are listed in the approximate order in which they appear.

Table 135: show class-of-service traffic-control-profile Output Fields

Field Name	Field Description
Traffic control profile	Name of the traffic control profile.
Index	Index number of the traffic control profile.
ATM Service	(MX Series routers with ATM Multi-Rate CE MIC) Configured category of ATM service. Possible values: <ul style="list-style-type: none"> cbr—Constant bit rate. rtvbr—Real time variable bit rate. nrtvbr—Non real time variable bit rate. ubr—Unspecified bit rate.
Maximum Burst Size	Configured maximum burst size, in cells.
Peak rate	Configured peak rate, in cps.

Table 135: show class-of-service traffic-control-profile Output Fields (*continued*)

Field Name	Field Description
Sustained rate	Configured sustained rate, in cps.
Shaping rate	Configured shaping rate, in bps. NOTE: (MX Series routers with ATM Multi-Rate CE MIC) Configured peak rate, in cps.
Shaping rate burst	Configured burst size for the shaping rate, in bytes. NOTE: (MX Series routers with ATM Multi-Rate CE MIC) Configured maximum burst rate, in cells.
Shaping rate priority high	Configured shaping rate for high-priority traffic, in bps.
Shaping rate priority medium	Configured shaping rate for medium-priority traffic, in bps.
Shaping rate priority low	Configured shaping rate for low-priority traffic, in bps.
Shaping rate excess high	Configured shaping rate for high-priority excess traffic, in bps.
Shaping rate excess low	Configured shaping rate for low-priority excess traffic, in bps.
Scheduler map	Name of the associated scheduler map.
Delay Buffer rate	Configured delay buffer rate, in bps.
Excess rate	Configured excess rate, in percent or proportion.
Excess rate high	Configured excess rate for high priority traffic, in percent or proportion.
Excess rate low	Configured excess rate for low priority traffic, in percent or proportion.
Guaranteed rate	Configured guaranteed rate, in bps or cps. NOTE: (MX Series routers with ATM Multi-Rate CE MIC) This value depends on the ATM service category chosen. Possible values: <ul style="list-style-type: none"> • cbr—Guaranteed rate is equal to the configured peak rate in cps. • rtvbr—Guaranteed rate is equal to the configured sustained rate in cps. • nrtvbr—Guaranteed rate is equal to the configured sustained rate in cps.
Guaranteed rate burst	Configured burst size for the guaranteed rate, in bytes.
adjust-minimum	Configured minimum shaping rate for an adjusted queue, in bps.

Table 135: show class-of-service traffic-control-profile Output Fields (*continued*)

Field Name	Field Description
overhead accounting mode	Configured shaping mode: Frame Mode or Cell Mode .
Overhead bytes	Configured byte adjustment value.

Sample Output

show class-of-service traffic-control-profile

```

user@host> show class-of-service traffic-control-profile
Traffic control profile: Profile1, Index: 57625
  Scheduler map: m1
  Delay Buffer rate: 500000
  Guaranteed rate: 1000000

Traffic control profile: Profile2, Index: 57624
  Scheduler map: m2
  Delay Buffer rate: 600000
  Guaranteed rate: 2000000

Traffic control profile: Profile3, Index: 57627
  Scheduler map: m3
  Delay Buffer rate: 800000
  Guaranteed rate: 3000000
  .Excess rate high: proportion 4

Traffic control profile: Profile4, Index: 57626
  Scheduler map: m4
  Delay Buffer rate: 750000
  Guaranteed rate: 4000000
  ..adjust-minimum 20000000

```

show class-of-service traffic-control-profile (MX Series routers with Clear Channel Multi-Rate CE MIC)

```

user@host> show class-of-service traffic-control-profile
Traffic control profile: at-vbr1, Index: 11395
  ATM Service: RTVBR
  Scheduler map: m3
  overhead accounting mode: Frame Mode
  Shaping rate: 1000 cps
  Shaping rate burst: 500 cells
  Delay Buffer rate: 2000 cps
  Guaranteed rate: 1000 cps

Traffic control profile: foo, Index: 38286
  ATM Service: UBR
  Scheduler map: m3
  overhead accounting mode: Frame Mode

```

show class-of-service traffic-control-profile (ACX Series routers with ATM IMA pseudowire interfaces)

```

user@host> show class-of-service traffic-control-profile
Traffic control profile: foo, Index: 38286
  ATM Service: RTVBR
  Shaping rate: 2000 cps

```

```
Shaping rate burst: 200 cells  
Scheduler map: <default>  
Delay Buffer rate: 1000 cps  
Guaranteed rate: 1700 cps
```

show dcbx

Syntax	show dcbx
Release Information	Command introduced in Junos OS Release 11.3 for the QFX Series.
Description	List DCBX status (enabled or disabled) and the interfaces on which DCBX is enabled.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • show dcbx neighbors on page 625 • Configuring DCBX Autonegotiation on page 431
Output Fields	Table 136 on page 624 lists the output fields for the show dcbx command. Output fields are listed in the approximate order in which they appear.

Table 136: show dcbx output fields

Field Name	Field Description
DCBX	Status of DCBX on the switch or for the specified interface: <ul style="list-style-type: none"> • Enabled—DCBX is enabled on the switch or on the specified interface • Disabled—DCBX is disabled on the switch or on the specified interface
Interface	Name of the interface

Sample Output

show dcbx

```

user@switch> show dcbx
DCBX                : Enabled
Interface           DCBX
xe-0/0/9.0          enabled
xe-0/0/32.0         enabled
xe-0/0/36.0         enabled

```

show dcbx neighbors

Syntax	show dcbx neighbors <interface interface-name> <terse>
Release Information	Command introduced in Junos OS Release 11.1 for the QFX Series. Command introduced in Junos OS Release 11.3 for EX Series switches.
Description	Display information about Data Center Bridging Capability Exchange protocol (DCBX) neighbor interfaces.
Options	none —Display information about all DCBX neighbor interfaces. interface-name —(Optional) Display information for the specified interface. terse —Display the specified level of output.
Required Privilege Level	view
Related Documentation	<ul style="list-style-type: none"> • Configuring DCBX Autonegotiation on page 431 • Example: Configuring DCBX Application Protocol TLV Exchange on page 389 • Example: Configuring an FCoE Transit Switch • Example: Configuring DCBX to Support an iSCSI Application • Understanding DCB Features and Requirements on page 156 • Understanding Data Center Bridging Capability Exchange Protocol for EX Series Switches • dcbx on page 471
List of Sample Output	show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode) on page 638 show dcbx neighbors interface (QFX Series, IEEE DCBX Mode) on page 640 show dcbx neighbors terse (QFX Series) on page 642 show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly) on page 642 show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application) on page 643 show dcbx neighbors (EX4500 Switch: Includes ETS) on page 644
Output Fields	Table 137 on page 625 lists the output fields for the show dcbx neighbors command. Output fields are listed in the approximate order in which they appear.

Table 137: show dcbx neighbors Output Fields

Field Name	Field Description
Interface	Name of the interface.

Table 137: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Parent Interface	Name of the link aggregation group (LAG) interface to which the DCBX interface belongs.
Active-application-map	Name of the application map applied to the interface.
Protocol-Mode	<p>(QFX Series) DCBX protocol mode the interface uses:</p> <ul style="list-style-type: none"> IEEE DCBX Version—The interface uses IEEE DCBX mode. DCBX Version 1.01—The interface uses DCBX version 1.01. <p>NOTE: On interfaces that use the IEEE DCBX mode, the show dcbx neighbors interface <i>interface-name</i> operational command does not include application, PFC, or ETS operational state in the output.</p>
Protocol-State	<p>(DCBX Version 1.01 only) DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> in-sync—The local interface received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface. ack-pending—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a state change message sent by the local interface.
Local-Advertisement	<p>(DCBX Version 1.01 only)</p> <p>Status of advertisements that the local interface sends to the peer.</p>
Operational version	Version of the DCBX standard used.
sequence-number	<p>Number of state change messages sent to the peer.</p> <p>If the interface Protocol-State value is in-sync, this number should match the acknowledge-id number in the Peer-Advertisement section.</p> <p>If the interface Protocol-State value is ack-pending, this number does not match the acknowledge-id number in the Peer-Advertisement section.</p>
acknowledge-id	<p>Number of acknowledge messages received from the peer.</p> <p>If the Protocol-State value is in-sync, this number should match the sequence-number value in the Peer-Advertisement section.</p> <p>If the Protocol-State value is ack-pending, this number does not match the sequence-number value in the Peer-Advertisement section.</p>

Table 137: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Peer-Advertisement	(DCBX Version 1.01 only) Status of advertisements that the peer sends to the local interface.
Operational version	Version of the DCBX standard used.
sequence-number	<p>Number of state change messages the peer sent to the local interface.</p> <p>If this number matches the acknowledge-id number in the Local-Advertisement field, this indicates that the local interface has acknowledged all of the peer's state change messages and is synchronized.</p> <p>If this number does not match the acknowledge-id number in the Local-Advertisement field, this indicates that the peer has not yet received an acknowledgment for a state change message from the local interface.</p>
acknowledge-id	<p>Number of acknowledge messages the peer has received from the local interface.</p> <p>If this number matches the sequence-number value in the Local-Advertisement field, this indicates that the peer has acknowledged all of the local interface's state change messages and is in synchronization.</p> <p>If this number does not match the sequence-number value in the Local-Advertisement field, this indicates that the peer has not yet sent an acknowledgment for a state change message from the local interface.</p>

Table 137: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Feature: PFC	Priority-based flow control (PFC) feature DCBX state information.
Protocol-State	(DCBX Version 1.01 only) DCBX protocol state synchronization status: <ul style="list-style-type: none"> • ack-pending—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface. • in-sync—The local interface received an acknowledge message from the peer to indicate that the peer received a PFC state change message sent by the local interface. • not-applicable—PFC autonegotiation is disabled.
Operational State	(DCBX Version 1.01 only) Operational state of the feature: enabled or disabled .
Local-Advertisement	Status of advertisements that the local interface sends to the peer.
Enable	(DCBX Version 1.01 only) State that the local interface advertises to the peer: <ul style="list-style-type: none"> • Yes—The feature is enabled. • No—The feature is disabled.
Willing	Willingness of the local interface to learn the PFC configuration from the peer using DCBX: <ul style="list-style-type: none"> • Yes—The local interface is willing to learn the PFC configuration from the peer. • No—The local interface is not willing to learn the PFC configuration from the peer.
Mac auth Bypass Capability	(IEEE DCBX only) (QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. This is not supported, so the only value seen in the local advertisement field is no .
Error	(DCBX Version 1.01 only) Configuration compatibility error status: <ul style="list-style-type: none"> • No—No error detected. Local and peer configuration are compatible. • Yes—Error detected. Local and peer configuration are not compatible.

Table 137: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Operational State	<p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> • Enabled—PFC is enabled on the interface • Disabled—PFC is disabled on the interface
Maximum Traffic Classes capable to support PFC	<p>Largest number of traffic classes the local interface supports for PFC:</p> <ul style="list-style-type: none"> • 6 (EX Series switches) • 8 (QFX Series)
Code Point	<p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>
Admin Mode	<p>PFC administrative state for each code point on the local interface:</p> <ul style="list-style-type: none"> • Enabled—PFC is enabled for the code point. • Disabled—PFC is disabled for the code point.
Operational Mode	<p>(QFX Series) PFC operational mode for each code point:</p> <ul style="list-style-type: none"> • Enable—PFC is enabled on the code point. • Disable—PFC is disabled on the code point.
Peer-Advertisement	<p>Status of advertisements that the peer sends to the local interface.</p>
Enable	<p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> • Yes—The feature is enabled. • No—The feature is disabled.
Willing	<p>Willingness of the peer to learn the PFC configuration from the local interface using DCBX:</p> <ul style="list-style-type: none"> • Yes—The peer is willing to learn the PFC configuration from the local interface. • No—The peer is not willing to learn the PFC configuration from the local interface.
Error	<p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> • No—No error detected. Local and peer configuration are compatible. • Yes—Error detected. Local and peer configuration are not compatible.

Table 137: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Operational State	<p>PFC operational state on the interface:</p> <ul style="list-style-type: none"> • Enabled—PFC is enabled on the interface • Disabled—PFC is disabled on the interface
Mac auth Bypass Capability	<p>(IEEE DCBX only)</p> <p>(QFX Series) Media access controller (MAC) authentication bypass provides access to devices based on MAC address authentication. Although the QFX Series does not support this feature, the connected peer might support it. This field reports the peer state:</p> <ul style="list-style-type: none"> • Yes—The connected peer supports MAC authentication bypass. • No—The connected peer does not support MAC authentication bypass.
Maximum Traffic Classes capable to support PFC	<p>Largest number of traffic classes the peer supports for PFC:</p> <ul style="list-style-type: none"> • 6 (EX Series switches) • 8 (QFX Series)
Code Point	<p>PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.</p>
Admin Mode	<p>PFC administrative state for each code point on the peer:</p> <ul style="list-style-type: none"> • Enabled—PFC is enabled for the code point. • Disabled—PFC is disabled for the code point.

Table 137: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Feature: Application	State information for the DCBX application.
Protocol-State	<p>(DCBX Version 1.01 only)</p> <p>DCBX protocol state synchronization status:</p> <ul style="list-style-type: none"> • in-sync—The local interface received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface. • ack-pending—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an FCoE state change message sent by the local interface. • not-applicable—The local interface is set to no-auto-negotiation (autonegotiation is disabled). If the interface is associated with an FCoE forwarding class, the interface advertises FCoE capability even if the connected peer does not advertise FCoE capability.
Local-Advertisement	<p>Status of advertisements that the local interface sends to the peer.</p> <p>If the local interface is set to no-auto-negotiation (autonegotiation is disabled), the local advertisement portion of the output is not shown.</p>
Enable	<p>(DCBX Version 1.01 only)</p> <p>State that the local interface advertises to the peer:</p> <ul style="list-style-type: none"> • Yes—The feature is enabled. • No—The feature is disabled.
Willing	<p>(DCBX Version 1.01 only)</p> <p>Willingness of the local interface to learn the FCoE interface state from the peer using DCBX:</p> <ul style="list-style-type: none"> • Yes—The local interface is willing to learn the FCoE interface state from the peer. • No—The local interface is not willing to learn the FCoE interface state from the peer.
Error	<p>(DCBX Version 1.01 only)</p> <p>Configuration compatibility error status:</p> <ul style="list-style-type: none"> • No—No error detected. The local and peer configuration are compatible. • Yes—Error detected. The local and peer configuration are not compatible.
Appl-Name	Name of the application:

Table 137: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Ethernet-Type	<p>(DCBX Version 1.01 only)</p> <p>Ethernet type (EtherType) of the application. For example, 0x8906 indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>
Socket-Number	<p>Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.</p>
Priority-Field or Priority-Map	<p>Priority assigned to the application.</p> <p>For EX Series switches, the priority of the FCoE application is determined by the PFC congestion notification profile that has been configured and associated with the FCoE interface. For other applications, the priority is based on the application map.</p>
Status	<p>(DCBX Version 1.01 only)</p> <p>Local status when autonegotiation is enabled:</p> <ul style="list-style-type: none"> • Enabled—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.) • Disabled—The local configuration and the peer configuration do not match. <p>NOTE: If there is a configuration mismatch in one application between the switch and the peer, all the other applications including FCoE are disabled.</p>
Peer-Advertisement	<p>Status of advertisements that the peer sends to the local interface.</p>
Enable	<p>(DCBX Version 1.01 only)</p> <p>State that the peer advertises to the local interface:</p> <ul style="list-style-type: none"> • Yes—The feature is enabled. • No—The feature is disabled.
Willing	<p>(DCBX Version 1.01 only)</p> <p>Willingness of the peer to learn the FCoE interface state from the local interface using DCBX:</p> <ul style="list-style-type: none"> • Yes—The peer is willing to learn the FCoE interface state from the local interface. • No—The peer is not willing to learn the FCoE interface state from the local interface.

