



Providing QoS for Triple Play Traffic on E Series Routers, Release 13.1.0



Published: 2012-03-16

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

Copyright © 2012, Juniper Networks, Inc. All rights reserved.

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.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Providing QoS for Triple Play Traffic on E Series Routers, Release 13.1.0

Copyright © 2012, 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	xiii
	Documentation and Release Notes	xiii
	Documentation Conventions	xiii
	Documentation Feedback	xv
	Requesting Technical Support	xv
	Self-Help Online Tools and Resources	xvi
	Opening a Case with JTAC	xvi
Part 1	Overview	
Chapter 1	How JunosE QoS Works	3
	QoS on the E Series Router Overview	3
	QoS Features	4
	QoS Audience	5
	QoS Terms	6
	QoS References	7
Chapter 2	Hardware and Memory Requirements	9
	QoS Platform Considerations	9
	Interface Specifiers	9
	Memory Requirements for Queue and Buffers	10
	Managing System Resources for Nodes and Queues	10
	Scaling Subscribers on the TFA ASIC with QoS	11
Chapter 3	Classifying, Queuing, and Dropping Triple Play Traffic	13
	Traffic Class and Traffic-Class Groups Overview	13
	Best-Effort Forwarding	13
	Traffic-Class Groups Overview	14
	Queuing and Buffer Management Overview	14
	Static Oversubscription	15
	Dynamic Oversubscription	15
	Color-Based Thresholding	15
	Dropping Behavior Overview	16
Chapter 4	QoS Statistics	17
	QoS Statistics Overview	17
	Rate Statistics	17
	Event Statistics	18
	Bulk Statistics Support for QoS Statistics	18

Chapter 5	Scheduling and Shaping Triple Play Traffic	21
	Scheduler Hierarchy Overview	21
	Shaping Rates, Assured Rates, and Relative Weights in a Scheduler Hierarchy	22
	Shared Shaping Overview	23
	How Shared Shaping Works	24
	Active Constituents for Shared Shaping	25
	Simple Shared Shaping Overview	25
	Bandwidth Allocation for Simple Shared Shaping	25
	Simple Shared Shaping on the Best-Effort Scheduler Node	25
	Simple Shared Shaping for Triple-Play Networks	26
	Simple Shared Shaping Algorithm Overview	27
	Simple Shared Shaper Algorithm Calculations	28
	QoS Parameter Overview	29
Chapter 6	Multicast Bandwidth and Byte Adjustment of Triple Play Traffic	31
	Hierarchical QoS Parameters Overview	31
	IP Multicast Bandwidth Adjustment for QoS Overview	31
	Byte Adjustment for ADSL and VDSL Traffic Overview	33
	Byte Adjustment for Cell Shaping of ADSL Traffic Overview	33
	Calculation and Example of Byte Adjustment for Cell Shaping	33
	Byte Adjustment for Frame Shaping of VDSL Traffic Overview	35
	System Calculation for Byte Adjustment of ADSL and VDSL Traffic	35
Chapter 7	Ethernet Interface Solutions for Triple Play Traffic	37
	Providing QoS for Ethernet Overview	37
	QoS Shaping Mode for Ethernet Interfaces Overview	37
	QoS for 802.3ad Link Aggregation Interfaces Overview	39
	Types of Load Balancing	39
	Munged QoS Profiles and Load Balancing	39
	802.3ad Link Aggregation and QoS Parameters	40
	QoS and Ethernet Link Redundancy	40
	Active Link Failure and QoS	40
	Administratively Disabling a Link and QoS	40
	Adding a New Link to the LAG and QoS	40
	Hashed Load Balancing for 802.3ad Link Aggregation Groups Overview	41
	Sample Scheduler Hierarchy for Hashed Load Balancing	41
	Subscriber Load Balancing for 802.3ad Link Aggregation Groups Overview	42
	S-VLANs and Subscriber Load Balancing	42
	PPPoE over VLANs and Subscriber Load Balancing	42
	PPPoE over Ethernet (No VLANs) and Subscriber Load Balancing	42
	MPLS over LAG and Subscriber Load Balancing	43
	Sample Scheduler Hierarchy for Subscriber Load Balancing	43
	Subscriber Allocation in 802.3ad Link Aggregation Groups	44
Part 2	Configuration	
Chapter 8	Configuration Overview	49
	Configuring QoS on the E Series Router	49

Chapter 9	Configuration Tasks for Classifying, Queuing, and Dropping Triple Play Traffic	51
	Configuring Traffic Classes That Define Service Levels	51
	Configuring Traffic-Class Groups That Define Service Levels	52
	Configuring Queue Profiles to Manage Buffers and Thresholds	53
	Configuring WRED	54
Chapter 10	Configuration Tasks for Scheduling and Shaping Triple Play Traffic	57
	Configuring a Scheduler Hierarchy	57
	Configuring Strict-Priority Scheduling	58
	Configuring Relative Strict-Priority Scheduling for Aggregate Shaping Rates	59
	Configuring a QoS Profile	61
	Attaching a QoS Profile to an Interface	62
	Attaching a QoS Profile to a Base Interface	62
	Attaching a QoS Profile to an ATM VP	63
	Attaching a QoS Profile to an S-VLAN	64
	Attaching a QoS Profile to a Port Type	64
Chapter 11	Configuration Tasks for QoS on Ethernet Interfaces	67
	Creating a QoS Interface Hierarchy for Bulk-Configured VLAN Subinterfaces with RADIUS	67
	Configuring the Scheduler Hierarchy for Hashed Load Balancing in 802.3ad Link Aggregation Groups	70
	Enabling Default Subscriber Load Balancing for 802.3ad Link Aggregation Groups	70
	Configuring the Scheduler Hierarchy for Subscriber Load Balancing in 802.3ad Link Aggregation Groups	71
	Configuring Load Rebalancing for 802.3ad Link Aggregation Groups	71
	Configuring Load-Rebalancing Parameters	72
	Configuring the System to Dynamically Rebalance the LAG	73
Chapter 12	Configuration Tasks for Hierarchical and Multicast Bandwidth Adjustment	75
	Configuring a Basic Parameter Definition for QoS Administrators	75
	Configuring a Parameter Definition to Calculate Hierarchical Instances	77
	Configuring a Parameter Definition for IP Multicast Bandwidth Adjustment	77
	Configuring a Parameter Definition to Shape Ethernet Traffic Using Cell Mode	79
	Configuring a Parameter Definition for QoS Downstream Rate	80
	Creating Parameter Instances	82
	Creating a Global Parameter Instance	82
	Creating a Parameter Instance for an Interface	82
	Creating a Parameter Instance for an ATM VP	83
	Creating a Parameter Instance for an S-VLAN	83
Chapter 13	Configuration Tasks for Statistics Collection	85
	Configuring Statistic Profiles for QoS	85
	Configuring Rate Statistics	85
	Configuring Event Statistics	86

Chapter 14	Examples	89
	Example: QoS Parameter Configuration for Controlling Subscriber	
	Bandwidth	89
	Procedure for QoS Administrators	90
	Procedure for QoS Clients	95
	Monitoring the Subscriber Configuration	97
	Complete Configuration Example	101
	QoS Administrator Configuration	101
	QoS Client Configuration	103
	Example: QoS Parameter Configuration for IP Multicast Bandwidth	
	Adjustment	104
	Monitoring the Configuration	107
	Complete Configuration Example	109
	Example: QoS Parameter Configuration for QoS Downstream Rate	110
	Complete Configuration Example	114
	Example: QoS Parameter Configuration for Hierarchical Parameters	115
	Procedure for QoS Administrators	116
	Procedure for QoS Clients	117
	Monitoring Hierarchical QoS Parameters	118
	Complete Configuration Example	119
	QoS Administrator Configuration	119
	QoS Client Configuration	119
Chapter 15	Configuration Commands	121
	average-length-exponent	122
	buffer-weight	123
	committed-length	124
	committed-threshold	125
	conformed-fraction	126
	conformed-length	127
	conformed-threshold	128
	controlled-interface-type	129
	drop-profile	130
	exceeded-fraction	131
	exceeded-length	132
	exceeded-threshold	133
	fabric-strict-priority	134
	fabric-weight	135
	instance-interface-type	136
	load-rebalance	137
	member-interface	138
	node	139
	qos-parameter	140
	qos-parameter-define	142
	qos-port-type-profile	144
	qos-profile	145
	qos-shaping-mode	146
	queue	147
	queue-profile	148

	scheduler-profile	149
	shaping-rate	150
	shared-shaping-constituent	152
	shared-shaping-rate	153
	traffic-class	155
	traffic-class-group	156
Part 3	Administration	
Chapter 16	Monitoring QoS Statistics	159
	Monitoring Queue Statistics for the Fabric	159
	Clearing QoS Statistics on the Egress Queue	160
	Clearing QoS Statistics on the Fabric Queue	160
Chapter 17	Monitoring QoS on Interfaces	161
	Monitoring the QoS Configuration of Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet Interfaces	161
	Monitoring the QoS Configuration of IEEE 802.3ad Link Aggregation Group Bundles	163
Chapter 18	Monitoring Traffic Classification and Queuing	165
	Monitoring Service Levels with Traffic Classes	165
	Monitoring Service Levels with Traffic-Class Groups	166
	Monitoring Queues and Buffers	167
	Monitoring Queue Thresholds	167
Chapter 19	Monitoring Scheduling	171
	Monitoring the Configuration of Scheduler Profiles	171
	Monitoring Shared Shapers	172
	Monitoring the QoS Scheduler Hierarchy	174
	Monitoring the Configuration of QoS Profiles	180
	Monitoring the QoS Profiles Attached to an Interface	182
	Monitoring a Scheduler Hierarchy on an Interface with QoS Profiles	184
Chapter 20	Monitoring QoS Parameters	185
	Monitoring QoS Parameter Definitions	185
	Monitoring QoS Parameter Instances	186
Chapter 21	Statistics-Gathering Commands	189
	clear egress-queue	190
	clear fabric-queue	191
	committed-drop-threshold	192
	conformed-drop-threshold	193
	exceeded-drop-threshold	194
	forwarding-rate-threshold	195
	rate-period	196
	statistics-profile	197
Chapter 22	Monitoring Commands	199
	show aaa qos downstream-rate	200
	show drop-profile	201
	show interfaces	202

	show qos interface-hierarchy	203
	show qos-parameter	204
	show qos-parameter-define	206
	show qos-profile	207
	show qos queue-thresholds	208
	show qos scheduler-hierarchy	209
	show qos shared-shaper	211
	show scheduler-profile	213
	show statistics-profile	214
	show traffic-class	215
	show traffic-class-group	216
Part 4	Troubleshooting	
Chapter 23	Troubleshooting Egress Queues	219
	Troubleshooting Memory and Processor Use for Egress Queue Rate Statistics and Events	219
Part 5	Index	
	Index	223

List of Figures

Part 1	Overview	
Chapter 2	Hardware and Memory Requirements	9
	Figure 1: Traffic Flow Through an E Series Router	9
Chapter 5	Scheduling and Shaping Triple Play Traffic	21
	Figure 2: QoS Scheduler Hierarchy	22
	Figure 3: Simple Shared Shaping over ATM	26
	Figure 4: Simple Shared Shaping over Ethernet	27
	Figure 5: Simple Shared Shaper Behavior Without Algorithm Controls	27
	Figure 6: Less Conservative Simple Shared Shaper Behavior	28
	Figure 7: More Liberal Simple Shared Shaper Behavior	28
Chapter 6	Multicast Bandwidth and Byte Adjustment of Triple Play Traffic	31
	Figure 8: Byte Adjustment Calculation for Ethernet and ATM Encapsulations	34
Chapter 7	Ethernet Interface Solutions for Triple Play Traffic	37
	Figure 9: 802.3ad Link Aggregation Scheduler Hierarchy	41
	Figure 10: Subscriber LoadBalanced Scheduler Hierarchy for Port 0	43
	Figure 11: Subscriber LoadBalanced Scheduler Hierarchy for Port 1	44
	Figure 12: Subscriber Allocation and Load Balancing	44
Part 2	Configuration	
Chapter 10	Configuration Tasks for Scheduling and Shaping Triple Play Traffic	57
	Figure 13: Sample Strict-Priority Scheduling Hierarchy	59
	Figure 14: Sample Relative Strict-Priority Scheduler Hierarchy	60
Chapter 14	Examples	89
	Figure 15: Physical Network Topology	89
	Figure 16: QoS Scheduler Hierarchy	90
	Figure 17: Scheduler Hierarchy with QoS Adjustment for IP Multicast	104
	Figure 18: Hierarchical Parameters Scheduler Hierarchy	116

List of Tables

	About the Documentation	xiii
	Table 1: Notice Icons	xiv
	Table 2: Text and Syntax Conventions	xiv
Part 1	Overview	
Chapter 1	How JunosE QoS Works	3
	Table 3: QoS Features	4
	Table 4: QoS Terminology	6
Chapter 2	Hardware and Memory Requirements	9
	Table 5: Egress Memory and Region Size on ASIC Line Modules	10
Chapter 5	Scheduling and Shaping Triple Play Traffic	21
	Table 6: Comparison of Simple and Compound Shared Shaping	24
Chapter 6	Multicast Bandwidth and Byte Adjustment of Triple Play Traffic	31
	Table 7: Header Lengths for Ethernet Encapsulation	34
	Table 8: Header Lengths for ATM Encapsulation	34
	Table 9: Byte Adjustment Values for Frame and Cell Shaping Modes	35
Chapter 7	Ethernet Interface Solutions for Triple Play Traffic	37
	Table 10: Operational Shaping Modes	38
Part 2	Configuration	
Chapter 11	Configuration Tasks for QoS on Ethernet Interfaces	67
	Table 11: Load Balancing Algorithm Parameters	72
Chapter 14	Examples	89
	Table 12: Shaping Rate and Shaping Mode	110
Part 3	Administration	
Chapter 16	Monitoring QoS Statistics	159
	Table 13: show fabric-queue Output Fields	159
Chapter 17	Monitoring QoS on Interfaces	161
	Table 14: show interfaces Output Fields	162
	Table 15: show interfaces lag members Output Fields	163
Chapter 18	Monitoring Traffic Classification and Queuing	165
	Table 16: show traffic-class Output Fields	166
	Table 17: show traffic-class-group Output Fields	167

	Table 18: show qos queue-thresholds Output Fields	170
Chapter 19	Monitoring Scheduling	171
	Table 19: show scheduler-profile Output Fields	172
	Table 20: show qos shared-shaper Output Fields	173
	Table 21: show qos scheduler-hierarchy Output Fields	179
	Table 22: show qos-profile Output Fields	181
	Table 23: show qos interface-hierarchy Output Fields	183
Chapter 20	Monitoring QoS Parameters	185
	Table 24: show qos-parameter-define Output Fields	185
	Table 25: show qos-parameter Output Fields	187

About the Documentation

- Documentation and Release Notes on page xiii
- Documentation Conventions on page xiii
- Documentation Feedback on page xv
- Requesting Technical Support on page xv

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>.

Documentation Conventions

Table 1 on page xiv 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.

Table 2 on page xiv defines text conventions used in this guide and the command syntax conventions used primarily in the *JunosE Command Reference Guide*. For more information about command syntax, see the *JunosE System Basics Configuration Guide, Chapter 2, Command Line Interface*.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Text Conventions		
Bold text like this	Represents names of commands and keywords in text.	<ul style="list-style-type: none"> Issue the clock source command. Specify the keyword exp-msg.
Bold text like this	Represents text that you must type.	host1(config)#traffic class low-loss1
Fixed-width text like this	Represents output on the terminal screen.	<pre>host1#show ip community list Community standard list 1 permit 0:100 0:200 0:300</pre>
<i>Italic text like this</i>	<ul style="list-style-type: none"> Emphasizes words. Identifies variables. Identifies chapter, appendix, and book names. 	<ul style="list-style-type: none"> There are two levels of access, <i>user</i> and <i>privileged</i>. <i>clusterId</i>, <i>ipAddress</i>. <i>Appendix A, System Specifications</i>.
+ (plus sign) linking key names	Indicates that you must press two or more keys simultaneously.	Press Ctrl+b.
Syntax Conventions in the Command Reference Guide		
Plain text like this	Represents keywords.	<code>terminal length</code>
<i>Italic text like this</i>	Represents variables.	<i>mask</i> , <i>accessListName</i>

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

Convention	Description	Examples
(pipe symbol)	Indicates a choice between the keywords or variables on either side of the symbol. (Specifying a choice can be either optional or required.)	diagnostic line
[] (square brackets)	Enclose an optional choice of a single keyword or variable.	[internal external]
[]* (square brackets and the asterisk)	Enclose an optional choice of one or more keywords or variables.	[level1 level2 l1]*
{ } (braces)	Enclose a required choice of a single keyword or variable.	{ permit deny } { in out } { clusterId ipAddress }

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release 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: <https://www.juniper.net/alerts/>
- 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

- [How JunosE QoS Works on page 3](#)
- [Hardware and Memory Requirements on page 9](#)
- [Classifying, Queuing, and Dropping Triple Play Traffic on page 13](#)
- [QoS Statistics on page 17](#)
- [Scheduling and Shaping Triple Play Traffic on page 21](#)
- [Multicast Bandwidth and Byte Adjustment of Triple Play Traffic on page 31](#)
- [Ethernet Interface Solutions for Triple Play Traffic on page 37](#)

CHAPTER 1

How JunosE QoS Works

- [QoS on the E Series Router Overview on page 3](#)
- [QoS Features on page 4](#)
- [QoS Audience on page 5](#)
- [QoS Terms on page 6](#)
- [QoS References on page 7](#)

QoS on the E Series Router Overview

QoS is a suite of features that configure queuing and scheduling on the forwarding path of the Juniper Networks E Series Broadband Services Routers. QoS provides a level of predictability and control beyond the best-effort delivery that the router provides by default. Best-effort service provides packet transmission with no assurance of reliability, delay, jitter, or throughput.

QoS as developed for E Series routers conforms to the IETF Differentiated Services (DiffServ) model (RFCs 2597 and 2598). DiffServ networks classify packets into one of a small number of aggregated flows or traffic classes for which you can configure different QoS characteristics. The Juniper Networks QoS architecture extends DiffServ to support edge features such as high-density queuing.

The E Series router supports:

- IETF architecture for differentiated services
- Assured forwarding per-hop-behavior (PHB) groups
- Expedited forwarding PHB groups

The router supports configurable queuing and scheduling. It has an application-specific integrated circuit (ASIC) scheduler that supports thousands of queues in a hierarchical round-robin (HRR) scheduler. The scheduler allows the router to allocate separate queues for each forwarding interface. Separate queues enable fair access to buffers and bandwidth for each subscriber connected to the router.

Allocating queues per interface allows an Internet service provider (ISP) to shape an individual subscriber's traffic flows to specified rates independent of the underlying Layer 2 network type.

- Related Documentation**
- [Configuring QoS on the E Series Router on page 49](#)

QoS Features

Table 3 on page 4 describes the major QoS features supported on the E Series router.

Table 3: QoS Features

Feature	Description
Best effort	Default traffic class for packets being forwarded across the device. Packets that are not assigned to a specific traffic class are assigned to the best-effort traffic class.
Differentiated services	<ul style="list-style-type: none"> • Assured forwarding—See RFC 2597. • Expedited forwarding—See RFC 2598.
Drop profile	Template that specifies active queue management in the form of WRED behavior of an egress queue.
Port shaping	Shapes the aggregate traffic through a port or channel to a rate that is less than the line or port rate.
QoS parameters	Creates a queuing architecture without the numeric subscriber rates and weights in scheduler profiles. You then use the same QoS and scheduler profiles across all subscribers who use the same services but at different bandwidths, reducing the total number of QoS profiles and scheduler profiles required.
QoS port-type profile	QoS profile that is automatically attached to ports of the corresponding type if you do not explicitly attach a QoS profile.
QoS profile	Collection of QoS commands that specify queue profiles, drop profiles, scheduler profiles, and statistics profiles in combination with interface types.
Queue profile	Template that specifies the buffering and tail-dropping behavior of an egress queue.
Rate shaping	<p>Mechanism that throttles the rate at which an interface can transmit packets.</p> <p><i>Note:</i> Rate shaping as presented in policy management in releases before JunosE Release 4.0 is deprecated and converted to QoS profiles and scheduler profiles.</p>
Relative strict-priority scheduling	Provides strict-priority scheduling within a shaped aggregate rate. For example, it lets you provide 1 Mbps of aggregate bandwidth to a subscriber, with up to 500 Kbps of the bandwidth for low-latency traffic. If there is no strict-priority traffic, the low-latency traffic can use up to the full aggregate rate of 1 Mbps.
Scheduler profile	Configures the bandwidth at which queues drain as a function of relative weight, assured rate, and shaping rate.

Table 3: QoS Features (*continued*)

Feature	Description
Shared rate shaping	Mechanism for shaping a logical interface's aggregate traffic to a rate when the traffic for that logical interface is queued through more than one scheduler hierarchy.
Statistics profile	Template that specifies rate statistics and event-gathering characteristics.
Strict-priority scheduling	Designates the traffic class (queue) that receives top priority for transmission of its packets through a port. It is implemented with a special strict-priority scheduler node that is stacked directly above the port.
Traffic class	<p>A chassis-wide grouping of queues and buffers that support transmission of a designated set of traffic across the chassis, from ingress line module, through the switch fabric, and onto the egress line module.</p> <p>The router supports up to eight traffic classes, and therefore up to eight queues per logical interface.</p>
Traffic-class group	<p>Separate hierarchy of scheduler nodes and queues over a port. A traffic-class group uses one level of the scheduler hierarchy, level 1.</p> <p>Traffic classes belong to the default group unless they are specifically assigned to a named group. All queues are stacked in a single scheduler hierarchy above the physical port. When you configure a traffic class inside a group, its queues are stacked separately. The most common reason for creating separate scheduler hierarchies is to implement strict priority scheduling for all queues in the group.</p> <p>The router supports up to four traffic-class groups. A traffic class cannot belong to more than one group.</p>
WRED	Signals end-to-end protocols such as TCP that the router is becoming congested along a particular egress path. The intent is to trigger TCP congestion avoidance in a random set of TCP flows before congestion becomes severe and causes tail dropping on a large number of flows.

QoS Audience

This topic collection contains configuration information for two types of QoS users: QoS administrators and QoS clients.

QoS administrators are responsible for implementing a QoS queuing architecture by defining drop profiles, queue profiles, scheduler profiles, QoS profiles, and QoS parameter definitions.

QoS clients are responsible for configuring services for individual subscribers by creating parameter instances. The parameter instances that QoS clients can create depend on the settings defined in parameter definitions by the QoS administrator.

Related Documentation

- QoS Parameter Audience

QoS Terms

Table 4 on page 6 defines terms used in this discussion of QoS.

Table 4: QoS Terminology

Term	Description
Assured rate	Bandwidth guaranteed until the node below in the scheduler hierarchy is oversubscribed.
Best effort	Network forwards as many packets as possible in as reasonable a time as possible. This is the default per-hop behavior (PHB) for packet transmission.
Best-effort queue	For a logical interface, the queue associated with the best-effort traffic class for that logical interface,
Best-effort scheduler node	The scheduler node associated with a logical interface and traffic class group pair, and where the traffic class group contains the best-effort traffic class. Also known as best-effort node.
CDV	Cell delay variation. Measures the difference between a cell's expected and actual transfer delay. Determines the amount of jitter.
CDVT	Cell delay variation tolerance. Specifies the acceptable tolerance of CDV (jitter).
Effective weight	The result of a weight or an assured rate. Users configure the scheduler node by specifying either an assured rate or a weight within a scheduler profile. An assured rate, in bits per second, is translated into a weight. The resultant weight is referred to as an effective weight.
Group node	A scheduler node associated with a {port interface, traffic-class group} pair. Because the logical interface is the port, only one such scheduler node can exist for each traffic-class group above the port. This node aggregates all traffic for traffic classes in the group.
HAR	Hierarchical assured rate. Dynamically adjusts bandwidth for scheduler nodes.
HRR	Hierarchical round-robin. Allocates bandwidth to queues in proportion to their weights.
Latency	Delay in the transmission of a packet through a network from beginning to end.
Proprietary QoS Management Information Base (MIB)	Supported on the E Series router.

Table 4: QoS Terminology (*continued*)

Term	Description
Queue	First-in-first-out (FIFO) set of buffers that control packets on the data path.
QoS port-type profile	Supplies the QoS information for forwarding interfaces stacked above ports of the associated interface type.
QoS profile attachment	Applies the rules in the QoS profile to a specific interface.
Rate shaping	Allows you to throttle a queue to a specified rate.
RED	Random early detection congestion avoidance technique.
Scheduler hierarchy	A hierarchical, tree-like arrangement of scheduler nodes and queues. The router supports up to three levels of scheduler nodes stacked above a port. The port scheduler is at level 0, with two levels of scheduler nodes at levels 1 and 2. A final level of queues is stacked above the nodes.
Scheduler node	An element within the hierarchical scheduler that implements bandwidth controls for a group of queues. Queues are stacked above scheduler nodes in a hierarchy. The root node is associated with a channel or physical port.
Shaping rate	Bandwidth in a queue or node can be throttled to a specified rate.
Shared shaper constituent	All nodes and queues that are associated with a logical interface that is being shared shaped are considered potential constituents of the shared shaper.
Weight	Specifies the relative weight for queues in the traffic class.
WRED	Weighted random early detection congestion avoidance technique.

QoS References

For more information about QoS, see the following resources:

- RFC 2474—Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers (December 1998)
- RFC 2475—An Architecture for Differentiated Services (December 1998)
- RFC 2597—Assured Forwarding PHB Group (June 1999)
- RFC 2598—An Expedited Forwarding PHB (June 1999)
- RFC 2698—A Two Rate Three Color Marker (September 1999)
- RFC 2990—Next Steps for the IP QoS Architecture (November 2000)

- RFC 2998—A Framework for Integrated Services Operation over Diffserv Networks (November 2000)
- RFC 3246—An Expedited Forwarding PHB (Per-Hop Behavior) (March 2002)
- RFC 3260—New Terminology and Clarifications for Diffserv (April 2002)
- DSL Forum Technical Report (TR)-059—DSL Evolution - Architecture Requirements for the Support of QoS-Enabled IP Services
- Floyd, S., and Jacobson, V. Random Early Detection for Congestion Avoidance. IEEE/ACM Transactions on Networking 1(4), August 1993

CHAPTER 2

Hardware and Memory Requirements


- [QoS Platform Considerations on page 9](#)
- [Memory Requirements for Queue and Buffers on page 10](#)
- [Managing System Resources for Nodes and Queues on page 10](#)
- [Scaling Subscribers on the TFA ASIC with QoS on page 11](#)

QoS Platform Considerations

QoS is supported on all E Series line modules except for the ES2 10G Uplink LM.

[Figure 1 on page 9](#) shows the traffic flow through the router.

Figure 1: Traffic Flow Through an E Series Router



For information about the modules supported on E Series routers:

- See the *ERX Module Guide* for modules supported on ERX7xx models, ERX14xx models, and the Juniper Networks ERX310 Broadband Services Router.
- See the *E120 and E320 Module Guide* for modules supported on the Juniper Networks E120 and E320 Broadband Services Routers.

Interface Specifiers

The majority of the configuration task examples in this topic collection use the *slot/port* format to specify an interface. However, the interface specifier format that you use depends on the router that you are using.

For ERX7xx models, ERX14xx models, and ERX310 routers, use the *slot/port* format. For example, the following command specifies an ATM interface on slot 0, port 1 of an ERX7xx model, ERX14xx model, or ERX310 router.

```
host1(config)#interface gigabitEthernet 0/1
```

For E120 and E320 routers, use the *slot/adapter/port* format, which includes an identifier for the bay in which the I/O adapter (IOA) resides. In the software, adapter 0 identifies the right IOA bay (E120 router) and the upper IOA bay (E320 router); adapter 1 identifies the left IOA bay (E120 router) and the lower IOA bay (E320 router). For example, the following command specifies a 10-Gigabit Ethernet interface on slot 5, adapter 0, port 0 of an E320 router.

```
host1(config)#interface tenGigabitEthernet 5/0/0
```

Related Documentation

- Interface Types and Specifiers.

Memory Requirements for Queue and Buffers

JunosE Software uses 128-byte buffers.

The egress memory available for queues available depends on the ASIC and the line module. [Table 5 on page 10](#) lists the egress memory.

Table 5: Egress Memory and Region Size on ASIC Line Modules

ASIC	Line Module	Egress Memory (MB)	Region Size (MB)
EFA	All EFA line modules	32	4
FFA	GE-2 and GE-HDE	64	8
	OC48	128	16
	ES2 4G LM	128	16
TFA	ES2 10G LM	96	12

Related Documentation

- Guidelines for Managing Queue Thresholds
- Guidelines for Managing Buffers
- *ERX Module Guide* and the *E120 and E320 Module Guide*

Managing System Resources for Nodes and Queues

The type of ASIC that each line module uses determines the system resources for nodes and queues.

Line modules with the EFA ASIC hardware provide 85,000 descriptors that are shared between all nodes and queues. Each line module supports a maximum of 49,000 nodes or queues per line module.

Line modules with the FFA ASIC hardware provide 2000 level 1 nodes or queues and 64,000 level 2 nodes or queues. The ES2 4G LM provides 2000 level 1 nodes or queues

and 128,000 level 2 nodes or queues. The router implicitly creates the level 2 node. Each line module supports a maximum of 64,000 nodes or queues per line module.

Line modules with the TFA ASIC hardware provide 96,000 descriptors that are shared between all nodes and queues. Each line module supports a maximum of 64,000 nodes or queues.

**Related
Documentation**

- [Scaling Subscribers on the TFA ASIC with QoS on page 11](#)
- [Managing System Resources for Shadow Nodes](#)
- [Memory Requirements for Queue and Buffers on page 10](#)
- *ERX Module Guide* and the *E120 and E320 Module Guide*

Scaling Subscribers on the TFA ASIC with QoS

The TFA ASIC on the ES2 10G LM supports a total of 32,000 nodes; however, it requires that each queue stack above a node at both level 1 and level 2, and it cannot skip a level in the scheduler hierarchy. The FFA ASIC also requires that each queue stack above a node at both level 1 and 2, but it also offers more nodes, so the scheduler hierarchy requirement is not as visible. The EFA ASIC does not require queues to stack above any level.

Because the TFA ASIC cannot skip a level in the hierarchy and also offers a smaller amount of nodes, scaling subscribers for triple-play configurations can exhaust node resources. For example, the ethernet-default QoS profile specifies both an IP and a VLAN node. Configuring 16,000 IP over VLAN subinterfaces consumes all 32,000 nodes, with no node resources remaining for other traffic-class groups. By carefully configuring queues on the TFA ASIC, you can scale up to 16,000 subscribers for multiple traffic-class groups in a triple-play configuration.

To conserve nodes on the TFA ASIC, you could apply one of the following configurations:

- If the configuration includes IP and VLANs, you can configure shapers within those queues to control service throughout. For example, in a triple-play environment with voice, video, and data service, you might want to limit the overall rate of traffic using a shared shaper.

At the same time, you might want to individually restrict the maximum rate of each of the services. To conserve node usage, attach shapers to the queue for each service, and attach the shared shaper to the best-effort queue. These queues must be at level 3 in the scheduler hierarchy. Typically, aggregation nodes such as an S-VLAN are placed at level 2. The VLAN queues then feed in to the S-VLAN nodes, which then feed to the level 1 nodes below.

If you do not create a QoS hierarchy with queues at level 3, the system adds phantom nodes to enforce this requirement. To display the hierarchy that is created for the subscriber on the line module, issue the **show qos scheduler-hierarchy** command.

- If the configuration includes S-VLANs, you could configure S-VLAN nodes in the default traffic-class group. Combining S-VLAN and VLAN nodes uses fewer resources than

when you combine IP and VLAN nodes. You can also configure additional S-VLAN nodes in other traffic-class groups.

In non-default traffic-class groups, you can configure a group node and VLAN queues. Although this apparently does not consume nodes, it does consume a hidden, phantom node for each queue, to satisfy the level requirement of the TFA ASIC.

Alternatively, use group nodes and shadow nodes.

We recommend that you configure an Ethernet shadow node in the group with the following QoS profile rule:

```
host1(config-qos-profile)#ethernet shadow-node group groupname
```

This rule stacks another node over the group node, so all VLAN queues are stacked above the single shadow node. No nodes are consumed in the traffic-class group.

**Related
Documentation**

- Managing System Resources for Shadow Nodes
- For QoS system maximums, see *JunosE Release Notes, Appendix A, System Maximums*
- [Monitoring the QoS Profiles Attached to an Interface on page 182](#)

CHAPTER 3

Classifying, Queuing, and Dropping Triple Play Traffic

- [Traffic Class and Traffic-Class Groups Overview on page 13](#)
- [Queuing and Buffer Management Overview on page 14](#)
- [Dropping Behavior Overview on page 16](#)

Traffic Class and Traffic-Class Groups Overview

A traffic class is a systemwide collection of buffers, queues, and bandwidth that you can allocate to provide a defined level of service to packets in the traffic class.

A traffic class corresponds to what the IETF DiffServ working group calls a traffic class in RFC 2597—Assured Forwarding PHB Group (June 1999).

Traffic classes are global to the router. Packets are:

- Classified into a traffic class on ingress or egress by input policies
- Queued on fabric queues that are specific to the traffic class
- Queued on the egress line module on queues that are specific to the traffic class
- Scheduled for transmission by the scheduler

Best-Effort Forwarding

The router has a default traffic class called best-effort. You cannot delete this class. You can add the best-effort class to a traffic-class group. The router assigns packets to the best-effort class in each of the following cases:

- You do not create any other traffic classes.
- Packets are not classified into a traffic class.
- Packets arrive at an egress line module that has no queues allocated for their traffic class.

Traffic-Class Groups Overview

You can put traffic classes into a group to create a hierarchy of scheduler nodes and queues. Organizing traffic into multiple traffic-class groups enables you to manage and shape traffic—by service class, for example—when the traffic classes are distributed across different VCs. A traffic-class group contains one or more traffic classes, but a particular traffic class can belong only to a single group—either the default group or one named group.

You can configure an auto-strict group and up to three extended traffic-class groups. You must put traffic classes that require strict-priority scheduling in the auto-strict group. You can optionally put traffic classes that need a separate round robin (for example, video) in an extended group.

A traffic class that is not contained in any named group is considered to belong to the default group. Traffic classes are placed in the default traffic-class group when the classes are configured—you can then move a class to another traffic-class group. When you delete a traffic-class from a named group, the class is automatically moved to the default traffic-class group. ATM VC nodes that are configured in the default group (which is the factory default configuration) receive backpressure from the segmentation and reassembly (SAR) feature in the default qos-mode-port node.

Traffic-class groups are global in scope by default. However, you might want to manage certain traffic classes through particular line modules. If you have already created a traffic-class group, you can subsequently specify a slot number to create a local instance of the group that is restricted to the module occupying that slot. Characteristics configured for the local group on the line module override those of the global group, for only that line module. Traffic classes in a globally scoped traffic-class group cannot belong to any other group. Traffic classes in a local traffic-class group cannot belong to any other group.