Table 137: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Error	(DCBX Version 1.01 only) Configuration compatibility error status: <ul style="list-style-type: none"> • No—No error detected. Local and peer configuration are compatible. • Yes—Error detected. Local and peer configuration are not compatible.
Appl-Name	Name of the application: <ul style="list-style-type: none"> • FCoE—Fibre Channel over Ethernet
Ethernet-Type	Ethernet type (EtherType) of the application. For example, 0x8906 indicates the EtherType for the FCoE application. Either the EtherType (for Layer 2 applications) or the Socket-Number (for Layer 4 applications) of the application is displayed in the output.
Socket-Number	Destination port socket number of the application, if applicable. Either the EtherType (for Layer 2 applications) or the Socket Number (for Layer 4 applications) of the application is displayed in the output.
Priority-Field or Priority-Map	Priority assigned to the application.
Status	(DCBX Version 1.01 only) Peer interface status: <ul style="list-style-type: none"> • Enabled—The application feature is enabled on both the local interface and the peer interface. (The local configuration and the peer configuration match.) • Disabled—The local configuration and the peer configuration do not match.

Table 137: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Feature: ETS	Enhanced Transmission Selection (ETS) DCBX state information.
Protocol-State	(DCBX Version 1.01 only) ETS protocol state synchronization status: <ul style="list-style-type: none"> • in-sync—The local interface received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface. • ack-pending—The local interface has not yet received an acknowledge message from the peer to indicate that the peer received an ETS state change message sent by the local interface.
Operational State	(DCBX Version 1.01 only) Operational state of the feature, enabled or disabled .
Local-Advertisement	Status of advertisements that the local interface sends to the peer.
Enable	(DCBX Version 1.01 only) State that the local interface advertises to the peer: <ul style="list-style-type: none"> • Yes—The feature is enabled. • No—The feature is disabled.
TLV Type	(IEEE DCBX only) Type of ETS TLV: <ul style="list-style-type: none"> • Configuration—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration. • Recommendation—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration. • Recommendation-or-Configuration—Advertises both TLVs.
Willing	Willingness of the local interface to learn the ETS state from the peer using DCBX (EX Series switches always advertise No for this field): <ul style="list-style-type: none"> • Yes—Local interface is willing to learn the ETS state from the peer. • No—Local interface is not willing to learn the ETS state from the peer.
Credit Based Shaper	

Table 137: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
	(IEEE DCBX only) Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always No .
Error	(DCBX Version 1.01 only) Configuration error status: <ul style="list-style-type: none"> • No—No error. This should always be the switch ETS error state. • Yes—Error detected.
Maximum Traffic Classes capable to support PFC	(DCBX Version 1.01 only) Largest number of traffic classes the local interface supports for PFC.
Maximum Traffic Classes supported	(IEEE DCBX only) Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)
Code Point	PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.
Priority-Group	Class-of-service (CoS) priority group (forwarding class set) identification number.
Percentage B/W	Configured minimum percentage of link bandwidth allocated to the priority group. Only explicitly configured values appear in this output column. If the link bandwidth is the default percentage, it is not shown. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.)
Transmission Selection Algorithm	(IEEE DCBX only) The transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is ETS .
Peer-Advertisement	Status of advertisements that the peer sends to the local interface.
Enable	

Table 137: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
	(DCBX Version 1.01 only) State that the peer advertises to the local interface: <ul style="list-style-type: none"> • Yes—The feature is enabled. • No—The feature is disabled.
TLV Type	(IEEE DCBX only) Type of ETS TLV: <ul style="list-style-type: none"> • Configuration—Advertises the Configuration TLV, which communicates the local ETS configuration to the peer but does not ask the peer to use the configuration. • Recommendation—Advertises the Recommendation TLV, which communicates the local ETS configuration to the peer, and if the peer is “willing,” configures the peer interface to match the local ETS configuration. • Configuration/Recommendation—Advertises both TLVs.
Willing	Willingness of the peer to learn the ETS state from the local interface using DCBX: <ul style="list-style-type: none"> • Yes—Peer is willing to learn the ETS state from the local interface. • No—Peer is not willing to learn the ETS state from the local interface.
Credit Based Shaper	(IEEE DCBX only) Alternative method of flow control to buffer-to-buffer credit. The QFX Series does not support a credit-based shaper, so the value of this field is always No .
Error	(DCBX Version 1.01 only) Configuration error status of the peer: <ul style="list-style-type: none"> • No—No error in peer ETS TLV. • Yes—Error in peer ETS TLV.
Maximum Traffic Classes capable to support PFC	(DCBX Version 1.01 only) Largest number of traffic classes the local interface supports for PFC.
Maximum Traffic Classes supported	(IEEE DCBX only) Largest number of traffic classes the local interface supports for ETS. (EX Series switches support only one traffic class for ETS. However, a different value might be shown for this field.)
Code Point	

Table 137: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
	PFC code point, which is specified in the 3-bit class-of-service field in the VLAN header.
Priority-Group	CoS priority group (forwarding class set) identification number.
Percentage B/W	Configured minimum percentage of link bandwidth allocated to the priority group. (EX Series switches allocate 100% of link bandwidth to the default priority group, group 7.)
Transmission Selection Algorithm	(IEEE DCBX only) Transmission selection algorithm used by the interface. The QFX Series supports ETS but does not support using the credit-based shaper algorithm, so the only value shown in this field is ETS .
PFC	(QFX Series, terse option only) DCBX TLV advertisement state for PFC: <ul style="list-style-type: none"> • Disabled—PFC configuration matches the configuration on the connected peer and PFC is disabled • Enabled—PFC configuration matches the configuration on the connected peer and PFC is enabled • Not Advt—Interface does not advertise PFC to the connected peer
ETS	(terse option only) Local DCBX TLV advertisement state for ETS: <ul style="list-style-type: none"> • Advt—Interface advertises ETS TLVs • Disabled—ETS is disabled on the interface (interface does not advertise ETS)
ETS Rec	(terse option only) DCBX TLV peer advertisement state for ETS (state received from the connected DCBX peer): <ul style="list-style-type: none"> • Advt—Peer interface advertises ETS TLVs • Not Advt—Peer interface does not advertise ETS <p>NOTE: When the DCBX mode is DCBX version 1.01, no peer information is displayed.</p>

Table 137: show dcbx neighbors Output Fields (*continued*)

Field Name	Field Description
Version	<p>(terse option only) The DCBX version used on the interface and whether the DCBX version was autonegotiated or explicitly configured:</p> <ul style="list-style-type: none"> • IEEE—The interface uses IEEE DCBX. • 1.01—The interface uses DCBX version 1.01. <p>When the DCBX version used is the result of autonegotiation, the term (Auto) appears next to the version. For example, IEEE (Auto) indicates that the interface autonegotiated with the connected peer to use IEEE DCBX. Autonegotiation is enabled by default.</p>

Sample Output

show dcbx neighbors interface (QFX Series, DCBX Version 1.01 Mode)

```

user@switch> show dcbx neighbors interface xe-0/0/0
Interface : xe-0/0/0.0 - Parent Interface: ae0.0
Active-application-map: app-map-1
Protocol-State: in-sync
Protocol-Mode: DCBX Version 1.01

Local-Advertisement:
  Operational version: 1
  sequence-number: 130, acknowledge-id: 102

Peer-Advertisement:
  Operational version: 1
  sequence-number: 102, acknowledge-id: 130

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
  Enable: Yes, Willing: No, Error: No
  Maximum Traffic Classes capable to support PFC: 8

Code Point      Admin Mode      Operational Mode
000             Disabled       Disable
001             Disabled       Disable
010             Disabled       Disable
011             Enabled        Enable
100             Enabled        Enable
101             Disabled       Disable
110             Disabled       Disable
111             Disabled       Disable

Peer-Advertisement:
  Enable: Yes, Willing: No, Error: No
  Maximum Traffic Classes capable to support PFC: 8

Code Point      Admin Mode
000             Disabled

```

001	Disabled
010	Disabled
011	Enabled
100	Enabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

App1-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00001110	Enabled
iSCSI		3260	10000000	Enabled

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

App1-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906	N/A	00001110	Enabled

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0
101	1
110	1
111	7

Priority-Group	Percentage B/W
0	40%
1	5%

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 8

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0
101	1
110	1

111	7
Priority-Group	Percentage B/W
0	40%
1	5%

show dcbx neighbors interface (QFX Series, IEEE DCBX Mode)

user@switch> **show dcbx neighbors interface xe-0/0/0**

Interface : xe-0/0/0.0 - Parent Interface: ae0.0

Active-application-map: app-map-1

Protocol-Mode: IEEE-DCBX Version

Feature: PFC

Local-Advertisement:

Willing: No

Mac auth Bypass Capability: No

Operational State: Enabled

Maximum Traffic Classes capable to support PFC: 8

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Enabled
101	Disabled
110	Disabled
111	Disabled

Peer-Advertisement:

Willing: No

Mac auth Bypass Capability: No

Operational State: Enabled

Maximum Traffic Classes capable to support PFC: 8

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Enabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application

Local-Advertisement:

App1-Name	Ethernet-Type	Socket-Number	Priority-field
FCoE	0x8906		00001110
iSCSI		3260	10000000

Peer-Advertisement:

App1-Name	Ethernet-Type	Socket-Number	Priority-field
-----------	---------------	---------------	----------------

FCoE	0x8906	N/A	00001110
------	--------	-----	----------

Feature: ETS

Local-Advertisement:

TLV Type: Configuration/Recommendation

Willing: No

Credit Based Shaper: No

Maximum Traffic Classes supported: 3

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0
101	1
110	1
111	7

Priority-Group	Percentage B/W
0	40%
1	5%

Priority-Group	Transmission Selection Algorithm
0	Enhanced Transmission Selection
1	Enhanced Transmission Selection

Peer-Advertisement:

TLV Type: Configuration

Willing: No

Credit Based Shaper: No

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0
101	1
110	1
111	7

Priority-Group	Percentage B/W
0	40%
1	5%

Priority-Group	Transmission Selection Algorithm
0	Enhanced Transmission Selection
1	Enhanced Transmission Selection

Peer-Advertisement:

TLV Type: Recommendation

Code Point	Priority-Group
000	0
001	7
010	7
011	7
100	0

101	1
110	1
111	7
Priority-Group	Percentage B/W
0	40%
1	5%
Priority-Group	Transmission Selection Algorithm
0	Enhanced Transmission Selection
1	Enhanced Transmission Selection

show dcbx neighbors terse (QFX Series)

```

user@switch> show dcbx neighbors terse
Interface Parent PFC ETS ETS Version
Interface Rec
xe-0/0/8.0 - Enabled Advt Advt IEEE (Auto)
xe-0/0/9.0 - Disabled Disabled 1.01
xe-0/0/11.0 ae0.0 Enabled Advt Advt IEEE (Auto)
xe-0/0/12.0 ae0.0 Enabled Advt Advt IEEE (Auto)
xe-0/0/32.0 - Enabled Advt Not Advt IEEE
xe-0/0/36.0 - Not Advt Advt Advt IEEE

```

show dcbx neighbors (EX4500 Switch: FCoE Interfaces on Both Local and Peer with PFC Configured Compatibly)

```

user@switch> show dcbx neighbors interface xe-0/0/14

Interface : xe-0/0/14.0 - Parent Interface: ae0.0
Protocol-State: in-sync

Local-Advertisement:
  Operational version: 0
  sequence-number: 6, acknowledge-id: 6

Peer-Advertisement:
  Operational version: 0
  sequence-number: 6, acknowledge-id: 6

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:
  Enable: Yes, Willing: No, Error: No
  Maximum Traffic Classes capable to support PFC: 6

Code Point      Admin Mode
000             Disabled
001             Disabled
010             Disabled
011             Enabled
100             Disabled
101             Disabled
110             Disabled
111             Disabled

```

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Disabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No <<< Error bit will not be set as there is no miss configuration between local and peer.