Related Documentation

- [Configuring Traffic Classes That Define Service Levels on page 51](#)
- [Configuring Traffic-Class Groups That Define Service Levels on page 52](#)

Queuing and Buffer Management Overview

A queue is a set of first-in, first-out (FIFO) buffers that buffer packets on the data path. QoS associates queues with a traffic class/interface pair. For example, if you create 4000 IP interfaces and configure each interface with four traffic classes, then 16,000 queues are created. For specific information about the maximum number of QoS queues supported, see *JunosE Release Notes, Appendix A, System Maximums*.

The E Series router dynamically manages the shared memory on egress line modules to provide a good balance between sharing the memory among queues and protecting an individual queue's claim on its fair share of the egress memory.

When egress packet memory is in high demand and aggregate utilization of the packet memory is high, queue lengths are set to lengths that strictly partition egress memory into per-queue memory sections. This conservative buffer-management strategy reserves a fair share of buffers for each queue, so that high bandwidth consumers cannot starve

out moderate traffic consumers by allocating all the shared memory resource for themselves.

When egress packet memory is in low demand, a more liberal buffer management strategy is used to provide active queues with more access to the shared memory resource.

The router dynamically varies queue lengths for all queues as the real-time demand on the egress packet memory changes. You can configure limits to prevent the router from setting queue lengths too low or too high.

Static Oversubscription

The router uses static oversubscription to vary queue thresholds based on the number of queues currently configured, which is relatively static. Static oversubscription is based on the assumption that, when a few queues are configured, many of the queues are likely to be active at the same time. When a large number of queues are configured, fewer queues are likely to be active at the same time.

When few queues are configured, buffer memory is strictly partitioned between queues to ensure that buffers are available for all queues. As the number of configured queues increases, buffer memory is increasingly oversubscribed to allow more buffer sharing. Reserving buffer space for all queues when many are expected to be idle is unnecessary and wasteful.

Dynamic Oversubscription

The router uses dynamic oversubscription to vary queue thresholds based on the amount of egress buffer memory in use. The router divides egress buffer memory into eight regions.

The size of the region depends on the ASIC type. For more information, see [“Memory Requirements for Queue and Buffers” on page 10](#).

When buffer memory is in low demand, queues are given large amounts of buffer memory. As the demand for buffer memory increases, queues are given progressively smaller amounts of buffer memory.

Color-Based Thresholding

Packets within the router are tagged with a drop precedence:

- Committed—Green
- Conformed—Yellow
- Exceeded—Red

When the queue fills above the exceeded threshold, the router drops red packets, but still queues yellow and green packets. When the queue fills above the conformed drop threshold, the router queues only green packets.



NOTE: All color-based thresholds vary in proportion to the dynamic queue length.

- Related Documentation**
- [Configuring Queue Profiles to Manage Buffers and Thresholds on page 53](#)
 - Guidelines for Managing Queue Thresholds
 - Guidelines for Managing Buffers
 - RED and WRED Overview

Dropping Behavior Overview

Drop profiles control the dropping behavior of a set of egress queues. They define the range within the queue where random early detection (RED) operates, the maximum percentage of packets to drop, and sensitivity to bursts of packets. Weighted random early detection (WRED) is an extension to RED that enables you to assign different RED drop profiles to each color of traffic.

The purpose of RED and WRED is to signal end-to-end protocols, such as TCP, that the router is becoming congested along a particular egress path. The intent is to trigger TCP congestion avoidance in a random set of TCP flows before congestion becomes severe and causes tail dropping on a large number of flows. Tail dropping can lead to TCP slow-starts, and tail dropping on a large number of flows results in global synchronization.

By default, tail dropping occurs when the length of a queue exceeds a threshold. Drop profiles allow you to employ active queue management by specifying RED and WRED parameters to be applied to an egress queue.

Congestion of an egress queue occurs when the rate of traffic destined for the queue exceeds the rate of traffic draining from the queue; the queue fills to its limit, and any further traffic destined to it must be discarded until there is room in the queue. RED and WRED monitor average queue length over time to detect incipient congestion.

You can combine drop profiles and queue profiles within a queue rule of a QoS profile to specify up to 256 unique queuing behaviors within the router. You can then associate these queuing behaviors in any combination with any of the egress queues.

- Related Documentation**
- [Queuing and Buffer Management Overview on page 14](#)

CHAPTER 4

QoS Statistics

- [QoS Statistics Overview on page 17](#)

QoS Statistics Overview

Statistics profiles enable you to gather statistics for the rate at which packets are forwarded out of a queue and for the rate at which committed, conformed, or exceeded packets are dropped. Statistics profiles also enable you to use events to monitor the rate statistics. You can then use **show** commands to view the results of the statistics gathering.

You can create up to 250 statistics profiles on the E Series Broadband Services Routers. The profiles are referenced by a queue rule within a QoS profile.

Statistics cannot be collected on failover queues.

When you create a statistics profile, you specify the time period over which statistics are gathered. To gather event statistics, you configure the thresholds for triggering rate-event reporting.

- Rate period—Time period, in seconds, over which statistics are gathered. For example, a 30-second rate period results in rate statistics being gathered over 30-second time segments.
- Forwarding rate threshold—Threshold for forwarding rate events. A forwarding-rate event is counted whenever the forwarding rate exceeds the specified threshold.
- Committed drop threshold—Threshold above which committed drop rate events are counted.
- Conformed drop threshold—Threshold above which conformed drop rate events are counted.
- Exceeded drop threshold—Threshold above which exceeded drop rate events are counted.

Rate Statistics

You can configure the E Series router to gather statistics for the rate at which queues forward and drop packets.

Queue rate statistics measure the forwarding and drop rates of each queue in bits per second. All bytes in the Layer 2 encapsulation are included in the rate calculation. For example, rates for a queue on Ethernet include the Ethernet and VLAN encapsulations.

For ATM modules, you can optionally configure queue statistics and queue rates to include the cell encapsulation and padding. Cell encapsulation and padding are referred to as the *cell tax*. The QoS shaping mode that you set on ATM line modules determines whether queue rate statistics include cell tax.

- If the interface is configured with frame-based QoS shaping mode, the egress queue measures frame rate statistics; an ATM cell tax is not included.
- If the interface is configured with cell-based QoS shaping mode, the egress queue measures cell rate statistics; cell rates include ATM Adaptation Layer 5 (AAL5) encapsulation and cell padding.
- If the interface is configured with byte adjustment, the egress queue measures rate statistics that are adjusted to the byte adjustment value.



NOTE: If you change the QoS shaping mode value in the middle of a rate period, the gathered rates are a mixture of cell- and frame-based rates for that one rate period. The next rate period uses a rate based on the new QoS shaping mode setting.

Event Statistics

You can configure the E Series router to count the number of times that forwarding or drop rates exceed a specific threshold. Events can be useful when you are monitoring service level agreements. For example, you might count the number of times that the drop rate of a queue is nonzero.

Bulk Statistics Support for QoS Statistics

You can obtain queue-level QoS statistics for each logical interface by querying the SNMP MIB. However, using SNMP to obtain queue-level statistics consumes significant network bandwidth because SNMP polls large volumes of data frequently. As an alternative to using the SNMP MIB, you can use the bulkstats statistics application.

The bulk statistics application provides components to configure and organize network accounting data in a flexible manner. The application reduces the consumption of network bandwidth by collecting queue-level statistics and periodically transferring the data to a remote server. You can configure the bulk statistics schemas to export network accounting data. In particular, the QoS schema supports the export of queue-level QoS statistics on egress queues for various interface types.

Configuring QoS schemas helps service providers monitor their network and report congestion and oversubscription by obtaining queue-level statistics and configuration information for each logical interface.

For information about schemas and configuring a bulk statistics schema to export queue-level QoS statistics for egress queues on the router, see *JunosE System Basics Configuration Guide, Chapter 4, Configuring SNMP*.

**Related
Documentation**

- [Configuring Statistic Profiles for QoS on page 85](#)
- [Monitoring the Configuration of Statistics Profiles](#)
- [Troubleshooting Memory and Processor Use for Egress Queue Rate Statistics and Events on page 219](#)

CHAPTER 5

Scheduling and Shaping Triple Play Traffic

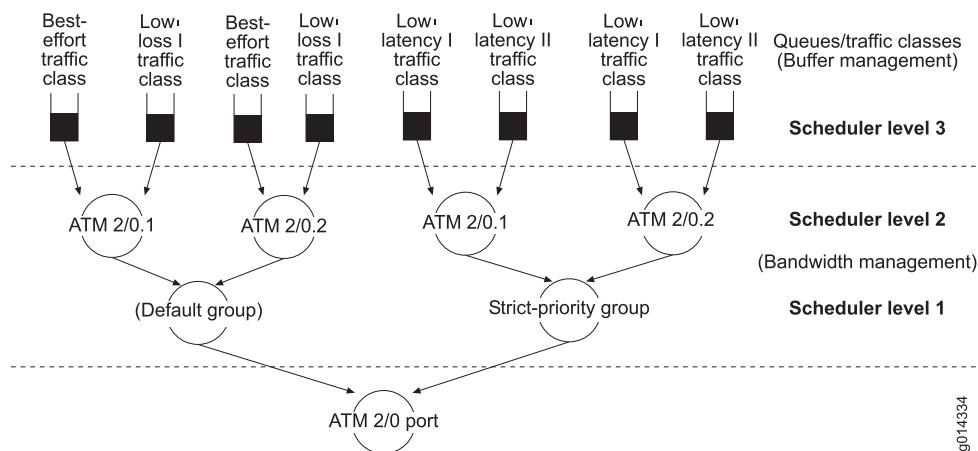
- [Scheduler Hierarchy Overview on page 21](#)
- [Shared Shaping Overview on page 23](#)
- [How Shared Shaping Works on page 24](#)
- [Simple Shared Shaping Overview on page 25](#)
- [Simple Shared Shaping Algorithm Overview on page 27](#)
- [QoS Parameter Overview on page 29](#)

Scheduler Hierarchy Overview

The egress line module scheduler is an HRR scheduler. [Figure 2 on page 22](#) is an example of a QoS scheduler's hierarchy.

As shown in [Figure 2 on page 22](#), the queues feeding a physical port are organized in a hierarchy. At each level in the hierarchy, the scheduler uses shaping rates, hierarchical or assured rates, and relative weights to determine the allocated bandwidth:

- The scheduler selects a first-level node based on the allocated bandwidth.
- The scheduler then selects a second-level node from the group of nodes that are stacked above the selected first-level node. This selection is also based on the allocated bandwidth.
- Finally, the scheduler selects a queue from the group of queues stacked above the second-level node.

Figure 2: QoS Scheduler Hierarchy


Shaping Rates, Assured Rates, and Relative Weights in a Scheduler Hierarchy

The scheduler supports hierarchical and static assured rates, relative weights, and shaping rates on all three levels of the hierarchy: first-level node, second-level node, and queue. The bandwidth delivered from a given node or queue is a function of the shaping rate and either the assured rate or relative weight:

- When the scheduler is not congested, the shaping rates determine which node or queue can claim the bandwidth. The shaping rate specifies the maximum bandwidth to the node or queue.
- When the scheduler is congested, either the hierarchical or static assured rate or the weight specifies the minimum bandwidth.
 - If the scheduler is configured to use a static assured rate and the assured rate is other than none (the default), it is used to determine the allocated bandwidth, and the weight setting is ignored. If the assured rate is zero, the weight setting is used to determine the bandwidth.

The static assured rate specifies the desired bandwidth. This rate is guaranteed until the bandwidth becomes oversubscribed.

- If the scheduler is configured to use hierarchical assured rate, the scheduler dynamically adjusts the amount of allocated bandwidth for service delivery based on the sum of the assured rates of all child nodes and queues.
- The assured rate also specifies that if bandwidth is over- or undersubscribed, all adjustments are made in proportion to the original assured-rate specification.

For example, if Node A is configured to receive 40 Mbps and Node B receives 20 Mbps, any available bandwidth above the subscribed total of 60 Mbps would be allocated to the two nodes at the same 2-to-1 ratio. Similarly, if the bandwidth were oversubscribed and only 30 Mbps were available, this amount would also be allocated to the two nodes at the 2-to-1 ratio, with Node A getting 20 Mbps and Node B getting 10 Mbps.



NOTE: For E Series ASIC modules, strict priority is supported only for a single first-level scheduler node.

When determining the shaping rate, the system includes all bytes in Layer 2 encapsulations. The packets that are included in the rate depend on the Layer 2 node that is specified in the QoS profile. For example, the shaping rate for an Ethernet node includes bytes from the Ethernet and VLAN encapsulations.

Related Documentation

- Static and Hierarchical Assured Rate Overview
- Rate Shaping and Port Shaping Overview
- [Shared Shaping Overview on page 23](#)
- [Configuring a Scheduler Hierarchy on page 57](#)

Shared Shaping Overview

In the JunosE Software QoS implementation, you configure a traffic-class group to create a separate scheduler hierarchy. Traffic classes in a traffic-class group are queued through a scheduler hierarchy dedicated to that group. QoS supports up to five user-configurable, named traffic-class groups. Traffic classes that do not belong to any named group belong to the default traffic-class group. With the factory default configuration, the best-effort traffic class is in the default traffic-class group.

Shared shaping is a mechanism for shaping a logical interface's aggregate traffic to a rate when the traffic for that logical interface is queued through more than one scheduler hierarchy. For example, a service provider can configure QoS for voice, video, and data traffic on a single ATM VC. The video traffic and the voice traffic are placed in separate scheduler hierarchies from the data traffic to provision the low latency that is required for voice traffic and the higher bandwidth that is required for video traffic.

In this scenario, the data traffic needs to be dynamically shaped so that its rate matches the bandwidth available after the voice and video bandwidth requirements are met. When less voice and video traffic is being forwarded, then the data traffic can expand to fill the line rate.

When determining a shared shaping rate, the system includes all bytes in Layer 2 encapsulations. The packets that are included in the rate depend on the node specified. For example, rates for an Ethernet node include the Ethernet and VLAN encapsulations.

Shared shaping is typically enabled on the access-facing line module, but you can enable the feature for any interface type recognized by QoS, on any line module and any E Series Broadband Services Routers.

Related Documentation

- [Simple Shared Shaping Overview on page 25](#)
- Compound Shared Shaping Overview

How Shared Shaping Works

You can configure the shared-shaping rate on either the best-effort scheduler node or the best-effort queue for the logical interface. The router also locates the queues in named traffic-class groups that are associated with the logical interface and shapes that set of queues to the shared rate. The shared-shaping rate is the total bandwidth for the logical interface.

A typical configuration places the low-latency voice traffic in the auto-strict-priority traffic-class group and video traffic in a separate extended traffic-class group. The data traffic is usually queued in the best-effort traffic class in the default traffic-class group.

The constraints of both the legacy hierarchical scheduler and the shared shaper affect the bandwidth of scheduler objects. The shared shaper limits the bandwidth even when the port or VP is not congested. When the port or VP is congested, the legacy scheduler is dominant. For example, when a heavily oversubscribed VP becomes congested, the legacy hierarchical scheduler may limit the VP bandwidth to a lower rate, so that shared shaping of excess bandwidth does not apply.

When determining the shared-shaping rate, the system includes all bytes in Layer 2 encapsulations. The packets that are included in the rate depend on the Layer 2 node that is specified in the QoS profile. For example, the shaping rate for an Ethernet node includes bytes from the Ethernet and VLAN encapsulations.

Two types of shared shaping are available, depending on your hardware. *Simple* shared shaping can shape the best-effort node or queue associated with a logical interface to a shared rate. *Compound* shared shaping is a hardware-assisted mode that controls bandwidth for all scheduler objects associated with the subscriber logical interface.

[Table 6 on page 24](#) compares the two types of shared shaping that are available.

Table 6: Comparison of Simple and Compound Shared Shaping

Shared Shaper	Advantages
Simple	<ul style="list-style-type: none"> Simple shared shaping is useful for triple-play configurations, because it manages voice and video queues in addition to data queues so that the shared rate cannot be exceeded. You can use line modules that have any ASIC hardware.
Compound	<ul style="list-style-type: none"> Compound shared shaping is useful for triple-play configurations, because it manages voice and video queues in addition to data queues so that the shared rate cannot be exceeded. Compound shared shaping responds to changes in traffic rates more rapidly than simple shared shaping, in the order of milliseconds. You can use line modules with the EFA2 ASIC or the TFA ASIC.

Active Constituents for Shared Shaping

When you specify a shared-shaping rate on a best-effort node or queue, QoS shapes the aggregate of traffic for the logical interface that owns the best-effort queue or node. QoS locates the queues and nodes owned by that logical interface and applies the shared shaper to them. The nodes and queues owned by the interface are called the *constituents* of the shared-shaper instance. For example, if the logical interface type is VC, the constituents are all VC objects: VC nodes and VC queues. A shared-shaping rule in a profile can apply to up to eight constituents.

Active constituents are actively controlled by the shared-shaper mechanism. *Inactive* constituents are indirectly controlled. For example, when ATM VC queues are stacked above an ATM VC node, the ATM VC node might be an active constituent. In this case, the queues stacked above the node are shaped to the shared rate indirectly by the hierarchical scheduler. If the ATM VC queues are the active constituents, then the ATM VC node is inactive.

- Related Documentation**
- [Simple Shared Shaping Overview on page 25](#)
 - Compound Shared Shaping Overview
 - Constituent Selection for Shared Shaping Overview

Simple Shared Shaping Overview

Simple shared shaping shapes the best-effort node or queue associated with a logical interface to a shared rate.

Bandwidth Allocation for Simple Shared Shaping

Once per second, the simple shared shaper calculates the combined rate of the voice and video queues for the logical interface, and shapes the best-effort queue for the data traffic to the shared rate minus the video and voice queue rates. The bandwidth for the voice and video queues is determined by the configuration of the hierarchical scheduler. The shared shaper does not actively manage the video and voice queues.

Simple Shared Shaping on the Best-Effort Scheduler Node

If you have a second traffic class for data in addition to the best-effort data traffic class, configure shared shaping on the best-effort scheduler node. In this scenario, two weighted queues are stacked above the best-effort scheduler node, one for the best-effort traffic class and the other for the second data traffic class. If you configure the shared-shaping rate on the best-effort queue, then the shared shaper can have a tendency to starve the best-effort queue in favor of the second data queue. If you instead configure the shared-shaping rate on the best-effort node, the hierarchical scheduler allocates bandwidth between multiple data queues based on their relative weight and assured rate.

If you are configuring VP shared shaping, configure shared shaping on the best-effort scheduler node for the VP. Shaping the best-effort scheduler node for the VP has the

effect of shaping all the VC best-effort queues for that VP. This enables you to retain the advantages of per-VC queuing in the hierarchical scheduler.

If you are configuring VC shared shaping and the SAR is operating in low-CDV mode, we recommend you configure the shared-shaping rate on the best-effort scheduler node for the VP or VC. The router sets the SAR shaper for the VC or VP to match the shared-shaping rate on VC and VP nodes in the hierarchical scheduler; this is usually the desired behavior. A shared shaper configured on the best-effort queue does not trigger the matching shaper in the SAR.

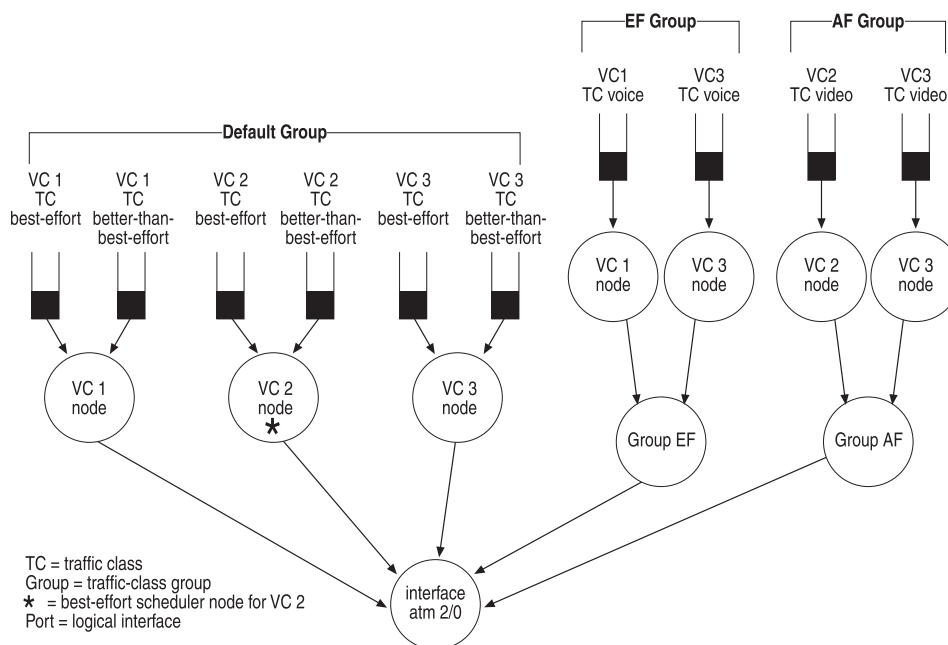
Simple Shared Shaping for Triple-Play Networks

Simple shared shaping enables you to shape the logical interface to a single rate for triple-play networks.

In [Figure 3 on page 26](#), the AF traffic-class group contains the video traffic class. The EF traffic-class group contains the voice traffic class. The best-effort and better-than-best-effort traffic classes remain outside any traffic-class group. Because the voice, video, and data queues are stacked in separate scheduler hierarchies, you must use the shared shaper to shape the logical interface aggregate to a single rate.

In this example, VC 1 is configured for voice and data. VC 2 is configured for data and video. VC 3 is configured for data, voice, and video. The shared shaper is configured on the best-effort node or queue for VC 1; the corresponding voice queue for VC 1 shares the configured rate.

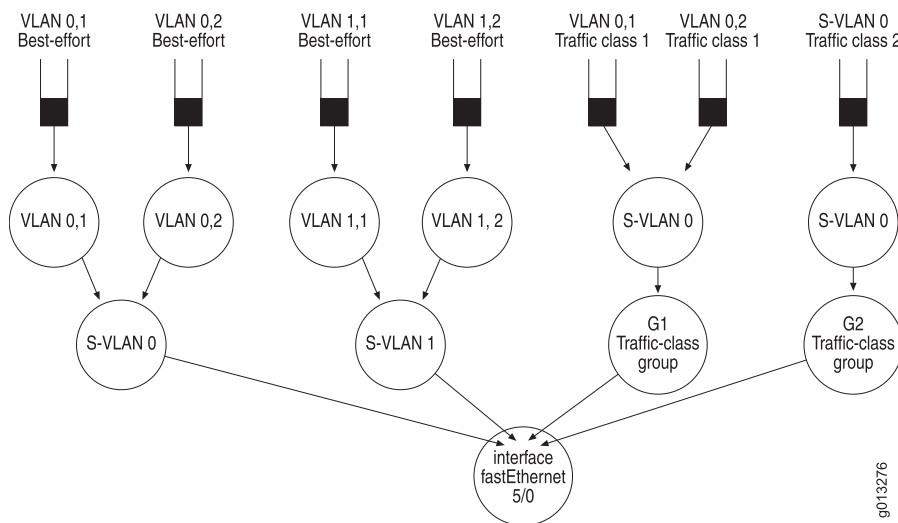
Figure 3: Simple Shared Shaping over ATM



In a typical triple-play network configuration over Ethernet, individual subscribers are represented on the B-RAS by VLANs and DSLAMs by SVLANs. [Figure 4 on page 27](#)

illustrates how to shape the subscriber aggregate of voice, video, and data to a single rate in Ethernet.

Figure 4: Simple Shared Shaping over Ethernet



- Related Documentation**
- [Shared Shaping Overview on page 23](#)
 - [Configuring Simple Shared Shaping](#)
 - [Constituent Selection for Shared Shaping Overview](#)


Simple Shared Shaping Algorithm Overview

You can configure variables within the simple shared shaper algorithm to control the minimum dynamic rate for all simple shared shapers on the router.

Configuring variables in the simple shared shaper algorithm is useful for IPTV configurations. Without limiting the dynamic rate, best-effort data traffic can be starved for a few seconds when a video stream starts. The minimum dynamic rate defined by shared shaper algorithm variables applies to best-effort traffic only.

[Figure 5 on page 27](#) shows a two-constituent simple shared shaper consisting of best-effort and video traffic. The sum of the best-effort and video traffic is shaped to the configured shared-shaping rate.


Figure 5: Simple Shared Shaper Behavior Without Algorithm Controls



When the video stream starts in the example displayed by [Figure 5 on page 27](#), the shared shaper reacts by drastically reducing best-effort traffic because it must avoid saturating downstream queues. In some cases, best-effort traffic is throttled for a few seconds. When the video stream stops, best-effort traffic can continually consume more bandwidth, up to the shared-shaping rate.


By controlling the minimum dynamic rate in the simple shared shaper algorithm, you can configure the less conservative simple shared shaping behavior displayed in [Figure 6 on page 28](#). In this example, as the video traffic starts, the best-effort rate is reduced less drastically, and best-effort traffic is not starved.

Figure 6: Less Conservative Simple Shared Shaper Behavior



You can also configure the more liberal simple shared shaper behavior that is displayed in [Figure 7 on page 28](#). In this example, the initial over-limit video traffic is ignored. When the video traffic stops, the system immediately allows best-effort traffic to consume the available bandwidth.

Figure 7: More Liberal Simple Shared Shaper Behavior



Simple Shared Shaper Algorithm Calculations

The simple shared shaper algorithm performs the following tasks to calculate the dynamic rate:

1. Calculates the new measured rate.
2. Calculates the virtual output queue length (VOQL).
3. Calculates the new dynamic rate.
4. Uses the larger value of the new dynamic rate (from Step 3) and a minimum dynamic rate.

- Related Documentation**
- Variables of the Simple Shared Shaper Algorithm
 - Configuring Simple Shared Shaper Algorithm Variables

QoS Parameter Overview

Using QoS parameters, you can configure a queuing architecture without specifying the numeric subscriber rates and weights in scheduler profiles. You then use the same QoS and scheduler profiles across all subscribers who use the same services but at different bandwidths, reducing the total number of QoS profiles and scheduler profiles required.

Using QoS parameters, you can specify the following attributes of a scheduler node or queue without specifying the numeric value explicitly in the scheduler profile:

- Shaping rate
- Shared-shaping rate
- Assured rate
- Scheduler weight

- Related Documentation**
- QoS Parameter Audience
 - QoS Parameter Terms

CHAPTER 6

Multicast Bandwidth and Byte Adjustment of Triple Play Traffic

- [Hierarchical QoS Parameters Overview on page 31](#)
- [IP Multicast Bandwidth Adjustment for QoS Overview on page 31](#)
- [Byte Adjustment for ADSL and VDSL Traffic Overview on page 33](#)

Hierarchical QoS Parameters Overview

You use hierarchical parameters in applications where you want the system to add instances associated with child interfaces and associate the sum with a parent interface. For example, to shape an S-VLAN to 50 percent of the sum of the shaping rates of the VLANs stacked above the S-VLAN, you specify *explicit* instances of the parameter associated with the VLANs, and the system creates an *implicit* instance of the parameter associated with the S-VLAN. The parameter maintains the value of the sum of the explicit instances.

The most common use of hierarchical parameters is in combination with the IP multicast bandwidth adjustment application.

For example, you create a hierarchical parameter that controls a VLAN. The hierarchical parameter has two explicit parameter instances on two IP interfaces, with values of 1 Mbps and 3 Mbps. Therefore, an implicit parameter instance is created at the VLAN interface with a value of 4 Mbps.

Related Documentation

- [Configuring a Parameter Definition to Calculate Hierarchical Instances on page 77](#)
- For information about the IP multicast bandwidth adjustment application, see [IP Multicast Bandwidth Adjustment for QoS Overview on page 31](#)

IP Multicast Bandwidth Adjustment for QoS Overview

You can associate the IP multicast bandwidth adjustment application (**ip-multicast**) with a parameter definition. Before you begin, you must define a multicast bandwidth map and the QoS adjustment for a virtual router.

You use the IP multicast bandwidth adjustment application to set the shared-shaping rate for a subscriber when a downstream DSLAM is replicating a multicast frame for

multiple downstream transmissions on a subscriber circuit. In this case, the router does not schedule the multicast traffic on a subscriber VLAN, but limits the scheduled non-multicast traffic on the subscriber VLAN so that the total of non-multicast and multicast traffic at the DSLAM is less than the subscriber shared-shaping rate.

To implement this, the IP multicast bandwidth adjustment application tracks the bandwidth of multicast flows based on IGMP joins and leaves. When the QoS administrator configures a QoS parameter with the IP multicast bandwidth adjustment application, the application automatically configures an instance of that parameter for each subscriber that is receiving multicast traffic. The value of the parameter instance is equal to the multicast bandwidth for a subscriber at a specific time. The shared-shaping rate of the VLAN node can be configured using a parameter expression such as `max-subscriber-bandwidth - ip-multicast-bandwidth`.

In a typical IP multicast bandwidth adjustment configuration, the shaping rate or shared-shaping rate is determined by calculating the total subscriber bandwidth of the logical interface minus the ip-multicast bandwidth. To enable the IP multicast QoS adjustment, you must:

- Define a qos-parameter using the **qos-parameter-define** command with the application **ip-multicast** and the **hierarchical** keyword.

```
host1(config)# qos-parameter-define ipm application ip-multicast hierarchical
host1(config-qos-parameter-define)#
```

- Reference the ipm parameter within a scheduler profile. For example:

```
host1(config)#scheduler-profile totalSubscriberBw
host1(config-scheduler-profile)#shared-shaping-rate 10000000 - ipm auto
```

This scheduler profile contains an expression for the shared-shaping rate that limits the shared-shaping rate to 10 Mbps less the rate of any IP multicast traffic.

- Reference the scheduler profile within a QoS profile rule. For example:

```
host1(config)#qos-profile subscriber
host1(config-qos-profile)#vlan node scheduler-profile totalSubscriberBw
```

This QoS profile rule limits a subscriber with vlan to the rate specified in the totalSubscriberBw scheduler profile.

QoS clients do not need to create a parameter instance to activate the IP multicast bandwidth adjustment application. The system automatically creates explicit instances based on IGMP joins and leaves.

When a subscriber logs in, the QoS scheduler hierarchy is created with the vlan configured for shared shaping, based on the expression `10000000 - ipm`. If no multicast traffic is being transmitted, there is no ipm parameter instance with the vlan.

To calculate the subscriber bandwidth from the total subscriber bandwidth, you must create a global parameter instance using the **ip-multicast** keyword and set the value to 0.

To ensure the system can locate an instance of the ipm parameter for subscribers that are not receiving multicast traffic, you must create a global parameter with a value of 0:


```
host1(config)# qos-parameter ipm 0
```

If you do not create the global parameter instance, the expression result is undefined for these subscribers and the shared shaping rate is not set.

By configuring a global parameter instance of 0, the value is applied to all the interfaces that reference the parameter. QoS overrides the global ipm parameter instance with the value specified in the bandwidth map for a specific IP interface on which IGMP joins.

Related Documentation

- Guidelines for Configuring IP Multicast Adjustment for QoS
- *JunosE Multicast Routing Configuration Guide*
- Scheduler Profiles and Parameter Expressions for QoS Administrators

Byte Adjustment for ADSL and VDSL Traffic Overview

You can associate a parameter definition with a byte adjustment application to adjust the shaping rates for ADSL and VDSL traffic on E Series Broadband Services Routers.

The byte adjustment differs for interfaces with cell shaping mode and frame shaping mode. For ADSL traffic, JunosE Software supports a byte adjustment application (**qos-byte-adjustment**) to adjust rates for cell shaping mode. For VDSL traffic, JunosE Software supports a frame byte-adjustment application (**qos-frame-byte-adjustment**) to adjust rates for frame shaping mode.

Frame is the default shaping mode for Ethernet interfaces on E Series routers. To configure the cell shaping mode, issue the **qos-shaping-mode** command or by specifying the **qos-cell-mode** application with a parameter definition.

Byte Adjustment for Cell Shaping of ADSL Traffic Overview

Managing the bandwidth of downstream ATM traffic to Ethernet interfaces is difficult because of the different layer 2 encapsulations. To reduce the number of packet drops in the Ethernet network, you can use the byte adjustment applications to account for the different encapsulations.

To adjust the shaping rates to account for different layer 2 encapsulations as well as the ATM cell pad, header, and trailer on interfaces, apply a parameter with the cell byte-adjustment application (**qos-byte-adjustment**).

When you apply a parameter with the **qos-byte-adjustment** application to an interface with frame shaping mode, you adjust shaping rates to account for different layer 2 encapsulations only.

Calculation and Example of Byte Adjustment for Cell Shaping

The system counts the bytes transmitted to track the shaping rate. Instead of counting the actual packet size, the system uses the CPE packet size. You can configure the byte adjustment so that the shaping rate matches the CPE bandwidth.

By default, the byte adjustment is set to 0. If the overhead between the access node and CPE is 0, you do not need to configure the byte adjustment value.

Figure 8 on page 34 displays an example of an Ethernet encapsulation and an ATM encapsulation.

Figure 8: Byte Adjustment Calculation for Ethernet and ATM Encapsulations

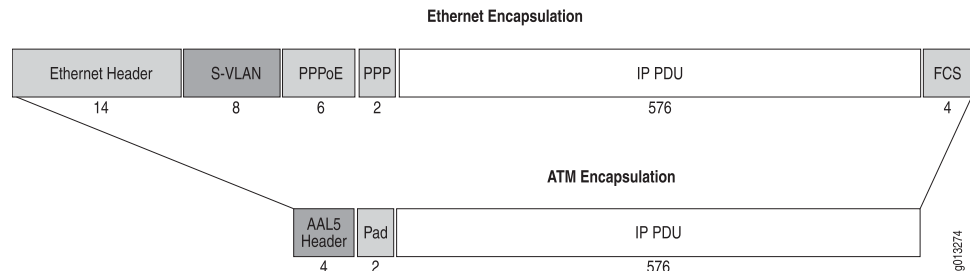


Table 7 on page 34 lists the header lengths for the Ethernet encapsulation, which represents the CPE protocol overhead. The hierarchy is PPPoE over S-VLAN over Ethernet.

Table 7: Header Lengths for Ethernet Encapsulation

Header	Number of Bytes
EnetHeader	14 bytes (6-SA, 6-DA, 2-ethertype)
Vstack	8 bytes (2-vmanTci, 2-ethertype, 2-vlanTci, 2-ethertype)
PppoeHeader	6 bytes (1-version/type, 1-code, 2-session id, 2-length)
Ppp	2 bytes (2-protocol id)
FCS	4 bytes
Total	34 bytes

Table 8 on page 34 lists the header lengths for the ATM encapsulation, which represents the B-RAS protocol overhead. The interface stack is PPPoA over ATM 1483 with LLC Mux. The ATM AAL5 trailer is considered cell tax and is not part of the byte adjustment calculation.