Appl-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00001000	Enabled

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Status	Appl-Name	Ethernet-Type	Socket-Number	Priority-Map
Enabled	FCoE	0x8906		00001000

show dcbx neighbors (EX4500 Switch: DCBX Interfaces on Local and Peer Are Configured Compatibly with iSCSI Application)

user@switch> show dcbx neighbors interface xe-0/0/14

Interface : xe-0/0/14.0 - Parent Interface: ae0.0

Protocol-State: in-sync

Active-application-map: iscsi-map

Local-Advertisement:

Operational version: 0

sequence-number: 9, acknowledge-id: 12

Peer-Advertisement:

Operational version: 0

sequence-number: 12, acknowledge-id: 9

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Disabled
101	Disabled
110	Disabled
111	Disabled

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

Code Point	Admin Mode
000	Disabled
001	Disabled
010	Disabled
011	Enabled
100	Disabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Appl-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00001000	Enabled
iscsi		3260	00100000	Enabled

Peer-Advertisement:

Enable: Yes, Willing: No, Error: No

Appl-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00001000	Enabled
iscsi		3260	00100000	Enabled

show dcbx neighbors (EX4500 Switch: Includes ETS)

user@switch> show dcbx neighbors interface xe-0/0/3

Interface : xe-0/0/3.0
 Protocol-State: in-sync
 Active-application-map: map_iscsi

Local-Advertisement:

Operational version: 0

sequence-number: 1, acknowledge-id: 5

Peer-Advertisement:

Operational version: 0

sequence-number: 5, acknowledge-id: 1

Feature: PFC, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes capable to support PFC: 6

Code Point	Admin Mode
000	Enabled
001	Enabled
010	Disabled
011	Disabled
100	Disabled
101	Disabled
110	Disabled
111	Disabled

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

Maximum Traffic Classes capable to support PFC: 8

Code Point	Admin Mode
000	Enabled
001	Disabled
010	Disabled
011	Disabled
100	Enabled
101	Disabled
110	Disabled
111	Disabled

Feature: Application, Protocol-State: in-sync

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

App1-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		00000001	Enabled
iscsi		3260	00000010	Enabled

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

App1-Name	Ethernet-Type	Socket-Number	Priority-Map	Status
FCoE	0x8906		0001000	Enabled
iscsi		3260	00010000	Enabled

Feature: ETS, Protocol-State: in-sync

Operational State: Enabled

Local-Advertisement:

Enable: Yes, Willing: No, Error: No

Maximum Traffic Classes supported : 3

Code Point	Priority-Group
000	7
001	7
010	7
011	7
100	7
101	7
110	7
111	7

Priority-Group	Percentage B/W
7	100%

Peer-Advertisement:

Enable: Yes, Willing: Yes, Error: No

Maximum Traffic Classes supported : 8

Code Point	Priority-Group
000	0
001	1
010	0
011	0
100	2
101	0
110	0
111	0

Priority-Group	Percentage B/W
0	30%
1	40%
2	30%

show interfaces queue

Syntax	<pre>show interfaces queue <aggregate remaining-traffic> <both-ingress-egress> <egress> <forwarding-class forwarding-class> <ingress> <interface-name interface-name> <l2-statistics></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>both-ingress-egress, egress, and ingress options introduced in Junos OS Release 7.6.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p> <p>l2-statistics option introduced in Junos OS Release 12.1.</p>
Description	Display class-of-service (CoS) queue information for physical interfaces.
Options	<p>none—Show detailed CoS queue statistics for all physical interfaces.</p> <p>aggregate—(Optional) Display the aggregated queuing statistics of all logical interfaces that have traffic-control profiles configured. (Not on the QFX Series.)</p> <p>both-ingress-egress—(Optional) On Gigabit Ethernet Intelligent Queuing 2 (IQ2) PICs, display both ingress and egress queue statistics. (Not on the QFX Series.)</p> <p>egress—(Optional) Display egress queue statistics.</p> <p>forwarding-class forwarding-class—(Optional) Forwarding class name for this queue. Shows detailed CoS statistics for the queue associated with the specified forwarding class.</p> <p>ingress—(Optional) On Gigabit Ethernet IQ2 PICs, display ingress queue statistics. (Not on the QFX Series.)</p> <p>interface-name interface-name—(Optional) Show detailed CoS queue statistics for the specified interface.</p> <p>l2-statistics—(Optional) Display Layer 2 statistics for MLPPP, FRF.15, and FRF.16 bundles</p> <p>remaining-traffic—(Optional) Display the remaining-traffic queue statistics of all logical interfaces that have traffic-control profiles configured.</p>

Overhead for Layer 2 Statistics

Transmitted packets and transmitted byte counts are displayed for the Layer 2 level with the addition of encapsulation overheads applied for fragmentation, as shown in [Table 138 on page 648](#). Others counters, such as packets and bytes queued (input) and drop counters, are displayed at the Layer 3 level. In the case of link fragmentation and interleaving (LFI) for which fragmentation is not applied, corresponding Layer 2 overheads are added, as shown in [Table 138 on page 648](#).

Table 138: Layer 2 Overhead, Transmitted Packets/Bytes

Protocol	Fragmentation		LFI
	First fragmentation	Second to n fragmentations	
	Bytes	Bytes	
MLPPP (Long)	13	12	8
MLPPP (short)	11	10	8
MLFR (FRF15)	12	10	8
MFR (FRF16)	10	8	-
MCMLPPP(Long)	13	12	-
MCMLPPP(Short)	11	10	-

Layer 2 Statistics - Fragmentation Overhead Calculation

MLPPP/MC-MLPPP Overhead details:

=====

Fragment 1:

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
Inner PPP header           : 1 byte
HDLC flag and FCS bytes    : 4 bytes

```

Fragments 2 .. n :

```

Outer PPP header           : 4 bytes
Long or short sequence MLPPP header : 4 bytes or 2 bytes
HDLC flag and FCS bytes    : 4 bytes

```

MLFR (FRF15) Overhead details:

=====

Fragment 1:

```

Framereelay header        : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
Inner proto               : 2 bytes
HDLC flag and FCS         : 4 bytes

```

Fragments 2 ...n :

```

Framereelay header        : 2 bytes
Control,NLPID             : 2 bytes
Fragmentaion header       : 2 bytes
HDLC flag and FCS         : 4 bytes

```

MFR (FRF16) Overhead details:

=====

```

Fragment 1:
  Fragmentation header : 2 bytes
  Framereelay header   : 2 bytes
  Inner proto          : 2 bytes
  HDLC flag and FCS    : 4 bytes

Fragments 2 ...n :
  Fragmentation header : 2 bytes
  Framereelay header   : 2 bytes
  HDLC flag and FCS    : 4 bytes

```

Overhead with LFI

```

MLPPP(Long & short sequence):
=====
  Outer PPP header : 4 bytes
  HDLC flag and FCS : 4 bytes

MLFR (FRF15):
=====
  Framereelay header : 2 bytes
  Control,NLPID      : 2 bytes
  HDLC flag and FCS  : 4 bytes

```

The following examples show overhead for different cases:

- A 1000-byte packet is sent to a mlppp bundle without any fragmentation. At the Layer 2 level, bytes transmitted is 1013 in 1 packet. This overhead is for MLPPP long sequence encap.
- A 1000-byte packet is sent to a mlppp bundle with a fragment threshold of 250byte. At the Layer 2 level, bytes transmitted is 1061 bytes in 5 packets.
- A 1000-byte LFI packet is sent to an mlppp bundle. At the Layer 2 level, bytes transmitted is 1008 in 1 packet.

remaining-traffic—(Optional) Display the queuing statistics of all logical interfaces that do not have traffic-control profiles configured. (Not on the QFX Series.)

Additional Information

For rate-limited interfaces hosted on Modular Interface Cards (MICs) or Modular Port Concentrators (MPCs), rate-limit packet-drop operations occur *before* packets are queued for transmission scheduling. For such interfaces, the statistics for queued traffic do not include the packets that have already been dropped due to rate limiting, and consequently the displayed statistics for queued traffic are the same as the displayed statistics for transmitted traffic.



NOTE: For rate-limited interfaces hosted on other types of hardware, rate-limit packet-drop operations occur *after* packets are queued for transmission scheduling. For these other interface types, the statistics for queued traffic include the packets that are later dropped due to rate limiting, and consequently the displayed statistics for queued traffic equals the sum of the statistics for transmitted and rate-limited traffic.

On M Series routers (except for the M320 and M120 routers), this command is valid only for a PIC installed on an enhanced Flexible PIC Concentrator (FPC).

Queue statistics for aggregated interfaces are supported on the M Series and T Series routers only. Statistics for an aggregated interface are the summation of the queue statistics of the child links of that aggregated interface. You can view the statistics for a child interface by using the **show interfaces statistics** command for that child interface.

When you configure tricolor marking on a 10-port 1-Gigabit Ethernet PIC, for queues 6 and 7 only, the output does not display the number of queued bytes and packets, or the number of bytes and packets dropped because of RED. If you do not configure tricolor marking on the interface, these statistics are available for all queues.

For the 4-port Channelized OC12 IQE PIC and 1-port Channelized OC48 IQE PIC, the **Packet Forwarding Engine Chassis Queues** field represents traffic bound for a particular physical interface on the PIC. For all other PICs, the **Packet Forwarding Engine Chassis Queues** field represents the total traffic bound for the PIC.

For Gigabit Ethernet IQ2 PICs, the **show interfaces queue** command output does not display the number of tail-dropped packets. This limitation does not apply to Packet Forwarding Engine chassis queues.

When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (under the **Packet Forwarding Engine Chassis Queues** field) shows the prefragmentation values.

The behavior of the **egress** queues for the **Routing Engine-Generated Traffic** is not same as the configured queue for MLPPP and MFR configurations.

For information about how to configure CoS, see the *Junos OS Network Interfaces Library for Routing Devices*. For related CoS operational mode commands, see the [CLI Explorer](#).

Required Privilege Level	view
List of Sample Output	show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC) on page 655 show interfaces queue (Aggregated Ethernet on a T320 Router) on page 656 show interfaces queue (Fast Ethernet on a J4300 Router) on page 658 show interfaces queue (Gigabit Ethernet on a T640 Router) on page 658 show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC) on page 659 show interfaces queue (Gigabit Ethernet IQ2 PIC) on page 663 show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC) on page 666 show interfaces queue ingress (Gigabit Ethernet IQ2 PIC) on page 668 show interfaces queue egress (Gigabit Ethernet IQ2 PIC) on page 669 show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC) on page 670 show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode) on page 673 show interfaces queue (QFX Series) on page 683 show interfaces queue l2-statistics (lsq interface) on page 684
Output Fields	Table 139 on page 651 lists the output fields for the show interfaces queue command. Output fields are listed in the approximate order in which they appear.

Table 139: show interfaces queue Output Fields

Field Name	Field Description
Physical interface	Name of the physical interface.
Enabled	State of the interface. Possible values are described in the "Enabled Field" section under <i>Common Output Fields Description</i> .
Interface index	Physical interface's index number, which reflects its initialization sequence.
SNMP ifIndex	SNMP index number for the interface.
Forwarding classes supported	Total number of forwarding classes supported on the specified interface.
Forwarding classes in use	Total number of forwarding classes in use on the specified interface.
Ingress queues supported	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues supported on the specified interface.
Ingress queues in use	On Gigabit Ethernet IQ2 PICs only, total number of ingress queues in use on the specified interface.
Output queues supported	Total number of output queues supported on the specified interface.
Output queues in use	Total number of output queues in use on the specified interface.
Egress queues supported	Total number of egress queues supported on the specified interface.
Egress queues in use	Total number of egress queues in use on the specified interface.
Queue counters (Ingress)	CoS queue number and its associated user-configured forwarding class name. Displayed on IQ2 interfaces. <ul style="list-style-type: none"> • Queued packets—Number of queued packets. • Transmitted packets—Number of transmitted packets. • Dropped packets—Number of packets dropped by the ASIC's RED mechanism.
Burst size	(Logical interfaces on IQ PICs only) Maximum number of bytes up to which the logical interface can burst. The burst size is based on the shaping rate applied to the interface.
The following output fields are applicable to both interface component and Packet Forwarding component in the show interfaces queue command:	
Queue	Queue number.
Forwarding classes	Forwarding class name.

Table 139: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
Queued Packets	<p>Number of packets queued to this queue.</p> <p>NOTE: For Gigabit Ethernet IQ2 interfaces, the Queued Packets count is calculated by the Junos OS interpreting one frame buffer as one packet. If the queued packets are very large or very small, the calculation might not be completely accurate for transit traffic. The count is completely accurate for traffic terminated on the router.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see "Additional Information" on page 649.</p>
Queued Bytes	<p>Number of bytes queued to this queue. The byte counts vary by interface hardware. For more information, see Table 140 on page 654.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic does not include traffic dropped due to rate limiting. For more information, see "Additional Information" on page 649.</p>
Transmitted Packets	<p>Number of packets transmitted by this queue. When fragmentation occurs on the egress interface, the first set of packet counters shows the postfragmentation values. The second set of packet counters (displayed under the Packet Forwarding Engine Chassis Queues field) shows the prefragmentation values.</p> <p>NOTE: For Layer 2 statistics, see "Overhead for Layer 2 Statistics" on page 647</p>
Transmitted Bytes	<p>Number of bytes transmitted by this queue. The byte counts vary by interface hardware. For more information, see Table 140 on page 654.</p> <p>NOTE: On MX Series routers, this number can be inaccurate when you issue the command for a physical interface repeatedly and in quick succession, because the statistics for the child nodes are collected infrequently. Wait ten seconds between successive iterations to avoid this situation.</p> <p>NOTE: For Layer 2 statistics, see "Overhead for Layer 2 Statistics" on page 647</p>
Tail-dropped packets	Number of packets dropped because of tail drop.
RL-dropped packets	<p>Number of packets dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic is not included in the queued traffic statistics. For more information, see "Additional Information" on page 649.</p>
RL-dropped bytes	<p>Number of bytes dropped due to rate limiting.</p> <p>For rate-limited interfaces hosted on MICs or MPCs only, this statistic is not included in the queued traffic statistics. For more information, see "Additional Information" on page 649.</p>

Table 139: show interfaces queue Output Fields (*continued*)

Field Name	Field Description
RED-dropped packets	<p>Number of packets dropped because of random early detection (RED).</p> <ul style="list-style-type: none"> (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, the total number of dropped packets is displayed. On all other M Series routers, the output classifies dropped packets into the following categories: <ul style="list-style-type: none"> Low, non-TCP—Number of low-loss priority non-TCP packets dropped because of RED. Low, TCP—Number of low-loss priority TCP packets dropped because of RED. High, non-TCP—Number of high-loss priority non-TCP packets dropped because of RED. High, TCP—Number of high-loss priority TCP packets dropped because of RED. (J Series routers and MX Series routers with enhanced DPCs, and T Series routers with enhanced FPCs only) The output classifies dropped packets into the following categories: <ul style="list-style-type: none"> Low—Number of low-loss priority packets dropped because of RED. Medium-low—Number of medium-low loss priority packets dropped because of RED. Medium-high—Number of medium-high loss priority packets dropped because of RED. High—Number of high-loss priority packets dropped because of RED. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>
RED-dropped bytes	<p>Number of bytes dropped because of RED. The byte counts vary by interface hardware. For more information, see Table 140 on page 654.</p> <ul style="list-style-type: none"> (M Series and T Series routers only) On M320 and M120 routers and the T Series routers, only the total number of dropped bytes is displayed. On all other M Series routers, the output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> Low, non-TCP—Number of low-loss priority non-TCP bytes dropped because of RED. Low, TCP—Number of low-loss priority TCP bytes dropped because of RED. High, non-TCP—Number of high-loss priority non-TCP bytes dropped because of RED. High, TCP—Number of high-loss priority TCP bytes dropped because of RED. (J Series routers only) The output classifies dropped bytes into the following categories: <ul style="list-style-type: none"> Low—Number of low-loss priority bytes dropped because of RED. Medium-low—Number of medium-low loss priority bytes dropped because of RED. Medium-high—Number of medium-high loss priority bytes dropped because of RED. High—Number of high-loss priority bytes dropped because of RED. <p>NOTE: Due to accounting space limitations on certain Type 3 FPCs (which are supported in M320 and T640 routers), this field does not always display the correct value for queue 6 or queue 7 for interfaces on 10-port 1-Gigabit Ethernet PICs.</p>

Byte counts vary by interface hardware. [Table 140 on page 654](#) shows how the byte counts on the outbound interfaces vary depending on the interface hardware.

[Table 140 on page 654](#) is based on the assumption that outbound interfaces are sending IP traffic with 478 bytes per packet.

Table 140: Byte Count by Interface Hardware

Interface Hardware	Output Level	Byte Count Includes	Comments
Gigabit Ethernet IQ and IQE PICs	Interface	<p>Queued: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>Transmitted: 490 bytes per packet, representing 478 bytes of Layer 3 packet + 12 bytes</p> <p>RED dropped: 496 bytes per packet representing 478 bytes of Layer 3 packet + 18 bytes</p>	<p>The 12 additional bytes include 6 bytes for the destination MAC address + 4 bytes for the VLAN + 2 bytes for the Ethernet type.</p> <p>For RED dropped, 6 bytes are added for the source MAC address.</p>
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p>	—
Non-IQ PIC	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet. <p>T4000 routers with Type 5 FPCs :</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Inter frame Gap. • Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead 8 bytes preamble + 12 bytes Interframe Gap. <p>M Series routers:</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet. • Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead. <p>PTX Series Packet Transport Routers:</p> <ul style="list-style-type: none"> • Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes FCS + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN). • Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including 4 bytes CRC + the full Layer 1 overhead of the MAC header DA + SA + EtherType (non-VLAN). • RED dropped: 478 bytes of Layer 3 packet + 22 bytes special header. To the TQ, this packet has 4 bytes more than queued or transmitted. 	<p>The Layer 2 overhead is 14 bytes for non-VLAN traffic and 18 bytes for VLAN traffic.</p>

Table 140: Byte Count by Interface Hardware (*continued*)

Interface Hardware	Output Level	Byte Count Includes	Comments
IQ and IQE PICs with a SONET/SDH interface	Interface	<p>Queued: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>Transmitted: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p> <p>RED dropped: 482 bytes per packet, representing 478 bytes of Layer 3 packet + 4 bytes</p>	The additional 4 bytes are for the Layer 2 Point-to-Point Protocol (PPP) header.
	Packet forwarding component	<p>Queued: 478 bytes per packet, representing 478 bytes of Layer 3 packet</p> <p>Transmitted: 486 bytes per packet, representing 478 bytes of Layer 3 packet + 8 bytes</p>	For transmitted packets, the additional 8 bytes includes 4 bytes for the PPP header and 4 bytes for a cookie.
Non-IQ PIC with a SONET/SDH interface	Interface	<p>T Series, TX Series, T1600, and MX Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 478 bytes of Layer 3 packet. <p>M Series routers:</p> <ul style="list-style-type: none"> Queued: 478 bytes of Layer 3 packet. Transmitted: 483 bytes per packet, representing 478 bytes of Layer 3 packet + 5 bytes RED dropped: 478 bytes per packet, representing 478 bytes of Layer 3 packet 	For transmitted packets, the additional 5 bytes includes 4 bytes for the PPP header and 1 byte for the packet loss priority (PLP).
Interfaces configured with Frame Relay Encapsulation	Interface	The default Frame Relay overhead is 7 bytes. If you configure the Frame Check Sequence (FCS) to 4 bytes, then the overhead increases to 10 bytes.	
1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs	Interface	<p>Queued: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p> <p>Transmitted: 478 bytes of Layer 3 packet + the full Layer 2 overhead including CRC.</p>	The Layer 2 overhead is 18 bytes for non-VLAN traffic and 22 bytes for VLAN traffic.
4-port 1G IQ2 and IQ2-E PICs	Packet forwarding component	Queued: 478 bytes of Layer 3 packet.	
8-port 1G IQ2 and IQ2-E PICs	Packet forwarding component	Transmitted: 478 bytes of Layer 3 packet.	—

Sample Output

show interfaces queue (Rate-Limited Interface on a Gigabit Ethernet MIC in an MPC)

The following example shows queue information for the rate-limited interface ge-4/2/0 on a Gigabit Ethernet MIC in an MPC. For rate-limited queues for interfaces hosted on MICs or MPCs, rate-limit packet drops occur prior to packet output queuing. In the

command output, the nonzero statistics displayed in the **RL-dropped packets** and **RL-dropped bytes** fields quantify the traffic dropped to rate-limit queue 0 output to 10 percent of 1 gigabyte (100 megabits) per second. Because the RL-dropped traffic is not included in the **Queued** statistics, the statistics displayed for queued traffic are the same as the statistics for transmitted traffic.