Table 8: Header Lengths for ATM Encapsulation

Header	Number of Bytes
ATM AAL5 LLC	4 bytes
PPP	2 bytes (2-protocol id)
Total	6 bytes

The byte adjustment calculation for these encapsulations is:

Byte Adjustment for Frame Shaping of VDSL Traffic Overview

Packet fragmentation can occur at a DSLAM because of the associated segment header that is added for VDSL2 in frame mode. Because the segment header is not included in the ANCP rate report, the forwarding rate on an E Series router can be higher than the DSLAM rate, which can result in packet loss.

You can use a QoS parameter expression with the frame byte-adjustment application to reduce the forwarding rate so that it matches the rate at the DSLAM. To adjust rates for interfaces with frame shaping mode, apply the frame byte-adjustment application (**qos-frame-byte-adjustment**).

When you apply a parameter with the **qos-byte-adjustment** application to an interface with frame shaping mode, you adjust shaping rates to account for different layer 2 encapsulations only.

System Calculation for Byte Adjustment of ADSL and VDSL Traffic

You can create parameter instances for the cell byte-adjustment application and the frame byte-adjustment application on the same system. The system performs the byte adjustment calculation based on the shaping mode specified. The byte adjustment can have both a positive and negative value.

[Table 9 on page 35](#) lists the final byte adjustment value that the system uses depending on the configured shaping mode and the value that you configured for the byte adjustment applications.

Table 9: Byte Adjustment Values for Frame and Cell Shaping Modes

Shaping Mode on Port 0	Configured qos-frame-byte-adjustment Value	Configured qos-byte-adjustment Value	Final Byte Adjustment Value
Cell	Any value	-4	-4
Cell	Any value	Undefined	0
Frame	Undefined	Undefined	0
Frame	8	-4	8
Frame	Undefined	8	8

Related Documentation

- [Configuring a Parameter Definition to Adjust Cell Shaping Rates for ADSL Traffic](#)
- [Configuring a Parameter Definition to Adjust Frame Shaping Rates for VDSL Traffic](#)
- [Example: QoS Parameter Configuration for QoS Cell Mode and Byte Adjustment for Cell Shaping](#)
- [QoS Shaping Mode for Ethernet Interfaces Overview on page 37](#)

- Cell Shaping Mode Using QoS Parameters Overview
- QoS Downstream Rate Application Overview

CHAPTER 7

Ethernet Interface Solutions for Triple Play Traffic

- [Providing QoS for Ethernet Overview on page 37](#)
- [QoS Shaping Mode for Ethernet Interfaces Overview on page 37](#)
- [QoS for 802.3ad Link Aggregation Interfaces Overview on page 39](#)
- [Hashed Load Balancing for 802.3ad Link Aggregation Groups Overview on page 41](#)
- [Subscriber Load Balancing for 802.3ad Link Aggregation Groups Overview on page 42](#)

Providing QoS for Ethernet Overview

Managing the bandwidth of downstream ATM traffic to Ethernet interfaces is difficult because of different layer 2 encapsulations and the ATM cell pad, trailer, and header.

The SAR scheduler is not available for Ethernet interfaces. However, you can still configure the operational shaping mode to shape downstream ATM traffic based on either frames or cells. Configuring cell-based shaping enables you to reduce packet drops in the Ethernet network by adjusting shaping for the ATM cell pad, trailer, and header.

You can also use RADIUS to provide QoS on bulk-configured VLAN subinterfaces.

Related Documentation

- [QoS Shaping Mode for Ethernet Interfaces Overview on page 37](#)
- [Creating a QoS Interface Hierarchy for Bulk-Configured VLAN Subinterfaces with RADIUS on page 67](#)
- [QoS for 802.3ad Link Aggregation Interfaces Overview on page 39](#)

QoS Shaping Mode for Ethernet Interfaces Overview

The SAR scheduler is not available for Ethernet interfaces. However, you can still configure the operational shaping mode to shape ATM traffic based on either frames or cells by issuing the **qos-shaping-mode** command.

Frame shaping mode is the default for Ethernet interfaces on all E Series Broadband Services Routers. You can configure cell shaping mode for the following interfaces:

- Gigabit Ethernet interfaces on the GE-2 line module and the GE-HDE line module on ERX routers
- Gigabit Ethernet and 10-Gigabit Ethernet interfaces on the ES2 4G LM on E120 and E320 Broadband Services routers
- 10-Gigabit Ethernet interfaces on the ES2 10G LM on E120 and E320 routers

When you use cell shaping mode to configure the shaping or policing rate, the resulting traffic stream conforms exactly to the policing rates configured in downstream ATM switches. Using cell shaping also reduces the number of packet drops in the Ethernet network.

The setting for port 0 provides the default shaping mode for all ports on the same I/O module or IOA. Individual ports can have a specific shaping mode setting that overrides the setting for port 0.

If you do not configure the QoS shaping mode for a port, the shaping mode is calculated using the value for port 0 on the same I/O module or IOA. If the port's shaping mode is configured, the system uses the port's shaping mode.

Table 10 on page 38 lists the possible combinations of the **qos-shaping-mode** command and the resultant operational shaping mode.

Table 10: Operational Shaping Modes

qos-shaping-mode for Port 0	qos-shaping-mode for Other Ports	Operational Shaping Mode
Cell	Cell	Cell
Frame	Frame	Frame
Cell	Frame	Frame
Frame	Cell	Cell
Frame	No shaping mode	Frame
Cell	No shaping mode	Cell
No shaping mode	No shaping mode	Frame

To account for different layer 2 encapsulations, you can configure the byte adjustment application using QoS parameters. The byte adjustment is calculated differently for frame shaping mode than cell shaping mode.



NOTE: You can also use the QoS cell mode application with QoS parameters to configure the shaping mode for a port.

- Related Documentation**
- [Configuring the QoS Shaping Mode for Ethernet Interfaces](#)
 - [Byte Adjustment for ADSL and VDSL Traffic Overview on page 33](#)
 - [Cell Shaping Mode Using QoS Parameters Overview](#)

QoS for 802.3ad Link Aggregation Interfaces Overview

You can configure QoS for 802.3ad link aggregation interfaces. To ensure that QoS is applied properly to the interface column, you configure the QoS profile using either a hashed loadbalancing scheme or a subscriber loadbalancing scheme.

Types of Load Balancing

For hashed load balancing, you configure the scheduler hierarchy with Ethernet queues, and the system replicates them on each link within the link aggregation group (LAG). The system demultiplexes each packet to one of the active links in the LAG using a random hash generated by fields in the packet header. For example, when an IP packet is routed to a LAG, the hash algorithm is based on the IP Source Address and Destination Address in the IP header.

For subscriber load balancing, you configure the scheduler hierarchy with IP, VLAN, and S-VLAN queues and the system allocates them to individual ports in the LAG. The system demultiplexes each packet to an active link based on the subinterface underlying the egress interface. For example, when an IP packet is routed to an IP interface over a LAG, the system binds the underlying VLAN, PPPoE, or MPLS subinterface to one of the active links in the LAG. The packet is transmitted over the interface.

Most network operators configure QoS over 802.3ad LAGs using subscriber load balancing to take advantage of subscriber class-based queueing (SCBQ) features. However, configuring hashed load balancing is useful for achieving fine-grained distribution of multicast VLAN traffic or for any high bandwidth VLAN that does not require shared shaping.

To ensure that QoS is symmetrically applied to all the links, the router periodically rebalances the traffic on the LAG. You can control the loadbalancing parameters.

Munged QoS Profiles and Load Balancing

To determine whether to use hashed load balancing or subscriber load balancing, the system munges a QoS profile for a subscriber.

In typical Ethernet configurations, the munged QoS profile for a given subscriber interface comprises the accumulated rules of the QoS profiles attached below the subscriber interface in the interface column. Rules in higher-attached QoS profiles override or eclipse rules in lower-attached QoS profiles. For example, rules from specific interface attachments such as a VLAN override those from attachments at S-VLANs or ports.

When applying QoS to LAGs, the system uses a modified algorithm to munge QoS profile attachments. The system automatically builds the munged QoS profile using the rules in the QoS profile attached at the LAG interface.

For example, the munged QoS profile for VLAN 0,0 consists of the munge of:

- Attachment 1—QoS profile attached to the VLAN
- Attachment 2—QoS profile attached to the S-VLAN
- Attachment 3—QoS profile attached to the LAG

If there is no QoS profile attached to the LAG, the system locates the lag-default QoS profile indicated in the **qos-port-type-profile** command.

If the resulting QoS profile specifies only Ethernet queues, the system uses the hash algorithm to balance the links. If the resulting QoS profile specifies any VLAN, IP, or L2TP-Session queues, then the system uses subscriber load balancing.

802.3ad Link Aggregation and QoS Parameters

You can create parameter instances for IEEE 803.ad LAG interfaces. A parameter instance for LAG can control an Ethernet port or a node, but you cannot create parameter instances for the Ethernet interfaces within the LAG.

For example, a LAG instance can specify a shaping rate of 100 Mbps on an Ethernet port or a group node. The system shapes all Ethernet ports or group nodes to the same rate within the LAG. Using load balancing, the system strives to balance the traffic each link equally.

QoS and Ethernet Link Redundancy

You can configure Ethernet link redundancy for LAG interfaces. When you configure QoS for those links, be sure to consider the following behaviors.

Active Link Failure and QoS

When an active link fails, traffic that is hashed-load balanced is redirected onto the remaining active links in the LAG. Traffic that is hashed-load balanced might be lost on the disabled link, but from the moment of switchover, traffic arriving from the fabric on the egress line module is directed towards one of the remaining hashed load-balanced queues.

Subscriber loadbalanced traffic takes more time to reestablish on active links because of the amount of computation (approximately 3 ms per subscriber). During this time period, traffic directed to the disabled link might be lost.

Administratively Disabling a Link and QoS

When a link is administratively disabled, the system immediately redirects traffic from the link to other links in the LAG.

Adding a New Link to the LAG and QoS

When you add a new link to the LAG, the system immediately sends traffic that is hashed-load balanced to the link. Traffic that is subscriber-load balanced moves to the new link as new subscribers log in. The system automatically rebalances traffic to the new link based on the load rebalance configuration for the LAG.

Related Documentation

- [Hashed Load Balancing for 802.3ad Link Aggregation Groups Overview on page 41](#)
- [Subscriber Load Balancing for 802.3ad Link Aggregation Groups Overview on page 42](#)
- [Enabling Default Subscriber Load Balancing for 802.3ad Link Aggregation Groups on page 70](#)
- Ethernet Link Redundancy Overview
- Parameter Definition Attributes for QoS Administrators Overview
- Munged QoS Profile Overview
- *ERX Module Guide* and the *E120 and E320 Module Guide*

Hashed Load Balancing for 802.3ad Link Aggregation Groups Overview

To configure hashed load balancing, you configure a scheduler hierarchy with Ethernet queues and the system replicates the queues for each link within the LAG. The system shares the traffic equally across the links based on the distribution characteristics defined in the hash algorithm.

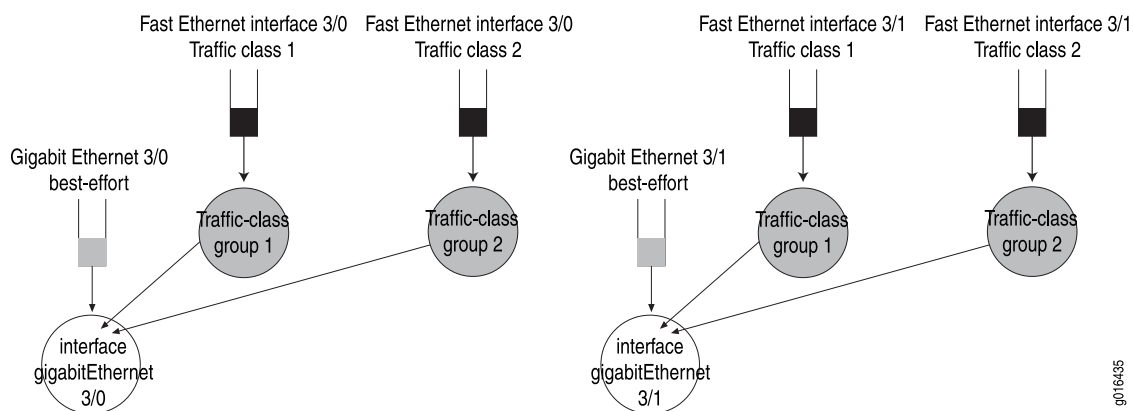
Because all traffic is carried in Ethernet queues, per-subscriber QoS features such as shared shaping for VLANs are not available.

Sample Scheduler Hierarchy for Hashed Load Balancing

Figure 9 on page 41 displays a sample 802.3ad link aggregation scheduler hierarchy that uses hashed load balancing.

The Gigabit Ethernet interfaces are on the same line module and are members of a LAG. The system dynamically balances traffic between the Ethernet queues on the two ports.

Figure 9: 802.3ad Link Aggregation Scheduler Hierarchy

**Related Documentation**

- [Configuring Load Rebalancing for 802.3ad Link Aggregation Groups on page 71](#)

Subscriber Load Balancing for 802.3ad Link Aggregation Groups Overview

To configure subscriber load balancing, you configure a scheduler hierarchy with nodes and queues for IP, VLANs, and S-VLANs. The system distributes those nodes and queues in the scheduler hierarchy over the ports within the LAG using a technique called *partitioning*.

Ethernet queues used for hashed load balancing are always present in the scheduler hierarchy.

To ensure that QoS is symmetrically applied to all the links, the router periodically rebalances the load within the LAG using a hash algorithm. You can control the loadbalancing parameters and configure the system to dynamically rebalance. Partitioning the Scheduler Hierarchy

The system then partitions the scheduler hierarchy by binding the IP, VLAN, L2TP session, and MPLS resources for each subscriber to a selected link within the LAG at the time the subscriber interface is configured.

S-VLANs and Subscriber Load Balancing

The system *clones* S-VLAN nodes and queues on each link in the LAG. The system clone S-VLANs so it can allocate subscribers that share a common S-VLAN ID to different links within the LAG. S-VLAN nodes and queues are the only resources that are cloned; the system always allocates nodes and queues for other interface types to a single selected link.

Cloning S-VLAN nodes enables fine-grained load balancing within the LAG because VLANs within the S-VLAN can be allocated to the link with the least traffic. However, cloned S-VLANs can introduce anomalous scheduling behavior. A shaped S-VLAN node within the LAG shapes traffic on a per-link basis. Shaping a LAG S-VLAN node to 2 Mbps on a LAG with 2 links can enable up to 4 Mbps of traffic (2 Mbps per link).

Shared shaping on an S-VLAN within a LAG has the same behavior; the LAG S-VLAN that is shared shaped to 10 Mbps on a LAG with 2 ports allows up to 20Mbps of traffic; 10 Mbps for each link.

PPPoE over VLANs and Subscriber Load Balancing

The system binds PPPoE subscribers stacked over a common VLAN to the same link within the LAG. Because the underlying VLAN node is allocated to a single link, the system allocates all traffic over that VLAN to that link.

PPPoE over Ethernet (No VLANs) and Subscriber Load Balancing

The system allocates subscribers to each link independently. There are no S-VLAN nodes to clone, and no related VLAN nodes that require allocation on the same link.

MPLS over LAG and Subscriber Load Balancing

For QoS purposes, the system considers base tunnels as logical interfaces, but does not consider stacked tunnels. The system assigns MPLS base tunnels stacked over VLANs to the link to which the VLAN is assigned.

Sample Scheduler Hierarchy for Subscriber Load Balancing

Figure 10 on page 43 displays the scheduler hierarchy for the Gigabit Ethernet interface in slot 3, port 0. Figure 11 on page 44 displays the scheduler hierarchy for the Gigabit Ethernet interface in slot 3, port 1.

The Ethernet queues are shown in gray; they are not bound to a link in the LAG and are replicated for each link in the LAG. These queues are used for subscribers with QoS profiles that indicate Ethernet queues, and for traffic classes other than best-effort, traffic class 1, and traffic class 2.

When partitioning the scheduler hierarchy that includes 1000 VLAN subinterfaces, the system binds 500 of the subinterfaces to port 0, and binds another 500 to port 1. The binding for a given VLAN subinterface is arbitrary.

The scheduler nodes for a given VLAN subinterface are always allocated to the same port within the LAG. In this example, the scheduler nodes for VLAN 0,0 are all allocated to Gigabit Ethernet interface in slot 3, port 0.

S-VLAN nodes and queues are cloned for each link in the LAG. S-VLAN nodes in each traffic-class group are shown identically on both ports.

Figure 10: Subscriber LoadBalanced Scheduler Hierarchy for Port 0

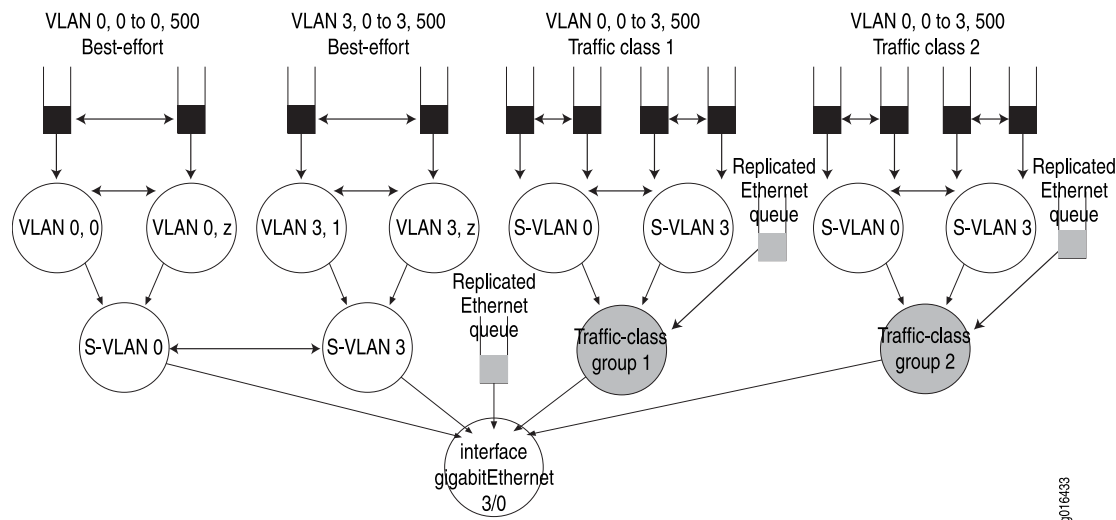
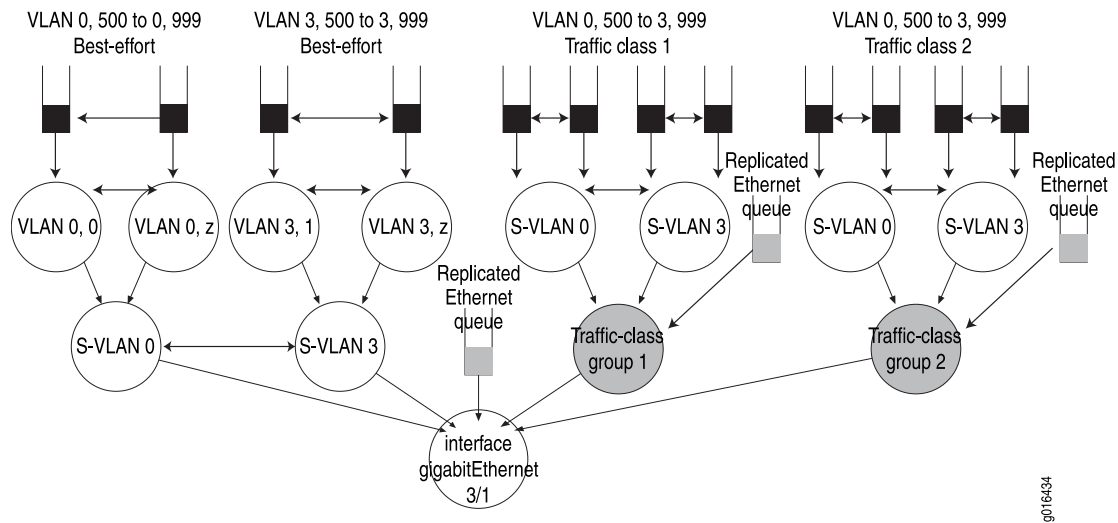


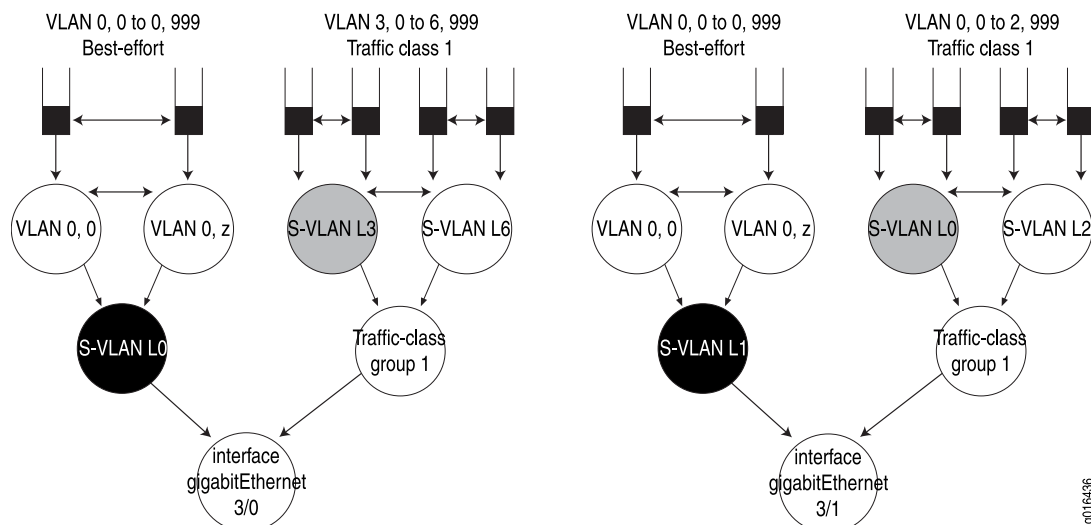
Figure 11: Subscriber LoadBalanced Scheduler Hierarchy for Port 1


Subscriber Allocation in 802.3ad Link Aggregation Groups

You can configure upper-layer subinterfaces over a LAG interface, including VLANs, PPPoE, and MPLS.

The system balances any upper-layer subinterfaces so that each active link in the LAG carries an equal number of upper-layer subinterfaces. For this purpose, the system counts each upper-layer interface as a single subscriber, regardless of the number of forwarding interfaces stacked above it.

Figure 12 on page 44 displays a sample allocation of subscribers.

Figure 12: Subscriber Allocation and Load Balancing


In an ideal QoS configuration, queues and nodes are stacked over a single port that corresponds to a LAG, with the port bandwidth equal to the sum of the overall port bandwidth.

However, the actual LAG behavior is different. No level 1 node or queue can exceed the bandwidth of a link. The relative weighting of queues and nodes results in proportional bandwidth allocation only within a link, but not across the entire LAG. Actual traffic might not be evenly balanced across links in the LAG, resulting in latency and loss on one link, while another link may be lightly loaded.

Even though relative weighting is different on a LAG, shaping and shared shaping in the partitioned scheduler hierarchy operate in the same way as a typical Ethernet configuration.

**Related
Documentation**

- [Configuring Load Rebalancing for 802.3ad Link Aggregation Groups on page 71](#)

PART 2

Configuration

- [Configuration Overview on page 49](#)
- [Configuration Tasks for Classifying, Queuing, and Dropping Triple Play Traffic on page 51](#)
- [Configuration Tasks for Scheduling and Shaping Triple Play Traffic on page 57](#)
- [Configuration Tasks for QoS on Ethernet Interfaces on page 67](#)
- [Configuration Tasks for Hierarchical and Multicast Bandwidth Adjustment on page 75](#)
- [Configuration Tasks for Statistics Collection on page 85](#)
- [Examples on page 89](#)
- [Configuration Commands on page 121](#)

CHAPTER 8

Configuration Overview

- [Configuring QoS on the E Series Router on page 49](#)

Configuring QoS on the E Series Router

Several of the tasks for configuring QoS on your E Series router are optional.

To configure QoS on your E Series router:

1. Create and configure a traffic class.
[See “Traffic Class and Traffic-Class Groups Overview” on page 13.](#)
2. (Optional) Create one or more traffic-class groups.
[See “Traffic Class and Traffic-Class Groups Overview” on page 13.](#)
3. (Optional) To configure nondefault buffer management, create a queue profile.
[See “Queuing and Buffer Management Overview” on page 14.](#)
4. (Optional) To configure RED or WRED, create a drop profile.
[See “Dropping Behavior Overview” on page 16.](#)
5. (Optional) To gather rate statistics, create a statistics profile.
[See “QoS Statistics Overview” on page 17.](#)
6. Configure a scheduler hierarchy with a scheduler profile.
[See “Scheduler Hierarchy Overview” on page 21.](#)
7. (Optional) Configure shaping:
 - Configure shaping and shared shaping using the scheduler profile.
[See Rate Shaping and Port Shaping Overview, “Simple Shared Shaping Overview” on page 25, and Compound Shared Shaping Overview.](#)
 - Configure shaping rates independent of the QoS profile and scheduler profile using QoS parameters.
[See Parameter Definition Attributes for QoS Administrators Overview.](#)
8. Create a QoS profile. QoS profiles reference queue, drop, statistics, and scheduler profiles.

See [“Queuing and Buffer Management Overview”](#) on page 14.

9. Attach the QoS profile to one or more interfaces, or specify the profile as a QoS port-type profile for a given interface type.

See [“Queuing and Buffer Management Overview”](#) on page 14.

CHAPTER 9

Configuration Tasks for Classifying, Queuing, and Dropping Triple Play Traffic

- [Configuring Traffic Classes That Define Service Levels on page 51](#)
- [Configuring Traffic-Class Groups That Define Service Levels on page 52](#)
- [Configuring Queue Profiles to Manage Buffers and Thresholds on page 53](#)
- [Configuring WRED on page 54](#)

Configuring Traffic Classes That Define Service Levels

The router supports up to eight global traffic classes. Each traffic class can appear in only one traffic-class group. If not explicitly added to a traffic-class group, the traffic class is considered to be ungrouped.

To configure a traffic class:

1. Create a traffic class by assigning a name that represents the type of service and enter Traffic Class Configuration mode.

```
host1(config)#traffic-class low-loss1
host1(config-traffic-class)#
```

The traffic class name can be up to 31 characters. It cannot include spaces.

2. (Optional) Specify strict-priority scheduling across the fabric for queues in the traffic class.

```
host1(config-traffic-class)#fabric-strict-priority
```

3. (Optional) For Juniper Networks ERX1440, E120, and E320 Broadband Services Routers, specify the relative weight for queues in the traffic class in the fabric.

```
host1(config-traffic-class)#fabric-weight 12
```

Fabric weight controls the bandwidth of fabric queues associated with the traffic class. It does not control the weight of egress queues associated with the traffic class. If multiple traffic classes are strict priority, the fabric weight determines which class gets more bandwidth.

The weight value is in the range 1–63. The default is 8. Zero is not a valid weight.

- Related Documentation**
- [Monitoring Traffic Classes and Traffic-Class Groups for Defined Levels of Service](#)
 - [fabric-strict-priority on page 134](#)
 - [fabric-weight on page 135](#)
 - [traffic-class on page 155](#)

Configuring Traffic-Class Groups That Define Service Levels

You can configure a traffic-class group and enter Traffic Class Group Configuration mode, from which you can add classes to or delete classes from the group.

Each traffic class can appear in only one traffic-class group. If not explicitly added to a traffic-class group, the traffic class is considered to be ungrouped.

To configure a traffic-class group:

1. Create a traffic-class group by assigning a name that represents the type of service and enter Traffic Class Group Configuration mode.

```
host1(config)#traffic-class-group assured slot 9 extended
host1(config-traffic-class-group)#
```

The traffic class name can be up to 31 characters. It cannot include spaces.

If you do not specify a keyword, the group is strict-priority by default.

You can use the **auto-strict-priority** keyword to explicitly configure a single traffic-class group with strict-priority scheduling, regardless of the scheduler profile associated with the group node.

You can use the **extended** keyword to configure up to three extended traffic-class groups. Scheduling for these groups is determined by the scheduler profile associated with the group node. If an explicitly configured strict-priority group exists, the scheduler for the extended groups may not specify strict-priority scheduling.

Use the **slot slotNumber** option to associate a pre-existing global traffic-class group with the module occupying that slot. Characteristics configured for the local group on the line module override those of the global group.

2. Add traffic classes to the traffic-class group.

```
host1(config-traffic-class-group)#traffic-class low-latency-traffic-class
```

- Related Documentation**
- [Configuring Traffic Classes That Define Service Levels on page 51](#)
 - [Monitoring Traffic Classes and Traffic-Class Groups for Defined Levels of Service](#)
 - [traffic-class on page 155](#)
 - [traffic-class-group on page 156](#)

Configuring Queue Profiles to Manage Buffers and Thresholds

A queue profile controls the buffering and dropping behavior of a set of egress queues by enabling you to set the buffer weight of the queue, the drop thresholds, and the constraints on queue lengths.

Set the queue lengths as follows:

- To oversubscribe buffer memory, set a minimum queue length.



NOTE: If the sum of the queue minimum lengths is greater than the amount of egress buffer memory, then the egress buffer memory is oversubscribed.

- To configure a minimal level of buffering or to limit the buffering in queues, set a maximum queue length. For example, if you want to control latency by configuring very small queues, set the maximum queue length to 256 bytes. The system queues no more than 256 bytes.

If you do not set the queue lengths, the router varies the queue length dynamically in the range 1 KB–7 MB.

- Create a queue profile and enter Queue Configuration mode.

```
host1(config)#queue-profile video
host1(config-queue)#
```

You can configure 16 queue profiles on an E Series router.

- (Optional) Set the buffer weight of the queue.

```
host1(config-queue)#buffer-weight 16
```

Queues with a buffer weight of 16 are twice as long as queues with a buffer weight of 8. The range is 1–63; the default is 8.

- (Optional) Set a minimum or maximum queue length for committed packets.

```
host1(config-queue)#committed-length 11000 15000
```

The range of minimum and maximum lengths is 0–1 GB. By default, there is no minimum or maximum length. The color for committed packets is green.

- (Optional) Set a minimum or maximum queue length for conformed packets.

```
host1(config-queue)#conformed-length 10000 14000
```

The range of minimum and maximum lengths is 0–1 GB. By default, there is no minimum or maximum length. The color for conformed packets is yellow.

- (Optional) Set a minimum or maximum queue length for exceeded packets.

```
host1(config-queue)#exceeded-length 9000 10000
```

The range of minimum and maximum lengths is 0–1 GB. By default, there is no minimum or maximum length. The color for exceeded packets is red.

6. (Optional) Set the conformed drop threshold as a percentage of the committed threshold.

```
host1(config-queue)#conformed-fraction 60
```

The range is 0–100 percent; the default is 50.

7. (Optional) Set the exceeded drop threshold as a percentage of the committed threshold.

```
host1(config-queue)#exceeded-fraction 40
```

The range is 0–100 percent; the default is 25.

Related Documentation

- [Queuing and Buffer Management Overview on page 14](#)
- [Guidelines for Managing Queue Thresholds](#)
- [Guidelines for Managing Buffers](#)
- [Memory Requirements for Queue and Buffers on page 10](#)
- [buffer-weight on page 123](#)
- [committed-length on page 124](#)
- [conformed-fraction on page 126](#)
- [conformed-length on page 127](#)
- [exceeded-fraction on page 131](#)
- [exceeded-length on page 132](#)
- [queue-profile on page 148](#)

Configuring WRED

The main difference between RED and WRED is that WRED deals with different colored packets. The router assigns a color to each packet. Committed means green, conformed means yellow, and exceeded means red.

Each line module supports a default drop profile and 15 configurable drop profiles.

WRED is not supported on the ES2 10G Uplink LM. On the ES2 10G LM, you must configure WRED in one of the 15 configurable drop profiles; you cannot configure its default drop profile.

To enable support for 32,000 subscribers with 128,000 QoS queues on ES2 10G ADV LMs, scheduler memory enhancements have reduced the number of QoS rate counters that are supported per egress queue from 7 to 5:

- 1 is used for forwarding events
- 3 are used for tail dropping behavior

- 1 is used for WRED functionality (an aggregate of all colors)

Each line module supports a default drop profile and 15 configurable drop profiles. On the ES2 10G ADV LM, you must configure WRED in one of the 15 configurable drop profiles; you cannot configure its default drop profile. Queue rate statistics measure the forwarding and drop rates of each queue in bits per second. Queue event statistics configure the E Series router to count the number of times that forwarding or drop rates exceed a specific threshold. To display information about the number of committed packets and bytes dropped by WRED for ES2 10G ADV LMs, see the number displayed in the Dropped by WRED committed field in the output of the **show ip interface** command. The Dropped by WRED confirmed and Dropped by WRED exceeded fields always display a value of zero because of the single counter being used for WRED functionality being calculated and displayed in the Dropped by WRED committed field of the output.

To configure WRED:

1. Create a drop profile and enter Drop Profile Configuration mode.

```
host1(config)#drop-profile internetDropProfile
host1(config-drop-profile)#
```

You can configure up to 16 drop profiles.

2. Set the average-length exponent, which specifies the exponent used to weight the average queue length over time, controlling WRED responsiveness.

```
host1(config-drop-profile)#average-length-exponent 9
```

- Specifying an average-length exponent enables the RED average queue length computation.
- A higher value smooths out the average and slows WRED reaction to congestion and decongestion, accommodating short bursts without dropping. Too large a value can smooth the average to the point that WRED does not react at all.
- A lower value speeds up WRED reaction. Too low a value can cause overreaction to short bursts, dropping packets unnecessarily.

3. (Optional) Set the minimum and maximum threshold for committed traffic.

```
host1(config-drop-profile)#committed-threshold percent 30 90 4
```

4. (Optional) Set the minimum and maximum threshold for conformed traffic.

```
host1(config-drop-profile)#conformed-threshold percent 25 90 5
```

5. (Optional) Set the minimum and maximum threshold for exceeded traffic.

```
host1(config-drop-profile)#exceeded-threshold percent 20 90 6
```

The thresholds specify a linear relationship between average queue length and drop probability.

You can express thresholds as either percentages of maximum queue size by including the keyword **percent**, or as absolute byte values by omitting the keyword.

**Related
Documentation**

- [Configuring RED](#)
- [Monitoring RED and WRED](#)
- [average-length-exponent on page 122](#)
- [committed-threshold on page 125](#)
- [conformed-threshold on page 128](#)
- [drop-profile on page 130](#)
- [exceeded-threshold on page 133](#)

CHAPTER 10

Configuration Tasks for Scheduling and Shaping Triple Play Traffic

- [Configuring a Scheduler Hierarchy on page 57](#)
- [Configuring Strict-Priority Scheduling on page 58](#)
- [Configuring Relative Strict-Priority Scheduling for Aggregate Shaping Rates on page 59](#)
- [Configuring a QoS Profile on page 61](#)
- [Attaching a QoS Profile to an Interface on page 62](#)

Configuring a Scheduler Hierarchy

When you configure a scheduler hierarchy, you configure the scheduler profile and assign attributes.