```
user@host> show interfaces queue ge-4/2/0
Physical interface: ge-4/2/0, Enabled, Physical link is Up
  Interface index: 203, SNMP ifIndex: 1054
  Forwarding classes: 16 supported, 4 in use
  Egress queues: 8 supported, 4 in use
  Queue: 0, Forwarding classes: best-effort
    Queued:
      Packets          :          131300649          141751 pps
      Bytes            :          11287964840        99793248 bps
    Transmitted:
      Packets          :          131300649          141751 pps
      Bytes            :          11287964840        99793248 bps
      Tail-dropped packets :              0              0 pps
      RL-dropped packets :          205050862          602295 pps
      RL-dropped bytes   :          13595326612      327648832 bps
      RED-dropped packets :              0              0 pps
        Low              :              0              0 pps
        Medium-low       :              0              0 pps
        Medium-high      :              0              0 pps
        High             :              0              0 pps
      RED-dropped bytes   :              0              0 bps
        Low              :              0              0 bps
        Medium-low       :              0              0 bps
        Medium-high      :              0              0 bps
        High             :              0              0 bps
    Queue: 1, Forwarding classes: expedited-forwarding
      Queued:
        Packets          :              0              0 pps
        Bytes            :              0              0 bps
```

show interfaces queue (Aggregated Ethernet on a T320 Router)

The following example shows that the aggregated Ethernet interface, **ae1**, has traffic on queues **af1** and **af12**:

```
user@host> show interfaces queue ae1
Physical interface: ae1, Enabled, Physical link is Up
  Interface index: 158, SNMP ifIndex: 33
  Forwarding classes: 8 supported, 8 in use
  Output queues: 8 supported, 8 in use
  Queue: 0, Forwarding classes: be
    Queued:
      Packets          :              5              0 pps
      Bytes            :              242              0 bps
    Transmitted:
      Packets          :              5              0 pps
      Bytes            :              242              0 bps
      Tail-dropped packets :              0              0 pps
      RED-dropped packets :              0              0 pps
      RED-dropped bytes   :              0              0 bps
    Queue: 1, Forwarding classes: af1
      Queued:
        Packets          :          42603765          595484 pps
```

```

Bytes                :          5453281920          609776496 bps
Transmitted:
Packets              :          42603765           595484 pps
Bytes                :          5453281920          609776496 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 2, Forwarding classes: ef1
Queued:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Transmitted:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 3, Forwarding classes: nc
Queued:
Packets              :              45              0 pps
Bytes                :             3930              0 bps
Transmitted:
Packets              :              45              0 pps
Bytes                :             3930              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 4, Forwarding classes: af11
Queued:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Transmitted:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 5, Forwarding classes: ef11
Queued:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Transmitted:
Packets              :              0              0 pps
Bytes                :              0              0 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 6, Forwarding classes: af12
Queued:
Packets              :          31296413          437436 pps
Bytes                :          4005940864          447935200 bps
Transmitted:
Packets              :          31296413          437436 pps
Bytes                :          4005940864          447935200 bps
Tail-dropped packets :              0              0 pps
RED-dropped packets  :              0              0 pps
RED-dropped bytes    :              0              0 bps
Queue: 7, Forwarding classes: nc2
Queued:
Packets              :              0              0 pps
Bytes                :              0              0 bps

```

```

Transmitted:
Packets      :                0                0 pps
Bytes        :                0                0 bps
Tail-dropped packets :                0                0 pps
RED-dropped packets :                0                0 pps
RED-dropped bytes  :                0                0 bps

```

show interfaces queue (Fast Ethernet on a J4300 Router)

```

user@host> show interfaces queue fe-4/0/0.0
Logical interface fe-4/0/0.0 (Index 71) (SNMP ifIndex 42)
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :                5240762                3404 pps
    Bytes        :            3020710354            15934544 bps
  Transmitted:
    Packets      :                5240762                3404 pps
    Bytes        :            3020710354            15934544 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                0                0 pps
    Low          :                0                0 pps
    Medium-low   :                0                0 pps
    Medium-high  :                0                0 pps
    High         :                0                0 pps
    RED-dropped bytes :                0                0 bps
    Low          :                0                0 pps
    Medium-low   :                0                0 pps
    Medium-high  :                0                0 pps
    High         :                0                0 pps
Queue: 1, Forwarding classes: af1
  Queued:
    Packets      :                2480391                1650 pps
    Bytes        :            1304685666            6945704 bps
  Transmitted:
    Packets      :                2478740                1650 pps
    Bytes        :            1303817240            6945704 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                1651                0 pps
    Low          :                0                0 pps
    Medium-low   :                0                0 pps
    Medium-high  :                0                0 pps
    High         :                1651                0 pps
    RED-dropped bytes :                868426                0 bps
    Low          :                0                0 pps
    Medium-low   :                0                0 pps
    Medium-high  :                0                0 pps
    High         :                868426                0 pps

```

show interfaces queue (Gigabit Ethernet on a T640 Router)

```

user@host> show interfaces queue
Physical interface: ge-7/0/1, Enabled, Physical link is Up
  Interface index: 150, SNMP ifIndex: 42
Forwarding classes: 8 supported, 8 in use
Output queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: be
  Queued:

```

```

Packets      :      13      0 pps
Bytes        :      622      0 bps
Transmitted:
Packets      :      13      0 pps
Bytes        :      622      0 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: af1
Queued:
Packets      :      1725947945      372178 pps
Bytes        :      220921336960      381110432 bps
Transmitted:
Packets      :      1725947945      372178 pps
Bytes        :      220921336960      381110432 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: ef1
Queued:
Packets      :      0      0 pps
Bytes        :      0      0 bps
Transmitted:
Packets      :      0      0 pps
Bytes        :      0      0 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: nc
Queued:
Packets      :      571      0 pps
Bytes        :      49318      336 bps
Transmitted:
Packets      :      571      0 pps
Bytes        :      49318      336 bps
Tail-dropped packets :      0      0 pps
RED-dropped packets :      0      0 pps
RED-dropped bytes  :      0      0 bps

```

show interfaces queue aggregate (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 aggregate
Physical interface: ge-2/2/9, Enabled, Physical link is Up
Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets      :      148450735      947295 pps
Bytes        :      8016344944      409228848 bps
Transmitted:
Packets      :      76397439      487512 pps
Bytes        :      4125461868      210602376 bps
Tail-dropped packets : Not Available
RED-dropped packets :      72053285      459783 pps
Low          :      72053285      459783 pps
Medium-low   :      0      0 pps
Medium-high  :      0      0 pps
High         :      0      0 pps
RED-dropped bytes  :      3890877444      198626472 bps

```

```

        Low                :                3890877444                198626472 bps
        Medium-low         :                0                0 bps
        Medium-high        :                0                0 bps
        High               :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets                :                0                0 pps
  Bytes                 :                0                0 bps
Transmitted:
  Packets                :                0                0 pps
  Bytes                 :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
  Low                   :                0                0 pps
  Medium-low            :                0                0 pps
  Medium-high           :                0                0 pps
  High                  :                0                0 pps
RED-dropped bytes :                0                0 bps
  Low                   :                0                0 bps
  Medium-low            :                0                0 bps
  Medium-high           :                0                0 bps
  High                  :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets                :            410278257                473940 pps
  Bytes                 :            22156199518            204742296 bps
Transmitted:
  Packets                :            4850003                4033 pps
  Bytes                 :            261900162            1742256 bps
Tail-dropped packets : Not Available
RED-dropped packets :            405425693                469907 pps
  Low                   :            405425693                469907 pps
  Medium-low            :                0                0 pps
  Medium-high           :                0                0 pps
  High                  :                0                0 pps
RED-dropped bytes :            21892988124            203000040 bps
  Low                   :            21892988124            203000040 bps
  Medium-low            :                0                0 bps
  Medium-high           :                0                0 bps
  High                  :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets                :                0                0 pps
  Bytes                 :                0                0 bps
Transmitted:
  Packets                :                0                0 pps
  Bytes                 :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
  Low                   :                0                0 pps
  Medium-low            :                0                0 pps
  Medium-high           :                0                0 pps
  High                  :                0                0 pps
RED-dropped bytes :                0                0 bps
  Low                   :                0                0 bps
  Medium-low            :                0                0 bps
  Medium-high           :                0                0 bps
  High                  :                0                0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort

```



```

Queued:
  Packets      :      76605230      485376 pps
  Bytes       :      5209211400    264044560 bps
Transmitted:
  Packets      :      76444631      484336 pps
  Bytes       :      5198235612    263478800 bps
Tail-dropped packets : Not Available
RED-dropped packets :      160475      1040 pps
  Low         :      160475      1040 pps
  Medium-low  :           0         0 pps
  Medium-high :           0         0 pps
  High        :           0         0 pps
RED-dropped bytes  :      10912300    565760 bps
  Low           :      10912300    565760 bps
  Medium-low    :           0         0 bps
  Medium-high   :           0         0 bps
  High         :           0         0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets      :           0         0 pps
  Bytes       :           0         0 bps
Transmitted:
  Packets      :           0         0 pps
  Bytes       :           0         0 bps
Tail-dropped packets : Not Available
RED-dropped packets :           0         0 pps
  Low         :           0         0 pps
  Medium-low  :           0         0 pps
  Medium-high :           0         0 pps
  High        :           0         0 pps
RED-dropped bytes  :           0         0 bps
  Low         :           0         0 bps
  Medium-low  :           0         0 bps
  Medium-high :           0         0 bps
  High        :           0         0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets      :      4836136      3912 pps
  Bytes       :      333402032    2139056 bps
Transmitted:
  Packets      :      3600866      1459 pps
  Bytes       :      244858888    793696 bps
Tail-dropped packets : Not Available
RED-dropped packets :      1225034      2450 pps
  Low         :      1225034      2450 pps
  Medium-low  :           0         0 pps
  Medium-high :           0         0 pps
  High        :           0         0 pps
RED-dropped bytes  :      83302312    1333072 bps
  Low         :      83302312    1333072 bps
  Medium-low  :           0         0 bps
  Medium-high :           0         0 bps
  High        :           0         0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets      :           0         0 pps
  Bytes       :           0         0 bps
Transmitted:
  Packets      :           0         0 pps
  Bytes       :           0         0 bps
Tail-dropped packets : Not Available

```

RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

Queued:

Packets	:	77059796	486384 pps
Bytes	:	3544750624	178989576 bps

Transmitted:

Packets	:	77059797	486381 pps
Bytes	:	3544750670	178988248 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 1, Forwarding classes: expedited-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 2, Forwarding classes: assured-forwarding

Queued:

Packets	:	4846580	3934 pps
Bytes	:	222942680	1447768 bps

Transmitted:

Packets	:	4846580	3934 pps
Bytes	:	222942680	1447768 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps

```

      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : 0 0 pps
    RED-dropped packets : 0 0 pps
      Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
      High : 0 0 pps
    RED-dropped bytes : 0 0 bps
      Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
      High : 0 0 bps

```

show interfaces queue (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-7/1/3
Physical interface: ge-7/1/3, Enabled, Physical link is Up
  Interface index: 170, SNMP ifIndex: 70 Forwarding classes: 16 supported, 4 in use
  Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets : 418390039 10 pps
    Bytes : 38910269752 7440 bps
  Transmitted:
    Packets : 418390039 10 pps
    Bytes : 38910269752 7440 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps
    RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps
    RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
    Tail-dropped packets : Not Available
    RED-dropped packets : 0 0 pps

```

```

RED-dropped bytes      :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets               :                7055              1 pps
  Bytes                 :            451552              512 bps
Transmitted:
  Packets               :                7055              1 pps
  Bytes                 :            451552              512 bps
Tail-dropped packets : Not Available
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps
Forwarding classes: 16 supported, 4 in use Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
  Packets               :                1031              0 pps
  Bytes                 :            143292              0 bps
Transmitted:
  Packets               :                1031              0 pps
  Bytes                 :            143292              0 bps
Tail-dropped packets : Not Available
RL-dropped packets    :                0                0 pps
RL-dropped bytes      :                0                0 bps
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
  Packets               :                0                0 pps
  Bytes                 :                0                0 bps
Transmitted:
  Packets               :                0                0 pps
  Bytes                 :                0                0 bps
Tail-dropped packets : Not Available
RL-dropped packets    :                0                0 pps
RL-dropped bytes      :                0                0 bps
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
  Packets               :                0                0 pps
  Bytes                 :                0                0 bps
Transmitted:
  Packets               :                0                0 pps
  Bytes                 :                0                0 bps
Tail-dropped packets : Not Available
RL-dropped packets    :                0                0 pps
RL-dropped bytes      :                0                0 bps
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps
Queue: 3, Forwarding classes: network-control
Queued:
  Packets               :                77009             11 pps
  Bytes                 :            6894286             7888 bps
Transmitted:
  Packets               :                77009             11 pps
  Bytes                 :            6894286             7888 bps
Tail-dropped packets : Not Available
RL-dropped packets    :                0                0 pps
RL-dropped bytes      :                0                0 bps
RED-dropped packets   :                0                0 pps
RED-dropped bytes     :                0                0 bps

```

Packet Forwarding Engine Chassis Queues:

Queues: 4 supported, 4 in use

Queue: 0, Forwarding classes: best-effort

Queued:

Packets	:	1031	0 pps
Bytes	:	147328	0 bps

Transmitted:

Packets	:	1031	0 pps
Bytes	:	147328	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low, non-TCP	:	0	0 pps
Low, TCP	:	0	0 pps
High, non-TCP	:	0	0 pps
High, TCP	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low, non-TCP	:	0	0 bps
Low, TCP	:	0	0 bps
High, non-TCP	:	0	0 bps
High, TCP	:	0	0 bps

Queue: 1, Forwarding classes: expedited-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low, non-TCP	:	0	0 pps
Low, TCP	:	0	0 pps
High, non-TCP	:	0	0 pps
High, TCP	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low, non-TCP	:	0	0 bps
Low, TCP	:	0	0 bps
High, non-TCP	:	0	0 bps
High, TCP	:	0	0 bps

Queue: 2, Forwarding classes: assured-forwarding

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low, non-TCP	:	0	0 pps
Low, TCP	:	0	0 pps
High, non-TCP	:	0	0 pps
High, TCP	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low, non-TCP	:	0	0 bps
Low, TCP	:	0	0 bps
High, non-TCP	:	0	0 bps
High, TCP	:	0	0 bps

Queue: 3, Forwarding classes: network-control

Queued:

Packets	:	94386	12 pps
Bytes	:	13756799	9568 bps

Transmitted:

Packets	:	94386	12 pps
Bytes	:	13756799	9568 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low, non-TCP	:	0	0 pps
Low, TCP	:	0	0 pps
High, non-TCP	:	0	0 pps
High, TCP	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low, non-TCP	:	0	0 bps
Low, TCP	:	0	0 bps
High, non-TCP	:	0	0 bps
High, TCP	:	0	0 bps

show interfaces queue both-ingress-egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 both-ingress-egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
  Interface index: 175, SNMP ifIndex: 121
  Forwarding classes: 8 supported, 4 in use
  Ingress queues: 4 supported, 4 in use
  Queue: 0, Forwarding classes: best-effort
    Queued:
      Packets      : Not Available
      Bytes        :                0                0 bps
    Transmitted:
      Packets      :                254                0 pps
      Bytes        :            16274                0 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :                0                0 pps
      RED-dropped bytes  :                0                0 bps
  Queue: 1, Forwarding classes: expedited-forwarding
    Queued:
      Packets      : Not Available
      Bytes        :                0                0 bps
    Transmitted:
      Packets      :                0                0 pps
      Bytes        :                0                0 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :                0                0 pps
      RED-dropped bytes  :                0                0 bps
  Queue: 2, Forwarding classes: assured-forwarding
    Queued:
      Packets      : Not Available
      Bytes        :                0                0 bps
    Transmitted:
      Packets      :                0                0 pps
      Bytes        :                0                0 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :                0                0 pps
      RED-dropped bytes  :                0                0 bps
  Queue: 3, Forwarding classes: network-control
    Queued:
      Packets      : Not Available
      Bytes        :                0                0 bps
    Transmitted:
      Packets      :                0                0 pps
      Bytes        :                0                0 bps
      Tail-dropped packets : Not Available
      RED-dropped packets :                0                0 pps

```

```

    RED-dropped bytes      :                0                0 bps
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets                : Not Available
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                3                0 pps
    Bytes                  :               126                0 bps
    Tail-dropped packets   : Not Available
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets                : Not Available
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
    Tail-dropped packets   : Not Available
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets                : Not Available
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
    Tail-dropped packets   : Not Available
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets                : Not Available
    Bytes                  :                0                0 bps
  Transmitted:
    Packets                :                0                0 pps
    Bytes                  :                0                0 bps
    Tail-dropped packets   : Not Available
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets                :             80564692            0 pps
    Bytes                  :          3383717100            0 bps
  Transmitted:
    Packets                :             80564692            0 pps
    Bytes                  :          3383717100            0 bps
    Tail-dropped packets   :                0                0 pps
    RED-dropped packets    :                0                0 pps
    RED-dropped bytes      :                0                0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets                :             80564685            0 pps
    Bytes                  :          3383716770            0 bps
  Transmitted:
    Packets                :             80564685            0 pps

```

```

Bytes : 3383716770 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : 0 0 pps
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : 9397 0 pps
Bytes : 3809052 232 bps
Transmitted:
Packets : 9397 0 pps
Bytes : 3809052 232 bps
Tail-dropped packets : 0 0 pps
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

show interfaces queue ingress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 ingress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 288 0 pps
Bytes : 18450 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available

```



```

RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

show interfaces queue egress (Gigabit Ethernet IQ2 PIC)

```

user@host> show interfaces queue ge-6/2/0 egress
Physical interface: ge-6/2/0, Enabled, Physical link is Up
Interface index: 175, SNMP ifIndex: 121
Forwarding classes: 8 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 3 0 pps
Bytes : 126 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets : Not Available
Bytes : 0 0 bps
Transmitted:
Packets : 0 0 pps
Bytes : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
RED-dropped bytes : 0 0 bps

```

```

Packet Forwarding Engine Chassis Queues:
Queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
  Transmitted:
    Packets      :      80564692      0 pps
    Bytes        :      3383717100    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
  Transmitted:
    Packets      :      80564685      0 pps
    Bytes        :      3383716770    0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
  Transmitted:
    Packets      :      0      0 pps
    Bytes        :      0      0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets      :      9538      0 pps
    Bytes        :      3819840      0 bps
  Transmitted:
    Packets      :      9538      0 pps
    Bytes        :      3819840      0 bps
    Tail-dropped packets :      0      0 pps
    RED-dropped packets :      0      0 pps
    RED-dropped bytes  :      0      0 bps

```

show interfaces queue remaining-traffic (Gigabit Ethernet Enhanced DPC)

```

user@host> show interfaces queue ge-2/2/9 remaining-traffic
Physical interface: ge-2/2/9, Enabled, Physical link is Up
  Interface index: 238, SNMP ifIndex: 71
Forwarding classes: 16 supported, 4 in use
Ingress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      :      110208969      472875 pps
    Bytes        :      5951284434    204282000 bps
  Transmitted:
    Packets      :      110208969      472875 pps
    Bytes        :      5951284434    204282000 bps
    Tail-dropped packets : Not Available
    RED-dropped packets :      0      0 pps
    Low          :      0      0 pps

```

```

Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps
High            : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
Queued:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Transmitted:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low             : 0 0 pps
Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps
High            : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
Queued:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Transmitted:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low             : 0 0 pps
Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps
High            : 0 0 bps
Queue: 3, Forwarding classes: network-control
Queued:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Transmitted:
Packets         : 0 0 pps
Bytes           : 0 0 bps
Tail-dropped packets : Not Available
RED-dropped packets : 0 0 pps
Low             : 0 0 pps
Medium-low      : 0 0 pps
Medium-high     : 0 0 pps
High            : 0 0 pps
RED-dropped bytes : 0 0 bps
Low             : 0 0 bps
Medium-low      : 0 0 bps
Medium-high     : 0 0 bps

```

```

      High : 0 0 bps
Forwarding classes: 16 supported, 4 in use
Egress queues: 4 supported, 4 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets : 109355853 471736 pps
    Bytes : 7436199152 256627968 bps
  Transmitted:
    Packets : 109355852 471736 pps
    Bytes : 7436198640 256627968 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 2, Forwarding classes: assured-forwarding
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Transmitted:
    Packets : 0 0 pps
    Bytes : 0 0 bps
  Tail-dropped packets : Not Available
  RED-dropped packets : 0 0 pps
    Low : 0 0 pps
    Medium-low : 0 0 pps
    Medium-high : 0 0 pps
    High : 0 0 pps
  RED-dropped bytes : 0 0 bps
    Low : 0 0 bps
    Medium-low : 0 0 bps
    Medium-high : 0 0 bps
    High : 0 0 bps
Queue: 3, Forwarding classes: network-control
  Queued:
    Packets : 0 0 pps
    Bytes : 0 0 bps

```

```

Transmitted:
Packets          :                0                0 pps
Bytes            :                0                0 bps
Tail-dropped packets : Not Available
RED-dropped packets :                0                0 pps
  Low            :                0                0 pps
  Medium-low     :                0                0 pps
  Medium-high    :                0                0 pps
  High           :                0                0 pps
RED-dropped bytes  :                0                0 bps
  Low            :                0                0 bps
  Medium-low     :                0                0 bps
  Medium-high    :                0                0 bps
  High           :                0                0 bps

```

show interfaces queue (Channelized OC12 IQE Type 3 PIC in SONET Mode)

```

user@host> show interfaces queue t3-1/1/0:7
Physical interface: t3-1/1/0:7, Enabled, Physical link is Up

  Interface index: 192, SNMP ifIndex: 1948

  Description: full T3 interface connect to 6ce13 t3-3/1/0:7 for FR testing -
Lam

  Forwarding classes: 16 supported, 9 in use

  Egress queues: 8 supported, 8 in use

  Queue: 0, Forwarding classes: DEFAULT

  Queued:

    Packets          :                214886                13449 pps
    Bytes            :                9884756                5164536 bps

  Transmitted:

    Packets          :                214886                13449 pps
    Bytes            :                9884756                5164536 bps
    Tail-dropped packets :                0                0 pps
    RED-dropped packets :                0                0 pps
      Low            :                0                0 pps
      Medium-low     :                0                0 pps
      Medium-high    :                0                0 pps
      High           :                0                0 pps
    RED-dropped bytes  :                0                0 bps
      Low            :                0                0 bps
      Medium-low     :                0                0 bps

```

Medium-high	:	0	0 bps
-------------	---	---	-------

High	:	0	0 bps
------	---	---	-------

Queue: 1, Forwarding classes: REALTIME

Queued:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

Transmitted:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

Tail-dropped packets	:	0	0 pps
----------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

Low	:	0	0 pps
-----	---	---	-------

Medium-low	:	0	0 pps
------------	---	---	-------

Medium-high	:	0	0 pps
-------------	---	---	-------

High	:	0	0 pps
------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Low	:	0	0 bps
-----	---	---	-------

Medium-low	:	0	0 bps
------------	---	---	-------

Medium-high	:	0	0 bps
-------------	---	---	-------

High	:	0	0 bps
------	---	---	-------

Queue: 2, Forwarding classes: PRIVATE

Queued:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

Transmitted:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

Tail-dropped packets	:	0	0 pps
----------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

Low	:	0	0 pps
-----	---	---	-------

Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 3, Forwarding classes: CONTROL

Queued:

Packets	:	60	0 pps
Bytes	:	4560	0 bps

Transmitted:

Packets	:	60	0 pps
Bytes	:	4560	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 4, Forwarding classes: CLASS_B_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 5, Forwarding classes: CLASS_C_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps

High	:	0	0 bps
------	---	---	-------

Queue: 6, Forwarding classes: CLASS_V_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS_S_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps

High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Packet Forwarding Engine Chassis Queues:

Queues: 8 supported, 8 in use

Queue: 0, Forwarding classes: DEFAULT

Queued:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps

Transmitted:

Packets	:	371365	23620 pps
Bytes	:	15597330	7936368 bps

Tail-dropped packets	:	0	0 pps
----------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

Low	:	0	0 pps
-----	---	---	-------

Medium-low	:	0	0 pps
------------	---	---	-------

Medium-high	:	0	0 pps
-------------	---	---	-------

High	:	0	0 pps
------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Low	:	0	0 bps
-----	---	---	-------

Medium-low	:	0	0 bps
------------	---	---	-------

Medium-high	:	0	0 bps
-------------	---	---	-------

High	:	0	0 bps
------	---	---	-------

Queue: 1, Forwarding classes: REALTIME

Queued:

Packets	:	0	0 pps
---------	---	---	-------