To configure a scheduler hierarchy:

1. Configure a scheduler profile.
See [Configuring a Scheduler Profile for a Scheduler Node or Queue](#).
2. (Optional) Configure attributes in the scheduler profile.
 - Configure a shaping rate for rate shaping or port shaping.
See [Configuring Rate Shaping for a Scheduler Node or Queue](#) or [Configuring Port Shaping](#).
 - Configure an assured rate.
See [Configuring an Assured Rate for a Scheduler Node or Queue](#).
 - Configure the HRR weight.
See [Configuring the HRR Weight for a Scheduler Node or Queue](#).
 - Configure shared shaping.
See [Configuring Simple Shared Shaping](#) and [Configuring Compound Shared Shaping](#).
 - Configure implicit and explicit constituent selection.

See [Configuring Implicit Constituents for Simple or Compound Shared Shaping](#) and [Configuring Explicit Constituents for Simple or Compound Shared Shaping](#).

3. Reference the scheduler profile in a QoS profile and apply to an interface.

See ["Configuring a QoS Profile" on page 61](#) and ["Attaching a QoS Profile to an Interface" on page 62](#).

Related Documentation

- [Scheduler Hierarchy Overview on page 21](#)
- [Parameter Definition Attributes for QoS Administrators Overview](#)

Configuring Strict-Priority Scheduling

To configure strict-priority scheduling:

1. Configure the traffic classes.

```
host1(config)#traffic-class Low-loss-1
host1(config-traffic-class)#exit
host1(config)#traffic-class Low-latency-1
host1(config-traffic-class)#exit
host1(config)#traffic-class Low-latency-2
host1(config-traffic-class)#exit
```

2. Configure the auto-strict-priority traffic-class group, and add the traffic classes that must receive strict-priority scheduling to the group.

```
host1(config)#traffic-class-group Strict-priority auto-strict-priority
host1(config-traffic-class-group)#traffic-class Low-latency-1
host1(config-traffic-class-group)#traffic-class Low-latency-2
host1(config-traffic-class-group)#exit
```

3. Create a scheduler profile for strict-priority traffic and configure the shaping rate.

```
host1(config)#scheduler-profile strictPriorityBandwidth
host1(config-scheduler-profile)#shaping-rate 20000000
host1(config-scheduler-profile)#exit
```

4. Configure a QoS profile.

```
host1(config)#qos-profile Example-qos-profile
host1(config-qos-profile)#atm group default
host1(config-qos-profile)#atm group Strict-priority scheduler-profile
strictPriorityBandwidth
host1(config-qos-profile)#atm-vc node group default
host1(config-qos-profile)#atm-vc node group Strict-priority
host1(config-qos-profile)#atm-vc queue traffic-class best-effort
host1(config-qos-profile)#atm-vc queue traffic-class Low-loss-1
host1(config-qos-profile)#atm-vc queue traffic-class Low-latency-1
host1(config-qos-profile)#atm-vc queue traffic-class Low-latency-2
host1(config-qos-profile)#exit
```

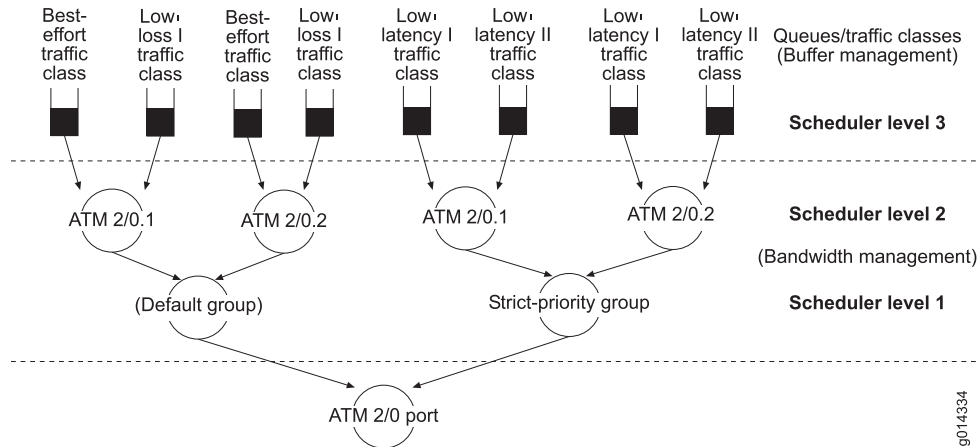
5. Attach the QoS profile to an interface.

```
host1(config)#interface atm 2/0
host1(config-if)#qos-profile Example-qos-profile
```

```
host1(config-if)#exit
host1(config)#
```

This configuration creates the hierarchy shown in [Figure 13 on page 59](#).

Figure 13: Sample Strict-Priority Scheduling Hierarchy



Related Documentation

- [Strict-Priority and Relative Strict-Priority Scheduling Overview](#)
- [Using Expressions for Bandwidth and Burst Values in a Scheduler Profile](#)
- [group](#)
- [node on page 139](#)
- [qos-profile on page 145](#)
- [queue on page 147](#)
- [scheduler-profile on page 149](#)
- [shaping-rate on page 150](#)
- [strict-priority](#)
- [traffic-class on page 155](#)
- [traffic-class-group on page 156](#)

Configuring Relative Strict-Priority Scheduling for Aggregate Shaping Rates

To configure relative strict priority scheduling for aggregate shaping rates:

1. Create a scheduler profile for the strict-priority queue.

```
host1(config)# scheduler-profile relativeStrict
host1(config-scheduler-profile)# shaping-rate 500000
host1(config-scheduler-profile)# weight 0
host1(config-scheduler-profile)# exit
```

Configuring the weight of 0 reduces latency and jitter.

2. Create a scheduler profile for the nonstrict best-effort queue.

```

host1(config)# scheduler-profile be
host1(config-scheduler-profile)# shaping-rate 1000000 burst 1
host1(config-scheduler-profile)# weight 8
host1(config-scheduler-profile)# exit

```



TIP: If you need to impose a shaping rate on the nonstrict queues to meet a functional requirement, you can specify a rate less than the aggregate rate. The key is that the burst size must be one, or small. The burst size determines the maximum-sized packet that can squeeze in front of a relative strict-priority packet in the round robin.

3. Create a scheduler profile for the aggregate bandwidth.

```

host1(config)# scheduler-profile vcAggregate
host1(config-scheduler-profile)# shaping-rate 1000000
host1(config-scheduler-profile)# exit

```

4. Create a QoS profile, configure node shaping for each queue, and add each of the queues to the QoS profile.

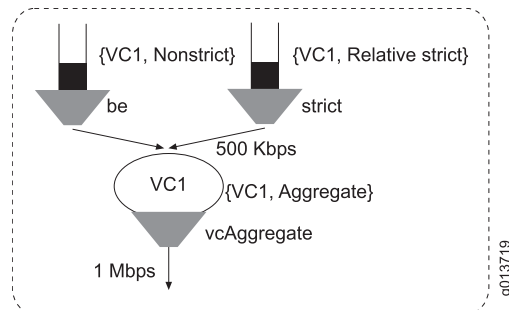
```

host1(config)# qos-profile relative-strict-aggregate
host1(config-qos-profile)# atm-vc node scheduler-profile vcAggregate
host1(config-qos-profile)# atm-vc queue traffic-class best-effort
scheduler-profile be
host1(config-qos-profile)# atm-vc queue traffic-class voice scheduler-profile
relativeStrict
host1(config-qos-profile)# exit
host1(config)#

```

This configuration creates the hierarchy shown in [Figure 14 on page 60](#).

Figure 14: Sample Relative Strict-Priority Scheduler Hierarchy



Related Documentation

- [Strict-Priority and Relative Strict-Priority Scheduling Overview](#)
- [Using Expressions for Bandwidth and Burst Values in a Scheduler Profile](#)
- [node on page 139](#)
- [qos-profile on page 145](#)
- [scheduler-profile on page 149](#)
- [shaping-rate on page 150](#)

- weight

Configuring a QoS Profile

Before you configure a QoS profile:

- Configure the traffic classes.
See “Configuring Traffic Classes That Define Service Levels” on page 51.
- Configure the queuing hierarchy.
See “Configuring Queue Profiles to Manage Buffers and Thresholds” on page 53.
- Configure the scheduler hierarchy and shaping with scheduler profiles.
See “Configuring a Scheduler Hierarchy” on page 57.

To configure a QoS profile:

1. Create a QoS profile and enter QoS Profile Configuration mode.

```
host1(config)#qos-profile qosp-vc-queuing
host1(config-qos-profile)#
```

2. (Optional) Configure a group node for each interface.

```
host1(config-qos-profile)#atm group groupA scheduler-profile scheduler1
statistics-profile statpro-1
```

When you configure a group node, you can also reference a default or named traffic-class group, a scheduler profile, or a statistics profile.

If you do not specify a traffic-class group, the group node defaults to the *default* group. Each traffic class can belong to only one traffic-class group (either the default group or a named group).

The router supports up to four traffic-class groups above a given port.

3. (Optional) Configure a scheduler node for interfaces.

```
host1(config-qos-profile)#atm node scheduler-profile scheduler1 group strict-priority
```

When you configure a scheduler node, you can also reference a default or named traffic-class group and a scheduler profile.

The scheduler profile supplies a relative weight and potentially a shaping rate to be applied at the scheduler node.



NOTE: You cannot associate a scheduler profile with a port-type interface unless you also specify the strict-priority group.

4. (Optional) Configure a queue for interfaces in the specified traffic class.

```
host1(config-qos-profile)#atm queue traffic-class strict-priority scheduler-profile
scheduler1 queue-profile queue1 drop-profile drop1
```

When you configure a queue, you can include any of the following profiles:

- The scheduler profile supplies a relative weight and potentially a shaping rate to be applied at the queue.
- The queue profile supplies threshold information for the queue if the router defaults are not appropriate.
- The drop profile supplies dropping behavior of a set of egress queues.

Each queue traffic class can appear in only one traffic-class group.

Related Documentation

- [Attaching a QoS Profile to an Interface on page 62](#)
- Supported Interface Types for QoS Profiles
- Configuring Shadow Nodes
- [Monitoring a Scheduler Hierarchy on an Interface with QoS Profiles on page 184](#)
- *JunosE Broadband Access Configuration Guide*
- group
- [node on page 139](#)
- [qos-profile on page 145](#)
- [queue on page 147](#)

Attaching a QoS Profile to an Interface

You can attach a QoS profile to the base of an interface hierarchy, to a specific ATM VP or S-VLAN, or to a port type.

Tasks to attach a QoS profile include:

- [Attaching a QoS Profile to a Base Interface on page 62](#)
- [Attaching a QoS Profile to an ATM VP on page 63](#)
- [Attaching a QoS Profile to an S-VLAN on page 64](#)
- [Attaching a QoS Profile to a Port Type on page 64](#)

Attaching a QoS Profile to a Base Interface

You can attach a QoS profile to an interface at the base of an interface hierarchy. Interface types below the attachment point cannot be referenced in the QoS profile.

To attach a profile to an interface:

1. Enter Interface Configuration mode for the interface.

```
host1(config)#interface gigabitEthernet 2/0
```

2. Attach a QoS profile to the interface.

```
host1(config-if)#qos-profile qosp-ethernet-queuing
```

Attaching a QoS Profile to an ATM VP

You can associate a QoS profile with all the ports of a certain interface type.

You can attach a QoS profile to an ATM VP. The profile applies to all VCs in the VP; for example, the profile specifies the scheduler hierarchy of scheduler nodes and queues for all VCs, IP interfaces, and L2TP sessions stacked above the VP.

To attach a profile to an ATM VP:

1. Enter Interface Configuration mode for the interface.

```
host1(config)#interface atm 1.0/1
```

2. Attach a QoS profile to the ATM VP.

```
host1(config-if)#atm-vp 50 qos-profile qosp-vp-strictbw
```

If you attempt to modify the QoS profile attached to an ATM VP that contains nonbroadcast multiaccess (NBMA) or multipoint interfaces from profileA to profileB by using the **atm-vp qos-profile** command for a specific VP on that interface, the command is configured correctly and no error message is displayed in the CLI interface. However, the shaping rate on the interfaces that are part of the ATM VP is not properly updated with the shaping rate specified in profileB. Instead, the multipoint interfaces remain configured with the shaping rate set in profileA.

To modify the QoS profile currently attached to ATM VPs that contain NBMA or multipoint interfaces from another profile, you must first remove the QoS profile attached to the interfaces by using the **no atm-vp qos-profile** command in Interface Configuration mode, and then attach the new QoS profile to the interfaces by using the **atm-vp qos-profile** command. This restriction exists because the mungeing of QoS profiles does not occur correctly if any of the attributes of ATM VPs with multipoint interfaces are modified.

If you modify the QoS profile attached to a point-to-point ATM interface from profileA to profileB by using the **qos-profile** command (or the **atm-vp qos-profile** command for a specific VP on the ATM interface) in Interface Configuration mode, the shaping rate is correctly configured on the interface and is modified with the value specified in profileB.

To modify the QoS profile attached to an ATM VP that contains an NBMA or a multipoint interface from profileA to profileB, perform the following steps. These steps assume that profileA and profileB have been previously configured on the router.

1. Enter Interface Configuration mode for the ATM VP.

```
host1(config)#interface atm 1/0
```

2. Remove the QoS profile, profileA, currently attached to the ATM VP that contains the NBMA interface.

```
host1(config-if)#no atm-vp 1 qos-profile profileA
```

3. Attach the new QoS profile, profileB, that you want to be attached to the ATM VP that contains the NBMA interface.

```
host1(config-if)#atm-vp 1 qos-profile profileB
```

Attaching a QoS Profile to an S-VLAN

You can attach a QoS profile to the specified S-VLAN ID assigned to a VLAN subinterface that is configured over an Ethernet interface.

The profile applies to all S-VLANs and VLANs in the interface stack; for example, the profile specifies the hierarchy of scheduler nodes and queues for all VLANs, IP interfaces stacked above the S-VLAN. However, you do not have to configure VLAN subinterfaces over the S-VLAN before you attach the QoS profile to the S-VLAN.

1. Specify the Ethernet interface and create the VLAN.

```
host1(config)#interface gigabitEthernet 3/0
host1(config-if)#encapsulation vlan
host1(config-if)#interface gigabitEthernet 3/0.1
```

2. Specify the S-VLAN ID.

```
host1(config-if)#svlan id 0 1
```

3. Attach the QoS profile to the S-VLAN.

```
host1(config-if)#svlan 1 qos-profile qosp-svlan-strictbw
```

Attaching a QoS Profile to a Port Type

By default, the router attaches a QoS port-type profile to all ATM, Ethernet, serial, or server ports. The port-type profile supplies QoS information for all forwarding interfaces stacked above all ports of the associated interface type.

Instead of using the default port-type profile, you can explicitly attach a QoS profile to a port. The QoS profile overrides the default QoS port-type profile. The QoS profile associates queue profiles, drop profiles, statistics profiles, and scheduler profiles with interface types, and it applies to all interfaces stacked above ports of the associated type.

To attach a QoS profile to a port type:

- Issue the **qos-port-type-profile** command from Global Configuration mode:

```
host1(config)#qos-port-type-profile atm qos-profile strict-priority
```

The interface type can be: atm, ethernet, lag, serial, or server-port.

A profile attached to a port must specify a queue for each forwarding interface type in the best-effort traffic class.

To restore the default port-type:

- Issue the **qos-port-type-profile** command and specify the server-default QoS profile from Global Configuration mode:

```
host1(config)#qos-port-type-profile server-port qos-profile server-default
```

Related Documentation

- Supported Interface Types for QoS Profiles

- [Configuring a QoS Profile on page 61](#)
- *JunosE Broadband Access Configuration Guide*
- atm-vp qos-profile
- atm vp-tunnel
- encapsulation vlan
- interface
- [qos-port-type-profile on page 144](#)
- [qos-profile on page 145](#)
- svlan id
- svlan qos-profile

Configuration Tasks for QoS on Ethernet Interfaces

- [Creating a QoS Interface Hierarchy for Bulk-Configured VLAN Subinterfaces with RADIUS on page 67](#)
- [Configuring the Scheduler Hierarchy for Hashed Load Balancing in 802.3ad Link Aggregation Groups on page 70](#)
- [Enabling Default Subscriber Load Balancing for 802.3ad Link Aggregation Groups on page 70](#)
- [Configuring the Scheduler Hierarchy for Subscriber Load Balancing in 802.3ad Link Aggregation Groups on page 71](#)
- [Configuring Load Rebalancing for 802.3ad Link Aggregation Groups on page 71](#)

Creating a QoS Interface Hierarchy for Bulk-Configured VLAN Subinterfaces with RADIUS

Bulk-configured VLAN subinterfaces are created dynamically, so you cannot apply a QoS profile directly to a VLAN subinterface. Instead, you can use subscriber service profiles and RADIUS to apply QoS profiles.