```

Bytes          :          0          0 bps
Transmitted:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
  Low          :          0          0 pps
  Medium-low   :          0          0 pps
  Medium-high  :          0          0 pps
  High         :          0          0 pps
RED-dropped bytes :          0          0 bps
  Low          :          0          0 bps
  Medium-low   :          0          0 bps
  Medium-high  :          0          0 bps
  High         :          0          0 bps
Queue: 2, Forwarding classes: PRIVATE
Queued:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Transmitted:
Packets        :          0          0 pps
Bytes          :          0          0 bps
Tail-dropped packets :          0          0 pps
RED-dropped packets :          0          0 pps
  Low          :          0          0 pps
  Medium-low   :          0          0 pps
  Medium-high  :          0          0 pps
  High         :          0          0 pps
RED-dropped bytes :          0          0 bps
  Low          :          0          0 bps
  Medium-low   :          0          0 bps

```

Medium-high	:	0	0 bps
-------------	---	---	-------

High	:	0	0 bps
------	---	---	-------

Queue: 3, Forwarding classes: CONTROL

Queued:

Packets	:	32843	0 pps
---------	---	-------	-------

Bytes	:	2641754	56 bps
-------	---	---------	--------

Transmitted:

Packets	:	32843	0 pps
---------	---	-------	-------

Bytes	:	2641754	56 bps
-------	---	---------	--------

Tail-dropped packets	:	0	0 pps
----------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

Low	:	0	0 pps
-----	---	---	-------

Medium-low	:	0	0 pps
------------	---	---	-------

Medium-high	:	0	0 pps
-------------	---	---	-------

High	:	0	0 pps
------	---	---	-------

RED-dropped bytes	:	0	0 bps
-------------------	---	---	-------

Low	:	0	0 bps
-----	---	---	-------

Medium-low	:	0	0 bps
------------	---	---	-------

Medium-high	:	0	0 bps
-------------	---	---	-------

High	:	0	0 bps
------	---	---	-------

Queue: 4, Forwarding classes: CLASS_B_OUTPUT

Queued:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

Transmitted:

Packets	:	0	0 pps
---------	---	---	-------

Bytes	:	0	0 bps
-------	---	---	-------

Tail-dropped packets	:	0	0 pps
----------------------	---	---	-------

RED-dropped packets	:	0	0 pps
---------------------	---	---	-------

Low	:	0	0 pps
-----	---	---	-------

Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 5, Forwarding classes: CLASS_C_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 6, Forwarding classes: CLASS_V_OUTPUT

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps
High	:	0	0 bps

Queue: 7, Forwarding classes: CLASS_S_OUTPUT, GETS

Queued:

Packets	:	0	0 pps
Bytes	:	0	0 bps

Transmitted:

Packets	:	0	0 pps
Bytes	:	0	0 bps
Tail-dropped packets	:	0	0 pps
RED-dropped packets	:	0	0 pps
Low	:	0	0 pps
Medium-low	:	0	0 pps
Medium-high	:	0	0 pps
High	:	0	0 pps
RED-dropped bytes	:	0	0 bps
Low	:	0	0 bps
Medium-low	:	0	0 bps
Medium-high	:	0	0 bps

High : 0 0 bps

show interfaces queue (QFX Series)

```

user@switch> show interfaces queue xe-0/0/15
Physical interface: xe-0/0/15, Enabled, Physical link is Up
Interface index: 49165, SNMP ifIndex: 539
Forwarding classes: 12 supported, 8 in use
Egress queues: 12 supported, 8 in use
Queue: 0, Forwarding classes: best-effort
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 3, Forwarding classes: fcoe
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
0 bps
Queue: 4, Forwarding classes: no-loss
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 7, Forwarding classes: network-control
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
    Tail-dropped packets : Not Available
    Total-dropped packets: 0 0 pps
    Total-dropped bytes  : 0 0 bps
Queue: 8, Forwarding classes: mcast
  Queued:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
  Transmitted:
    Packets      : 0 0 pps
    Bytes       : 0 0 bps
    Tail-dropped packets : Not Available

```

Total-dropped packets:	0	0 pps
Total-dropped bytes :	0	0 bps

show interfaces queue l2-statistics (lsq interface)

```

user@switch> show interfaces queue lsq-2/2/0.2 l2-statistics
Logical interface lsq-2/2/0.2 (Index 69) (SNMP ifIndex 1598)
Forwarding classes: 16 supported, 4 in use
Egress queues: 8 supported, 4 in use
Burst size: 0
Queue: 0, Forwarding classes: be
  Queued:
    Packets      :          1          0 pps
    Bytes        :        1001          0 bps
  Transmitted:
    Packets      :          5          0 pps
    Bytes        :        1062          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 1, Forwarding classes: ef
  Queued:
    Packets      :          1          0 pps
    Bytes        :        1500          0 bps
  Transmitted:
    Packets      :          6          0 pps
    Bytes        :        1573          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 2, Forwarding classes: af
  Queued:
    Packets      :          1          0 pps
    Bytes        :          512          0 bps
  Transmitted:
    Packets      :          3          0 pps
    Bytes        :          549          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
Queue: 3, Forwarding classes: nc
  Queued:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
  Transmitted:
    Packets      :          0          0 pps
    Bytes        :          0          0 bps
    Tail-dropped packets :          0          0 pps
    RED-dropped packets :          0          0 pps
    RED-dropped bytes  :          0          0 bps
=====

```


show pfe next-hop

List of Syntax	Syntax on page 685 Syntax (TX Matrix and TX Matrix Plus Routers) on page 685
Syntax	<pre>show pfe next-hop <interface <i>interface-name</i>></pre>
Syntax (TX Matrix and TX Matrix Plus Routers)	<pre>show pfe next-hop <fpc <i>slot</i>> <interface <i>interface-name</i>> <lcc <i>number</i>></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	Display Packet Forwarding Engine next-hop information.
Options	<p>none—Display all Packet Forwarding Engine next-hop information.</p> <p>fpc <i>slot</i>—(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.</p> <ul style="list-style-type: none"> On a TX Matrix router, if you specify the number of a T640 router by using the lcc <i>number</i> option (the recommended method), replace <i>slot</i> with a value from 0 through 7. Otherwise, replace <i>slot</i> with a value from 0 through 31. On a TX Matrix Plus router, if you specify the number of a T1600 router by using the lcc <i>number</i> option (the recommended method), replace <i>slot</i> with a value from 0 through 7. Otherwise, replace <i>slot</i> with a value from 0 through 31. On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 router by using the lcc <i>number</i> option (the recommended method), replace <i>slot</i> with a value from 0 through 7. Otherwise, replace <i>slot</i> with a value from 0 through 63. <p>For example, the following commands have the same result:</p> <pre>user@host> show pfe next-hop fpc 1 lcc 1 user@host> show pfe next-hop fpc 9</pre> <p>interface <i>interface-name</i>—(Optional) Display the Packet Forwarding Engine next-hop interface.</p> <p>lcc <i>number</i>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Packet Forwarding Engine next-hop interface for a specific T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display Packet Forwarding Engine next-hop interface for the router (or line-card chassis) that is connected to a TX Matrix Plus router.</p>

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

Required Privilege Level admin

Related Documentation

- *Routing Matrix with TXP-T1600 Configuration*
- *Routing Matrix with TXP-T1600-3D Configuration*
- *Routing Matrix with TXP-T4000-3D Configuration*
- *Overview of a Routing Matrix with TXP-Mixed-LCC-3D Configuration*

List of Sample Output

[show pfe next-hop on page 687](#)
[show pfe next-hop fpc \(TX Matrix Router\) on page 687](#)
[show pfe next-hop fpc \(TX Matrix Plus Router\) on page 687](#)

Sample Output

show pfe next-hop

```
user@host> show pfe next-hop
NextHop Info:
  ID      Type      Interface      Protocol      Encap      Next Hop Addr      MTU
  ----      -      -      -      -      -      -
    4      Mcast      -              IPv4          -      0.0.0.0              0
    5      Bcast      -              IPv4          -      -                    0
    7      Discard      -              IPv4          -      -                    0
    8      MDiscard      -              IPv4          -      -                    0
    9      Reject      -              IPv4          -      -                    0
   13      Local      -              IPv4          -      192.168.4.60          0
   14      Resolve      fxp0.0          IPv4      Unspecified      -                    0
   17      Local      -              IPv4          -      127.0.0.1              0
   18      Unicast      fxp0.0          IPv4      Unspecified      192.168.4.254          0
   21      Local      -              IPv4          -      11.1.0.1              0
   22      Unicast      at-0/1/0.0      IPv4      ATM SNAP      11.1.0.2              4482
  ...
```

show pfe next-hop fpc (TX Matrix Router)

```
user@host> show pfe next-hop fpc 1
Slot 1
NextHop Info:
  ID      Type      Interface      Next Hop Addr      Protocol      Encap      MTU
  ----      -      -      -      -      -      -
    5      Mcast      -              default            IPv4          -      0
    6      Bcast      -              -                  IPv4          -      0
    8      Discard      -              -                  IPv4          -      0
    9      MDiscard      -              -                  IPv4          -      0
   13      Mcast      -              default            IPV6          -      0
   17      MDiscard      -              -                  IPV6          -      0
   18      Reject      -              -                  IPV6          -      0
   24      Discard      -              -                  None          -      0
   68      Local      -              192.168.66.113      IPv4          -      0
   69      Resolve      fxp0.0          -                  IPv4      Unspecified      0
   70      Unicast      fxp0.0          192.168.71.254      IPv4      Unspecified      0
  256      Local      -              10.71.71.1          IPv4          -      0
  257      Local      -              127.0.0.1           IPv4          -      0
  258      Mcast.local..1      default            IPv4      Unspecified      0
  259      Bcast.local..1      -                  IPv4      Unspecified      0
  261      Discard.local..1      -                  IPv4      Unspecified      0
  262      MDiscard.local..1      -                  IPv4      Unspecified      0
  269      Mcast.local..1      default            IPV6      Unspecified      0
  271      Discard.local..1      -                  IPV6      Unspecified      0
  ...
```


show pfe next-hop fpc (TX Matrix Plus Router)

```
user@host> show pfe next-hop fpc 0
Slot 0
  ID      Type      Interface      Next Hop Addr      Protocol      Encap      MTU
  ----      -      -      -      -      -      -
   31      Mcast      -              default            IPv4          -      0
   32      Bcast      -              -                  IPv4          -      0
   34      Discard      -              -                  IPv4          -      0
   35      MDiscard      -              -                  IPv4          -      0
```

36	Reject	-	-	IPv4	-	0
39	Mcast	-	default	IPv6	-	0
42	Discard	-	-	IPv6	-	0
43	MDiscard	-	-	IPv6	-	0
44	Reject	-	-	IPv6	-	0
49	Receive	-	-	MPLS	-	0
50	Discard	-	-	MPLS	-	0
111	Mcast	.local..1	default	IPv4	Unspecified	0
112	Bcast	.local..1	-	IPv4	Unspecified	0
114	Discard	.local..1	-	IPv4	Unspecified	0
115	MDiscard	.local..1	-	IPv4	Unspecified	0
116	Reject	.local..1	-	IPv4	Unspecified	0
119	Mcast	.local..1	default	IPv6	Unspecified	0
122	Discard	.local..1	-	IPv6	Unspecified	0
123	MDiscard	.local..1	-	IPv6	Unspecified	0
124	Reject	.local..1	-	IPv6	Unspecified	0
191	Mcast	.local..2	default	IPv4	Unspecified	0
192	Bcast	.local..2	-	IPv4	Unspecified	0
194	Discard	.local..2	-	IPv4	Unspecified	0
195	MDiscard	.local..2	-	IPv4	Unspecified	0
196	Reject	.local..2	-	IPv4	Unspecified	0
322	Local	-	10.1.0.5	IPv4	-	0
323	Resolve	bcm0.0	-	IPv4	Unspecified	0
326	Local	-	129.0.0.5	IPv4	-	0
327	Resolve	bcm0.0	-	IPv4	Unspecified	0
328	Local	-	fe80::201:ff:fe01:5	IPv6	-	0
329	Receive	bcm0.0	ff02::1:ff01:5	IPv6	Unspecified	0
330	Receive	bcm0.0	fe80::	IPv6	Unspecified	0
331	Resolve	bcm0.0	-	IPv6	Unspecified	0
332	Local	-	fec0::a:1:0:5	IPv6	-	0
333	Receive	bcm0.0	ff02::1:ff00:5	IPv6	Unspecified	0
334	Receive	bcm0.0	fec0::	IPv6	Unspecified	0
335	Resolve	bcm0.0	-	IPv6	Unspecified	0
348	Local	-	192.168.178.4	IPv4	-	0
349	Resolve	em0.0	-	IPv4	Unspecified	0
350	Unicast	em0.0	192.168.178.126	IPv4	Unspecified	0
357	Local	-	fe80::201:1ff:fe01:5	IPv6	-	0
512	Local	-	10.255.178.11	IPv4	-	0
513	Local	-	127.0.0.1	IPv4	-	0
515	Local	-	abcd::10:255:178:11	IPv6	-	0
516	Local	-	fe80::200:ff:fe00:0	IPv6	-	0
517	Local	-	127.0.0.1	IPv4	-	0
518	Mcast	.local..3	default	IPv4	Unspecified	0
519	Bcast	.local..3	-	IPv4	Unspecified	0
521	Discard	.local..3	-	IPv4	Unspecified	0
522	MDiscard	.local..3	-	IPv4	Unspecified	0
523	Reject	.local..3	-	IPv4	Unspecified	0
531	Mcast	.local..3	default	IPv6	Unspecified	0
533	Discard	.local..3	-	IPv6	Unspecified	0
534	MDiscard	.local..3	-	IPv6	Unspecified	0
535	Reject	.local..3	-	IPv6	Unspecified	0
539	Mgroup	-	-	IPv4	-	0
540	Bcast	ge-15/0/3.0	-	IPv4	Ethernet	0
541	Receive	ge-15/0/3.0	14.2.1.0	IPv4	Ethernet	0
542	Local	-	14.2.1.1	IPv4	-	0
543	Resolve	ge-15/0/3.0	-	IPv4	Ethernet	0
544	Bcast	ge-31/0/4.0	-	IPv4	Ethernet	0

545	Receive	ge-31/0/4.0	14.1.1.0	IPv4	Ethernet	0
546	Local	-	14.1.1.1	IPv4	-	0
547	Resolve	ge-31/0/4.0	-	IPv4	Ethernet	0
548	Unicast	ge-31/0/4.0	14.1.1.2	IPv4	Ethernet	0
549	Unicast	ge-15/0/3.0	14.2.1.2	IPv4	Ethernet	0
550	Bcast	ae1.0	-	IPv4	Ethernet	0
551	Receive	ae1.0	11.1.1.0	IPv4	Ethernet	0
552	Local	-	11.1.1.1	IPv4	-	0
553	Resolve	ae1.0	-	IPv4	Ethernet	0
554	Aggreg.	ae1.0	-	IPv4	Ethernet	0
555	Unicast	ge-23/0/8.0	11.1.1.2	IPv4	Ethernet	0
556	Unicast	ge-7/0/9.0	11.1.1.2	IPv4	Ethernet	0
557	Aggreg.	ae1.0	-	MPLS	Ethernet	0
558	Unicast	ge-23/0/8.0	-	MPLS	Ethernet	0
559	Unicast	ge-7/0/9.0	-	MPLS	Ethernet	0
560	Aggreg.	ae1.0	-	MPLS	Ethernet	0
561	Unicast	ge-23/0/8.0	-	MPLS	Ethernet	0
562	Unicast	ge-7/0/9.0	-	MPLS	Ethernet	0

show pfe route

List of Syntax	Syntax on page 690 Syntax (EX Series Switch and QFX Series) on page 690 Syntax (TX Matrix and TX Matrix Plus Routers) on page 690
Syntax	<pre>show pfe route <<inet6 ip iso> <prefix prefix> <table <table-name> <index index> <prefix prefix>>> <mpls> <summary></pre>
Syntax (EX Series Switch and QFX Series)	<pre>show pfe route <<inet6 ip> <prefix prefix> <table <table-name> <index index> <prefix prefix>>> <mpls> <summary></pre>
Syntax (TX Matrix and TX Matrix Plus Routers)	<pre>show pfe route <fpc slot> <<inet6 ip iso> <prefix prefix> <table <table-name> <index index> <prefix prefix>>> <lcc number> <mpls> <summary></pre>
Release Information	<p>Command introduced before Junos OS Release 7.4.</p> <p>Command introduced in Junos OS Release 9.0 for EX Series switches.</p> <p>Command introduced in Junos OS Release 11.1 for the QFX Series.</p>
Description	<p>Display the routes in the Packet Forwarding Engine forwarding table. The Packet Forwarding Engine forwards packets between input and output interfaces.</p>
<div>  <p>NOTE: The Routing Engine maintains a master copy of the forwarding table. It copies the forwarding table to the Packet Forwarding Engine, which is the part of the router or switch responsible for forwarding packets. To display the routes in the Routing Engine forwarding table, use the <code>show route forwarding table</code> command. For more information, see the CLI Explorer.</p> </div>	
Options	<p>none—Display all Packet Forwarding Engine forwarding table information.</p> <p>fpc slot—(TX Matrix and TX Matrix Plus routers only) (Optional) Show the next hops for a Flexible PIC Concentrator (FPC) slot.</p> <ul style="list-style-type: none"> On a TX Matrix router, if you specify the number of a T640 router by using the lcc number option (the recommended method), replace slot with a value from 0 through 7. Otherwise, replace slot with a value from 0 through 31. On a TX Matrix Plus router, if you specify the number of a T1600 router by using the lcc number option (the recommended method), replace slot with a value from 0 through 7. Otherwise, replace slot with a value from 0 through 31.

- On a TX Matrix Plus router in the TXP-T1600-3D, TXP-T4000-3D, or TXP-Mixed-LCC-3D configuration, if you specify the number of a T1600 or T4000 router by using the **lcc number** option (the recommended method), replace **slot** with a value from 0 through 7. Otherwise, replace **slot** with a value from 0 through 63.

For example, the following commands have the same result:

```
user@host> show pfe route fpc 1 lcc 1
user@host> show pfe route fpc 9
```

index index—(Optional) Display table index.

inet6—(Optional) Display Packet Forwarding Engine IPv6 routes.

ip—(Optional) Display Packet Forwarding Engine IPv4 routes.

iso—(Optional) Display ISO version routing tables.

lcc number—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, the slot number of the T640 router (or line-card chassis) that houses the FPC. On a TX Matrix Plus router, the slot number of the router (line-card chassis) that houses the FPC.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

mpls—(Optional) Display Packet Forwarding Engine MPLS information.

prefix prefix—(Optional) IPv4 or IPv6 prefix for which to show table entries.

summary—(Optional) Display summary of Packet Forwarding Engine information.

table <table-name>—(Optional) Display table information.

Required Privilege Level

admin

Related Documentation

- *Routing Matrix with TXP-T1600 Configuration*
- *Routing Matrix with TXP-T1600-3D Configuration*
- *Routing Matrix with TXP-T4000-3D Configuration*
- *Overview of a Routing Matrix with TXP-Mixed-LCC-3D Configuration*

- List of Sample Output
- [show pfe route ip on page 692](#)
 - [show pfe route iso on page 692](#)
 - [show pfe route lcc summary \(TX Matrix Router\) on page 692](#)
 - [show pfe route lcc summary \(TX Matrix Plus Router\) on page 694](#)

Sample Output

show pfe route ip

```
user@host> show pfe route ip
```

```
IPv4 Route Table 0, default.0, 0x0:
Destination                NH IP Addr      Type      NH ID Interface
-----
default                    127.0.0.1      Discard    8
127.0.0.1                  127.0.0.1      Local      256
172.16/12                  192.168.71.254 Unicast    68 fxp0.0
192.168.0/18               192.168.71.254 Unicast    68 fxp0.0
192.168.40/22              192.168.71.254 Unicast    68 fxp0.0
192.168.64/18              192.168.71.254 Unicast    68 fxp0.0
192.168.64/21              192.168.71.254 Resolve    67 fxp0.0
192.168.71.249             192.168.71.249 Local      66
192.168.220.0/30           192.168.71.249 Resolve    303 fe-0/0/0.0
192.168.220.0              192.168.220.0 Receive    301 fe-0/0/0.0
224.0.0.1                  Mcast         5
255.255.255.255           Bcast         6

...
```

show pfe route iso

```
user@host# show pfe route iso
```

```
CLNS Route Table 0, CLNP.0, 0x0:
Destination                Type      NH ID Interface
-----
default                    Reject    60
47.0005.80ff.f800.0000.0108.0001.0102.5508.2159/152 Local     514
49.0001.00a0.c96b.c491/72 Local     536
```

show pfe route lcc summary (TX Matrix Router)

```
user@host> show pfe route lcc 2 summary
```

```
Slot 0
```

```
IPv4 Route Tables:
Index      Routes      Size(b)
-----
Default    43             3081
1          4              281
```

```
MPLS Route Tables:
Index      Routes      Size(b)
-----
Default    1             68
```

```
IPv6 Route Tables:
Index      Routes      Size(b)
-----
Default    9             717
```


1 5 389

Slot 1

IPv4 Route Tables:

Index	Routes	Size(b)
Default	43	3081
1	4	281

MPLS Route Tables:

Index	Routes	Size(b)
Default	1	68

IPv6 Route Tables:

Index	Routes	Size(b)
Default	9	717
1	5	389

Slot 16

IPv4 Route Tables:

Index	Routes	Size(b)
Default	41	2938
1	4	281

MPLS Route Tables:

Index	Routes	Size(b)
Default	1	68

IPv6 Route Tables:

Index	Routes	Size(b)
Default	9	717
1	5	389

Slot 17

IPv4 Route Tables:

Index	Routes	Size(b)
Default	41	2938
1	4	281

MPLS Route Tables:

Index	Routes	Size(b)
Default	1	68

IPv6 Route Tables:

Index	Routes	Size(b)
Default	9	717
1	5	389

show pfe route lcc summary (TX Matrix Plus Router)

user@host> show pfe route lcc 2 summary

Slot 0

IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	25	2266
1	9	815
2	6	545
3	5	453
4	15	1371
5	5	453
6	13	1187

MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	88
4	5	452

IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	7	697
1	13	1305
3	4	385
4	4	385
5	4	385
6	18	1833

Slot 6

IPv4 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	25	2266
1	9	815
2	6	545
3	5	453
4	15	1371
5	5	453
6	13	1187

MPLS Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	1	88
4	5	452

IPv6 Route Tables:

Index	Routes	Size(b)
-----	-----	-----
Default	7	697
1	13	1305
3	4	385
4	4	385

5	4	385
6	18	1833
...		

show pfe terse

List of Syntax	Syntax on page 696 Syntax (TX Matrix and TX Matrix Plus Router) on page 696 Syntax (MX Series Router) on page 696
Syntax	show pfe terse
Syntax (TX Matrix and TX Matrix Plus Router)	show pfe terse <lcc <i>number</i> scc> <sfc <i>number</i> >
Syntax (MX Series Router)	show pfe terse <all-members> <local> <member <i>member-id</i> >
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 9.0 for EX Series switches. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Packet Forwarding Engine status information.
Options	<p>none—Display brief information about the Packet Forwarding Engine.</p> <p>all-members—(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for all members in the Virtual Chassis configuration.</p> <p>lcc <i>number</i>—(TX Matrix and TX Matrix Plus routers only) (Optional) On a TX Matrix router, display Packet Forwarding Engine information for a T640 router (or line-card chassis) that is connected to a TX Matrix router. On a TX Matrix Plus router, display Packet Forwarding Engine information for the router (or line-card chassis) that is connected to a TX Matrix Plus router.</p> <p>Replace <i>number</i> with the following values depending on the LCC configuration:</p> <ul style="list-style-type: none">• 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.• 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.• 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.• 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix. <p>local—(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the local Virtual Chassis member.</p>

member *member-id*—(MX Series routers only) (Optional) Display Packet Forwarding Engine status information for the specified member of the Virtual Chassis configuration. Replace *member-id* with a value of 0 or 1.

scc—(TX Matrix routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix router (or switch-card chassis).

sfc—(TX Matrix Plus routers only) (Optional) Display Packet Forwarding Engine information for the TX Matrix Plus router (or switch-fabric chassis).

Required Privilege Level admin

List of Sample Output [show pfe terse \(TX Matrix Router\) on page 697](#)
[show pfe terse \(TX Matrix Plus Router\) on page 697](#)
[show pfe terse sfc \(TX Matrix Plus Router\) on page 697](#)

Sample Output

show pfe terse (TX Matrix Router)

```
user@host> show pfe terse
Slot Type Slot State Flags Uptime
0 SFM Present Online 0x0bf 01:25:42
2 SFM Present Online 0x0bf 01:25:40
0 FPC Present Online 0x102 01:25:57
1 FPC Present Online 0x102 01:25:55
2 FPC Present Online 0x102 01:25:53
```

show pfe terse (TX Matrix Plus Router)

```
user@host> show pfe terse
sfc0-re0:
-----
Slot Type Slot State Uptime
0 LCC Present Online 2d 05:26

lcc0-re0:
-----
Slot Type Slot State Uptime
0 GFPC Present Online 2d 05:25
1 GFPC Present Online 2d 05:25
```

show pfe terse sfc (TX Matrix Plus Router)

```
user@host> show pfe terse sfc 0
sfc0-re0:
-----
Slot Type Slot State Uptime
0 LCC Present Online 2d 05:25
```

show pfe version

Syntax	show pfe version <brief detail>
Release Information	Command introduced before Junos OS Release 7.4. Command introduced in Junos OS Release 11.1 for the QFX Series.
Description	Display Packet Forwarding Engine version information.
Options	brief detail —Display the specified level of output.
Required Privilege Level	admin
List of Sample Output	show pfe version brief on page 698 show pfe version detail on page 698

Sample Output

show pfe version brief

```
user@host> show pfe version brief
PFED release 11.1D0 built by builder on 2010-11-11 05:16:11 UTC
```

show pfe version detail

```
user@host> show pfe version detail
PFED release 11.1D0 built by builder on 2010-11-11 05:16:11 UTC

junos-core01.juniper.net:/volume/build/junos/rpd_feb11/11.1/development/20101111.0/obj-i386/
junos/usr.sbin/pfed
```

PART 4

Troubleshooting

- [Troubleshooting Procedures on page 701](#)

CHAPTER 10

Troubleshooting Procedures

- [Troubleshooting Dropped FCoE Traffic on page 701](#)
- [Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth on page 704](#)
- [Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth on page 705](#)
- [Troubleshooting Egress Queue Bandwidth Impacted by Congestion on page 706](#)
- [Troubleshooting an Unexpected Rewrite Value on page 707](#)
- [Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic on page 708](#)

Troubleshooting Dropped FCoE Traffic

Problem **Description:** Fibre Channel over Ethernet (FCoE) traffic for which you want guaranteed delivery is dropped.

Cause There are several possible causes of dropped FCoE traffic (the list numbers of the possible causes correspond to the list numbers of the solutions in the *Solution* section.):

1. Priority-based flow control (PFC) is not enabled on the FCoE priority (IEEE 802.1p code point) in both the input and output stanzas of the congestion notification profile.
2. The FCoE traffic is not classified correctly at the ingress interface. FCoE traffic should either use the default **fcoe** forwarding class and classifier configuration (maps the **fcoe** forwarding class to IEEE 802.1p code point 011) or be mapped to a lossless forwarding class and to the code point enabled for PFC on the input and output interfaces.
3. The congestion notification profile that enables PFC on the FCoE priority is not attached to the interface.
4. The forwarding class set (priority group) used for guaranteed delivery traffic does not include the forwarding class used for FCoE traffic.
5. Insufficient bandwidth has been allocated for the FCoE queue or for the forwarding class set to which the FCoE queue belongs.

6. If you are using Junos OS Release 12.2, the **fcoe** forwarding class has been explicitly configured instead of using the default **fcoe** forwarding class configuration (forwarding-class-to-queue mapping).



.....

NOTE: If you are using Junos OS Release 12.2, use the default forwarding-class-to-queue mapping for the lossless **fcoe** and **no-loss** forwarding classes. If you explicitly configure the lossless forwarding classes, the traffic mapped to those forwarding classes is treated as lossy (best effort) traffic and does *not* receive lossless treatment.

.....

7. If you are using Junos OS Release 12.3 or later and you are not using the default **fcoe** forwarding class configuration, the forwarding class used for FCoE is not configured with the **no-loss** packet drop attribute. In Junos OS 12.3 or later, explicit forwarding classes configurations must include the **no-loss** packet drop attribute to be treated as lossless forwarding classes.

Solution The list numbers of the possible solutions correspond to the list numbers of the causes in the *Cause* section.

1. Check the congestion notification profile (CNP) to see if PFC is enabled on the FCoE priority (the correct IEEE 802.1p code point) on both input and output interfaces. Use the **show class-of-service congestion-notification** operational command to show the code points that are enabled for PFC in each CNP.

If you are using the default configuration, FCoE traffic is mapped to code point 011 (priority 3). In this case, the input stanza of the CNP should show that PFC is enabled on code point 011, and the output stanza should show that priority 011 is mapped to flow control queue 3.

If you explicitly configured a forwarding class for FCoE traffic, ensure that:

- You specified the **no-loss** packet drop attribute in the forwarding class configuration
- The code point mapped to the FCoE forwarding class in the ingress classifier is the code point enabled for PFC in the CNP input stanza
- The code point and output queue used for FCoE traffic are mapped to each other in the CNP output stanza (if you are not using the default priority and queue, you must explicitly configure each output queue that you want to respond to PFC messages)

For example, if you explicitly configure a forwarding class for FCoE traffic that is mapped to output queue 5 and to code point 101 (priority 5), the output of the **show class-of-service congestion-notification** looks like:

```
Name: fcoe_p5_cnp, Index: 12183
Type: Input
Cable Length: 100 m
  Priority  PFC      MRU
  000      Disabled
  001      Disabled
  010      Disabled
  011      Disabled
  100      Disabled
  101      Enabled   2500
  110      Disabled
  111      Disabled
Type: Output
  Priority  Flow-Control-Queues
  101      5
```

2. Use the **show class-of-service classifier type ieee-802.1p** operational command to check if the classifier maps the forwarding class used for FCoE traffic to the correct IEEE 802.1p code point.
3. Ensure that the congestion notification profile and classifier are attached to the correct ingress interface. Use the operational command **show configuration class-of-service interfaces interface-name**.
4. Check that the forwarding class set includes the forwarding class used for FCoE traffic. Use the operational command **show configuration class-of-service forwarding-class-sets** to show the configured priority groups and their forwarding classes.

5. Verify the amount of bandwidth allocated to the queue mapped to the FCoE forwarding class and to the forwarding class set to which the FCoE traffic queue belongs. Use the **show configuration class-of-service schedulers *scheduler-name*** operational command (specify the scheduler for FCoE traffic as the *scheduler-name*) to see the minimum guaranteed bandwidth (**transmit-rate**) and maximum bandwidth (**shaping-rate**) for the queue.

Use the **show configuration class-of-service traffic-control-profiles *traffic-control-profile*** operational command (specify the traffic control profile used for FCoE traffic as the *traffic-control-profile*) to see the minimum guaranteed bandwidth (**guaranteed-rate**) and maximum bandwidth (**shaping-rate**) for the forwarding class set.

6. Delete the explicit FCoE forwarding-class-to-queue mapping so that the system uses the default FCoE forwarding-class-to-queue mapping. Include the **delete forwarding-classes class fcoe queue-num 3** statement at the **[edit class-of-service]** hierarchy level to remove the explicit configuration. The system then uses the default configuration for the FCoE forwarding class and preserves the lossless treatment of FCoE traffic.
7. Use the **show class-of-service forwarding-class** operational command to display the configured forwarding classes. The *No-Loss* column shows whether lossless transport is enabled or disabled for each forwarding class. If the forwarding class used for FCoE traffic is not enabled for lossless transport, include the **no-loss** packet drop attribute in the forwarding class configuration (**set class-of-service forwarding-classes class *fcoe-forwarding-class-name* queue-num *queue-number* no-loss**).

See “[Example: Configuring CoS PFC for FCoE Traffic](#)” on page 241 for step-by-step instructions on how to configure PFC for FCoE traffic, including classifier, interface, congestion notification profile, PFC, and bandwidth scheduling configuration.

Related Documentation

- [show class-of-service congestion-notification on page 562](#)
- [show class-of-service forwarding-class-set on page 570](#)
- [Configuring CoS PFC \(Congestion Notification Profiles\) on page 417](#)
- [Example: Configuring CoS PFC for FCoE Traffic on page 241](#)
- [Overview of CoS Changes Introduced in Junos OS Release 12.2 on page 21](#)
- [Understanding CoS Flow Control \(Ethernet PAUSE and PFC\) on page 160](#)

Troubleshooting Egress Bandwidth That Exceeds the Configured Maximum Bandwidth

Problem **Description:** The maximum bandwidth of a queue when measured at the egress port exceeds the maximum bandwidth (shaping rate) configured for the queue.

Cause When you configure bandwidth for a queue or a priority group, the switch accounts for the configured bandwidth as data only. The switch does not rate-shape the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its maximum bandwidth calculations.

The measured egress bandwidth can exceed the configured maximum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.

Solution When you calculate the bandwidth requirements for queues on which you expect a significant amount of traffic with small packet sizes, consider the shaping rate as the maximum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined maximum data rate (shaping rate) and the preamble and IFG.

If the maximum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to the queue, reduce the shaping rate for that queue.

Related Documentation

- [shaping-rate on page 532](#)
- [Example: Configuring Maximum Output Bandwidth on page 305](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Understanding CoS Output Queue Schedulers on page 86](#)

Troubleshooting Egress Bandwidth That Exceeds the Configured Minimum Bandwidth

Problem **Description:** The minimum bandwidth of a queue or a priority group when measured at the egress port exceeds the minimum bandwidth configured for the queue (transmit-rate) or for the priority group (guaranteed-rate).

Cause When you configure bandwidth for a queue or a priority group, the switch accounts for the configured bandwidth as data only. The switch does not include the preamble and the interframe gap (IFG) associated with frames, so the switch does not account for the bandwidth consumed by the preamble and the IFG in its minimum bandwidth calculations.

The measured egress bandwidth can exceed the configured minimum bandwidth when small packet sizes (64 or 128 bytes) are transmitted because the preamble and the IFG are a larger percentage of the total traffic. For larger packet sizes, the preamble and IFG overhead are a small portion of the total traffic, and the effect on egress bandwidth is minor.



NOTE: The sum of the queue transmit rates in a priority group should not exceed the guaranteed rate for the priority group. (You cannot guarantee a minimum bandwidth for the queues that is greater than the minimum bandwidth guaranteed for the entire set of queues.)

Solution When you calculate the bandwidth requirements for queues and priority groups on which you expect a significant amount of traffic with small packet sizes, consider the transmit

rate and the guaranteed rate as the minimum bandwidth for the data only. Add sufficient bandwidth to your calculations to account for the preamble and IFG so that the port bandwidth is sufficient to handle the combined minimum data rate and the preamble and IFG.

If the minimum bandwidth measured at the egress port exceeds the amount of bandwidth that you want to allocate to a queue or to a priority group, reduce the transmit rate for that queue and reduce the guaranteed rate of the priority group that contains the queue.

- Related Documentation**
- [guaranteed-rate on page 496](#)
 - [transmit-rate on page 539](#)
 - [Example: Configuring Minimum Guaranteed Output Bandwidth on page 300](#)
 - [Example: Configuring Queue Schedulers on page 289](#)
 - [Understanding CoS Output Queue Schedulers on page 86](#)

Troubleshooting Egress Queue Bandwidth Impacted by Congestion

Problem **Description:** Congestion on an egress port causes egress queues to receive less bandwidth than expected. Egress port congestion can impact the amount of bandwidth allocated to queues on the congested port and, in some cases, on ports that are not congested.

Cause Egress queue congestion can cause the ingress port buffer to fill above a certain threshold and affect the flow to the queues on the egress port. One queue receives its configured bandwidth, but the other queues on the egress port are affected and do not receive their configured share of bandwidth.

Solution The solution is to configure a drop profile to apply weighted random early detection (WRED) to the queue or queues on the congested ports.

Configure a drop profile on the queue that is receiving its configured bandwidth. This queue is preventing the other queues from receiving their expected bandwidth. The drop profile prevents the queue from affecting the other queues on the port.

To configure a WRED profile using the CLI:

- Name the drop profile and set the drop start point, drop end point, minimum drop rate, and maximum drop rate for the drop profile:

```
[edit class-of-service]
user@switch# set drop-profile drop-profile-name interpolate fill-level percentage fill-level
percentage drop-probability 0 drop-probability percentage
```

- Related Documentation**
- [drop-profile on page 476](#)
 - [Example: Configuring WRED Drop Profiles on page 279](#)
 - [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)
 - [Understanding CoS WRED Drop Profiles on page 149](#)

Troubleshooting an Unexpected Rewrite Value

Problem **Description:** Traffic from one or more forwarding classes on an egress port is assigned an unexpected rewrite value.



NOTE: For packets that carry both an inner VLAN tag and an outer VLAN tag, the rewrite rules rewrite only the outer VLAN tag.

Cause If you configure a rewrite rule for a forwarding class on an egress port but you do not configure a rewrite rule for every forwarding class on that egress port, then the forwarding classes that do not have a configured rewrite rule are assigned random rewrite values.

For example:

1. Configure forwarding classes **fc1**, **fc2**, and **fc3**.
2. Configure rewrite rules for forwarding classes **fc1** and **fc2**, but not for forwarding class **fc3**.
3. Assign forwarding classes **fc1**, **fc2**, and **fc3** to a port.

When traffic for these forwarding classes flows through the port, traffic for forwarding classes **fc1** and **fc2** is rewritten correctly. However, traffic for forwarding class **fc3** is assigned a random rewrite value.

Solution If any forwarding class on an egress port has a configured rewrite rule, then all forwarding classes on that egress port must have a configured rewrite rule. Configuring a rewrite rule for any forwarding class that is assigned a random rewrite value solves the problem.



TIP: If you want the forwarding class to use the same code point value assigned to it by the ingress classifier, specify that value as the rewrite rule value. For example, if a forwarding class has the IEEE 802.1 ingress classifier code point value 011, configure a rewrite rule for that forwarding class that uses the IEEE 802.1p code point value 011.



NOTE: There are no default rewrite rules. You can bind one rewrite rule for each type (DSCP and IEEE 802.1) to a given interface. A rewrite rule can contain multiple forwarding-class-to-rewrite-value associations.

1. Assign a rewrite value to a forwarding class. Add the new rewrite value to the same rewrite rule as the other forwarding classes on the port:

```
[edit class-of-service rewrite-rules]
```

```
user@switch# set (dscp | ieee-802.1) rewrite-name forwarding-class class-name loss-priority
priority code-point (alias | bits)
```

For example, if the other forwarding classes on the port use rewrite values defined in the rewrite rule **custom-rw**, the forwarding class **fcoe** is being randomly rewritten, and you want to use IEEE 802.1 code point **011** for the **fcoe** forwarding class:

```
[edit class-of-service rewrite-rules]
user@switch# set ieee-802.1 custom-rw forwarding-class fcoe loss-priority high code-point
011
```

2. Enable the rewrite rule on an interface if it is not already enabled on the desired interface:

```
[edit]
user@switch# set class-of-service interfaces interface-name unit unit rewrite-rules (dscp |
ieee-802.1) rewrite-rule-name
```

For example, to enable the rewrite rule **custom-rw** on interface **xe-0/0/24.0**:

```
[edit]
user@switch# set class-of-service interfaces xe-0/0/24 unit 0 rewrite-rules ieee-802.1
custom-rw
```

Related Documentation

- [interfaces on page 508](#)
- [rewrite-rules on page 525](#)
- [Defining CoS Rewrite Rules on page 427](#)
- [Monitoring CoS Rewrite Rules on page 548](#)

Troubleshooting a Port Reset on QFabric Systems When a Queue Stops Transmitting Traffic

Problem	Description: In QFabric systems, if any queue that contains outgoing packets does not transmit packets for 12 consecutive seconds, the port automatically resets.
Cause	Failure of a queue to transmit packets for 12 consecutive seconds may be due to: <ul style="list-style-type: none">• A strict-high priority queue consuming all of the port bandwidth• Several queues consuming all of the port bandwidth• Any queue or port receiving continuous priority-based flow control (PFC) or 802.3x Ethernet PAUSE messages (received PFC and PAUSE messages prevent a queue or a port, respectively, from transmitting packets because of network congestion)• Other conditions that prevent a queue from obtaining port bandwidth for 12 consecutive seconds
Solution	If the cause is a strict-high priority queue or other queues consuming all of the port bandwidth, you can use rate shaping to configure a maximum rate for the queues that are using all of the port bandwidth and preventing other queues from obtaining bandwidth on the port. You configure a maximum rate by creating a scheduler, using a scheduler

map to apply it to a forwarding class (which maps to an output queue), and applying the scheduler map to the port using a forwarding class set and a traffic control profile.

To configure rate shaping using the CLI:

1. Name the existing scheduler or create a scheduler and define the maximum bandwidth as a rate or as a percentage:

```
[edit class-of-service]
user@switch# set schedulers scheduler-name shaping-rate (rate | percent percentage)
```

2. Configure a scheduler map to associate the scheduler with the forwarding class (queue) that is consuming all of the port bandwidth:

```
[edit class-of-service]
user@switch# set scheduler-maps scheduler-map-name forwarding-class forwarding-class-name scheduler scheduler-name
```

3. Associate the scheduler map with a traffic control profile:

```
[edit class-of-service]
user@switch# set traffic-control-profiles traffic-control-profile-name scheduler-map scheduler-map-name
```

4. Associate the traffic control profile (and thus the scheduler map that contains the rate shaping queue scheduler) with a forwarding class set and apply them to the interface that is being reset:

```
[edit class-of-service]
user@switch# set interfaces interface-name forwarding-class-set fc-set-name output-traffic-control-profile traffic-control-profile-name
```

For example, a strict-high priority queue is using all of the bandwidth on interface **shpnode:xe-0/0/10** and preventing other queues from transmitting for 12 consecutive seconds. You decide to set a maximum rate of 7 Gbps on the strict-high priority queue to ensure that at least 3 Gbps of the port bandwidth is available to service other queues.

[Table 141 on page 709](#) shows the topology for this example:

Table 141: Components of the Rate Shaping Troubleshooting Example

Component	Settings
Affected interface	shpnode:xe-0/0/10
Scheduler (strict-high priority scheduler)	Name: shp-sched Shaping rate: 7g Priority: strict-high NOTE: This example assumes that the scheduler already exists and has been configured as strict-high priority, but that rate shaping to prevent the strict-high priority traffic from using all of the port bandwidth has not been applied.
Scheduler map	Name: shp-map Forwarding class to associate with the shp-sched scheduler: strict-high NOTE: This example assumes that a strict-high priority forwarding class has been configured and assigned the name strict-high .

Table 141: Components of the Rate Shaping Troubleshooting Example (*continued*)

Component	Settings
Traffic control profile	Name: shp-tcp NOTE: This example does not describe how to define a complete traffic control profile.
Forwarding class set	Name: shp-pg

To configure the scheduler, map it to the strict-high priority forwarding class, and apply it to interface **shpnode:xe-0/0/10** using the CLI:

1. Specify the scheduler for the strict-high priority queue (**shp-sched**) with a maximum bandwidth of 7 Gbps:

```
[edit class-of-service schedulers]
user@switch# set shp-sched shaping-rate 7g
```

2. Configure a scheduler map (**shp-map**) that associates the scheduler (**shp-sched**) with the forwarding class (**strict-high**):

```
[edit class-of-service scheduler-maps]
user@switch# set shp-map forwarding-class strict-high scheduler shp-sched
```

3. Associate the scheduler map **shp-map** with a traffic control profile (**shp-tcp**):

```
[edit class-of-service traffic-control-profiles]
user@switch# set shp-tcp scheduler-map shp-map
```

4. Associate the traffic control profile **shp-tcp** with a forwarding class set (**shp-pg**) and the affected interface (**shpnode:xe-0/0/10**):

```
[edit class-of-service]
user@switch# set interfaces shpnode:xe-0/0/10 forwarding-class-set shp-pg
output-traffic-control-profile shp-tcp
```

Related Documentation

- [Understanding CoS Output Queue Schedulers on page 86](#)
- [Defining CoS Queue Scheduling Priority on page 414](#)
- [Example: Configuring Queue Schedulers on page 289](#)
- [Example: Configuring Traffic Control Profiles \(Priority Group Scheduling\) on page 298](#)
- [Example: Configuring Forwarding Class Sets on page 286](#)
- [Example: Configuring CoS Hierarchical Port Scheduling \(ETS\) on page 220](#)