To create an interface hierarchy for bulk-configured VLAN subinterfaces:

1. Configure the bulk-configured VLAN subinterface.

```
host1(config)#interface gigabitEthernet 6/0/0
host1(config-if)#encapsulation vlan
host1(config-if)#auto-configure vlan
host1(config-if)#vlan bulk-config BulkConfig
host1(config-if)#profile vlan bulk-config BulkConfig vlanBulkProfile
host1(config-if)#vlan bulk-config BulkConfig vlan-range 1 3600
```

2. Configure the profiles and service profile for the bulk-configured VLAN subinterfaces and the IP upper-layer encapsulation.

```
host1(config-if)#profile vlanBulkProfile
host1(config-profile)#vlan auto-configure ip
host1(config-profile)#vlan profile ip ipProfile
host1(config-profile)#vlan service-profile vlanServiceProfile
host1(config-profile)#exit
host1(config-profile)#profile ipProfile
```

```
host1(config-profile)#ip unnumbered loopback 0
host1(config-profile)#exit
```

3. Configure an IP service profile.

```
host1(config)#ip service-profile vlanServiceProfile
host1(config-service-profile)#user-name "vlan@test"
host1(config-service-profile)#password 56789
host1(config-service-profile)#exit
```



TIP: Configure the service profile in the default virtual router or the virtual router in which RADIUS is configured.

4. Access the RADIUS server and assign values for the RADIUS attributes necessary for creating a QoS interface hierarchy, including the QoS profile name. For example:

- Juniper VSA Qos-Profile-Name [26-26]—QoS profile name
- (Optional) Juniper VSA Virtual-Router [26-1]—Virtual router name
- (Optional) IETF VSA [22]—Framed-Route

5. Verify that the attributes are being used by RADIUS.

The highlighted output from this debug log message shows the QoS profile, virtual router, and framed route attributes configured through RADIUS.

```
DEBUG 06/17/2007 14:50:19 radiusSendAttributes: ACCESS-REQUEST attributes (default)
```

```
DEBUG 06/17/2007 14:50:19 radiusSendAttributes:      username attr added: vlan@test
DEBUG 06/17/2007 14:50:19 radiusSendAttributes:      acct-session-id attr added: erxGigabitEthernet 2/1.100:100:0004194348
DE BUG 06/17/2007 14:50:19 radiusSendAttributes:      user-password attr added: <value withheld>
DEBUG 06/17/2007 14:50:19 radiusSendAttributes:      calling-station-id attr added: #ananke#E21#100
DEBUG 06/17/2007 14:50:19 radiusSendAttributes:      nas-port-type attr added: 15
DEBUG 06/17/2007 14:50:19 radiusSendAttributes:      nas-port attr added: 553648228
DEBUG 06/17/2007 14:50:19 radiusSendAttributes:      nas-port-id attr added: GigabitEthernet 2/1.100:100
DEBUG 06/17/2007 14:50:19 radiusSendAttributes:      nas-ip-address attr added: 172.26.27.50
DEBUG 06/17/2007 14:50:19 radiusSendAttributes:      nas-identifier attr added: ananke
DEBUG 06/17/2007 14:50:19 radiusAttributes: USER ATTRIBUTES: (vlan@test)
DEBUG 06/17/2007 14:50:19 radiusAttributes:      class attr: (binary data)
DEBUG 06/17/2007 14:50:19 radiusAttributes: total eap message attr length = 0
DEBUG 06/17/2007 14:50:19 radiusAttributes:      framed route attr: 40.40.41.0/30 0.0.0.0
DEBUG 06/17/2007 14:50:19 radiusAttributes:      ingress policy name (vsa) attr: test
DEBUG 06/17/2007 14:50:19 radiusAttributes:      ingress policy stats (vsa) attr: 1
DEBUG 06/17/2007 14:50:19 radiusAttributes:      egress policy name (vsa) attr: test
DEBUG 06/17/2007 14:50:19 radiusAttributes:      egress policy stats (vsa) attr: 1
DEBUG 06/17/2007 14:50:19 radiusAttributes:      qos profile name (vsa) attr: test
DEBUG 06/17/2007 14:50:19 radiusAttributes:      virtual router name (vsa) attr: server
```

6. Verify that the interface was created in the default virtual router.

```
host1:server# show ip interface brief
```

Interface	IP-Address	Status	Protocol	Description
-----	-----	-----	-----	-----
nu110	255.255.255.255/32	up	up	
loopback0	10.1.0.1/24	up	up	
GigabitEthernet6/0.100	Unnumbered	up	up	

7. Verify that the framed route is installed.

```
host1:server# show ip route
```

Prefix/Length	Type	Next Hop	Dst/Met	Interface
10.1.0.0/24	Connect	10.1.0.1	0/0	loopback0
40.40.41.0/30	Access	0.0.0.0 3/2	GigabitEthernet6/0/0.100	



TIP: When you initially create the user record for dynamic IP interfaces using VSA [22], you might not know the next hop. In this case, specify the value 0.0.0.0 for the next hop. The E Series router then assigns the subinterface associated with the user as the next hop in the routing table.

8. Verify that the correct QoS profile is attached to the VLAN subinterface.

```
host1:server#show qos interface-hierarchy interface gigabitEthernet
6/0/0.100
```

```
attachment@ ip GigabitEthernet6/0/0.100:
```

		t-class	interface	rule	traffic	scheduler	queue
qos profile	group	type	type	class	profile	profile	
test@GigabitEthernet6/0/0.100	vlan	node	default	default			

Related Documentation

- For information about bulk-configured VLAN subinterfaces, see *JunosE Link Layer Configuration Guide*
- Juniper Networks VSAs
- Understanding Subscriber Management
- auto-configure vlan
- encapsulation vlan
- interface gigabitEthernet
- ip service-profile
- profile
- profile vlan bulk-config
- vlan auto-configure
- vlan bulk-config
- vlan profile
- vlan service-profile
- show ip interface
- show ip route
- [show qos interface-hierarchy on page 203](#)

Configuring the Scheduler Hierarchy for Hashed Load Balancing in 802.3ad Link Aggregation Groups

The type of load balancing that the system performs depends on the configuration of the scheduler hierarchy in the QoS profile.

To configure the scheduler hierarchy for hashed load balancing:

1. Configure a QoS profile.

```
host1(config)#qos-profile hashed-lag
```

2. Configure the nodes and queues, including an Ethernet queue.

```
host1(config-qos-profile)#ethernet queue traffic-class best-effort
host1(config-qos-profile)#ethernet queue traffic-class tc1
host1(config-qos-profile)#ethernet queue traffic-class tc2
```

3. Create the LAG interface and attach the QoS profile.

```
host1(config)#interface lag lg1
host1(config-if)#qos-profile hashed-lag
```

Related Documentation

- [QoS for 802.3ad Link Aggregation Interfaces Overview on page 39](#)
- [Hashed Load Balancing for 802.3ad Link Aggregation Groups Overview on page 41](#)
- [interface lag](#)
- [node on page 139](#)
- [qos-profile on page 145](#)
- [queue on page 147](#)

Enabling Default Subscriber Load Balancing for 802.3ad Link Aggregation Groups

The factory default contents of the lag-default QoS profile include an Ethernet queue and the best-effort traffic class.

When you use the lag-default QoS profile, the system automatically sends traffic to the Ethernet queue and uses hash load balancing for the Ethernet queues.

To enable subscriber load balancing as the default behavior for all LAGs, issue the following command:

```
host1(config)#qos-port-type-profile lag qos-profile ethernet-default
```

Related Documentation

- [QoS for 802.3ad Link Aggregation Interfaces Overview on page 39](#)
- [qos-port-type-profile on page 144](#)

Configuring the Scheduler Hierarchy for Subscriber Load Balancing in 802.3ad Link Aggregation Groups

The type of load balancing that the system performs depends on the configuration of the scheduler hierarchy in the QoS profile.

To configure the scheduler hierarchy for subscriber load balancing:

1. Configure the QoS profile.

```
host1(config)#qos-profile subscriber-lag
```

2. Configure the queues and nodes for VLANs and S-VLANs.

```
host1(config-qos-profile)#vlan queue traffic-class best-effort
host1(config-qos-profile)#vlan queue traffic-class tc1
host1(config-qos-profile)#vlan node scheduler-profile subscriber
host1(config-qos-profile)#svlan node scheduler-profile svlan
host1(config-qos-profile)#svlan node group g1 scheduler-profile svlan
```

3. Create the LAG interface and assign member interfaces.

```
host1(config)#interface lag lg1
host1(config-if)#member-interface gigabitEthernet 3/0
host1(config-if)#member-interface gigabitEthernet 3/1
```

4. Attach the QoS profile to the LAG interface.

```
host1(config-if)#qos-profile subscriber-lag
```

Related Documentation

- [QoS for 802.3ad Link Aggregation Interfaces Overview on page 39](#)
- [Subscriber Load Balancing for 802.3ad Link Aggregation Groups Overview on page 42](#)
- [Enabling Default Subscriber Load Balancing for 802.3ad Link Aggregation Groups on page 70](#)
- [interface lag](#)
- [member-interface on page 138](#)
- [node on page 139](#)
- [qos-profile on page 145](#)
- [queue on page 147](#)

Configuring Load Rebalancing for 802.3ad Link Aggregation Groups

You can configure the parameters that the system uses to rebalance the links in a LAG. You can also configure the system to dynamically rebalance the links in the LAG.

Tasks to configure load rebalancing are:

- [Configuring Load–Rebalancing Parameters on page 72](#)
- [Configuring the System to Dynamically Rebalance the LAG on page 73](#)

Configuring Load–Rebalancing Parameters

To configure load–rebalancing parameters:

1. Specify the LAG interface.

```
host1(config)#interface lag lg1
```

2. Configure parameters that guide the system to rebalance.

```
host1(config-if)#load-rebalance period 120 start-threshold 20 percent stop-threshold
100 percent maximum-improvement 300
```

This example specifies that the system rebalance within 120 seconds, can accept imbalance in the LAG in the range 20–100 percent, and can move 300 subscribers to other ports during that time.

[Table 11 on page 72](#) describes the load balancing algorithm parameters that you can configure.

Table 11: Load Balancing Algorithm Parameters

Keyword	Description
period	Specifies the time period for rebalancing. For example, a period of 120 specifies that rebalancing occurs once every 2 minutes.
start-threshold	<p>Specifies the amount of imbalance in the LAG that triggers the algorithm to start rebalancing. The default is 0 percent. Optionally, you can specify one of the following units of measure:</p> <ul style="list-style-type: none"> • percent—Specifies that the amount of imbalance is measured as a percentage of the average load per link. The range is 0–100 percent. For example, the average load per link in a LAG is 500. Specifying start-threshold 5 percent indicates that the algorithm rebalances any link that deviates from the average load per link by 25 (5 percent of 500). • subscribers—Specifies that the amount of imbalance is measured by the number of subscribers from the average subscriber count in the LAG. The range is 0–10000. For example, specifying start-threshold 20 subscribers indicates that the algorithm rebalances any link with a subscriber count that differs from the average subscriber count by more than 20.

Table 11: Load Balancing Algorithm Parameters (*continued*)

Keyword	Description
stop-threshold	<p>Specifies the amount of imbalance in the LAG that triggers the algorithm to stop rebalancing. The algorithm continues rebalancing until this value is reached. The default is 0 percent. Optionally, you can specify one of the following units of measure:</p> <ul style="list-style-type: none"> • percent—Specifies that the amount of imbalance is measured as a percentage of the average load per link. The range is 0–100 percent. For example, the average load per link in a LAG is 500. Specifying the stop-threshold 2 percent command indicates that the algorithm stops within 10 of 500 (2 percent of 500). In this case, the algorithm stops when the links are at 510 and 490. • subscribers—Specifies that the amount of imbalance is measured by the number of subscribers. The range is 0–10000. For example, specifying stop-threshold 100 subscribers indicates that the algorithm continues until each link in the LAG is within 100 subscribers of the average subscriber count.
maximum-improvement	<p>Specifies the maximum number of links to rebalance in the LAG per period. The default is 100 percent. Optionally, you can specify one of the following units of measure:</p> <ul style="list-style-type: none"> • percent—Specifies that the maximum number of links is measured as a percentage of the total links. The range is 0–100 percent. For example, specifying maximum-improvement 1 percent indicates that the algorithm rebalances 10 links per period (1 percent of 1000). • subscribers—Specifies that the maximum number of links is measured by the number of subscribers. The range is 0–10000 subscribers. For example, specifying maximum-improvement 40 subscribers indicates that the algorithm rebalances 40 subscribers per period.

Configuring the System to Dynamically Rebalance the LAG

To configure the system to dynamically rebalance the LAG:

1. Specify the LAG interface.

```
host1(config)#interface lag lg1
```

2. Issue the load balance command with no keywords:

```
host1(config-if)#load-rebalance
```

Related Documentation

- [Configuring Load Rebalancing for 802.3ad Link Aggregation Groups on page 71](#)
- [QoS for 802.3ad Link Aggregation Interfaces Overview on page 39](#)
- [Enabling Default Subscriber Load Balancing for 802.3ad Link Aggregation Groups on page 70](#)
- [Configuring the Scheduler Hierarchy for Subscriber Load Balancing in 802.3ad Link Aggregation Groups on page 71](#)
- `interface lag`

- [load-rebalance on page 137](#)

Configuration Tasks for Hierarchical and Multicast Bandwidth Adjustment

- [Configuring a Basic Parameter Definition for QoS Administrators on page 75](#)
- [Configuring a Parameter Definition to Calculate Hierarchical Instances on page 77](#)
- [Configuring a Parameter Definition for IP Multicast Bandwidth Adjustment on page 79](#)
- [Configuring a Parameter Definition to Shape Ethernet Traffic Using Cell Mode on page 79](#)
- [Configuring a Parameter Definition for QoS Downstream Rate on page 80](#)
- [Creating Parameter Instances on page 82](#)

Configuring a Basic Parameter Definition for QoS Administrators

This section describes how to configure an individual parameter definition and how to associate it with an application.

Several of the following tasks are optional. Perform the required tasks and also any optional tasks that you need for your QoS parameter configuration.

To configure a parameter definition:

1. Create traffic classes.

```
host1(config)#traffic-class business-data
host1(config-traffic-class)#exit
host1(config)#traffic-class voice
host1(config-traffic-class)#exit
host1(config)#traffic-class video
```

2. Create a parameter definition.

- a. Specify the parameter definition name.

```
host1(config)#qos-parameter-define max-subscriber-bandwidth
host1(config-qos-parameter-define)#
```

- b. Specify the logical interface types for the nodes and queues controlled by this parameter.

```
host1(config-qos-parameter-define)#controlled-interface-type atm-vc
host1(config-qos-parameter-define)#controlled-interface-type vlan
```

You can specify up to four of the following controlled-interface types per parameter definition: atm, atm-vc, atm-vp, bridge, ethernet, fr-vc, ip, ip-tunnel, ipv6, l2tp-session, l2tp-tunnel, lsp, pppoe, serial, server-port, vlan.

- c. Specify the set of logical interfaces types upon which a QoS client can create instances of the parameter.

```
host1(config-qos-parameter-define)#instance-interface-type atm-vc
host1(config-qos-parameter-define)#instance-interface-type ip
```

You can specify up to four of the following controlled-interface types per parameter definition: atm, atm-vc, atm-vp, bridge, ethernet, fr-vc, ip, ip-tunnel, ipv6, lag, l2tp-session, l2tp-tunnel, lsp, pppoe, serial, server-port, svlan, vlan.

- d. (Optional) Specify the set of interface types that a QoS client can assign to a parameter instance to represent subscribers.

```
host1(config-qos-parameter-define)#subscriber-interface-type ip
```

You can specify up to four of the following subscriber-interface types: atm-vc, ip, ipv6, l2tp-session, vlan.

- e. (Optional) Define the range of values that a QoS client can assign to a parameter instance.

```
host1(config-qos-parameter-define)#range 64000 8000000
```

3. Reference the parameter within a scheduler profile parameter expression and configure an assured rate, shaping rate, shared-shaping rate, or weight.

```
host1(config)#scheduler-profile business-data
host1(config-scheduler-profile)#shaping-rate max-subscriber-bandwidth % 25
```

4. Add the scheduler profile to a QoS profile and configure the QoS profile.

```
host1(config)#qos-profile subscriber
host1(config-qos-profile)#atm-vc queue traffic-class business-data scheduler-profile
business-data
host1(config-qos-profile)#atm-vc queue traffic-class video scheduler-profile voice
host1(config-qos-profile)#atm-vc queue traffic-class voice scheduler-profile video
```

Related Documentation

- [Parameter Definition Attributes for QoS Administrators Overview](#)
- [Example: QoS Parameter Configuration for Controlling Subscriber Bandwidth on page 89](#)
- [Configuring a Scheduler Hierarchy on page 57](#)
- [Configuring a QoS Profile on page 61](#)
- [assured-rate](#)
- [controlled-interface-type on page 129](#)
- [instance-interface-type on page 136](#)
- [node on page 139](#)
- [qos-parameter-define on page 142](#)

- [qos-profile on page 145](#)
- [queue on page 147](#)
- range
- [scheduler-profile on page 149](#)
- [shaping-rate on page 150](#)
- [shared-shaping-rate on page 153](#)
- subscriber-interface-type
- [traffic-class on page 155](#)
- weight

Configuring a Parameter Definition to Calculate Hierarchical Instances

You can configure hierarchical parameters for applications where you want the system to add instances associated with child interfaces and associate the sum with a parent interface.

Hierarchical parameters have explicit instances that are associated with the logical interfaces of instance-interface types, as well as implicit instances that are associated with the logical interfaces of controlled-interface types. The system computes the values of an implicit instance as the sum of the values of the explicit instances stacked above the implicit instance.

To configure a hierarchical QoS parameter definition:

- Include the **hierarchical** keyword with the **qos-parameter-define** command.

```
host1(config)#qos-parameter-define max-subscriber-bandwidth
host1(config-qos-parameter-define)#
```

Related Documentation

- [Hierarchical QoS Parameters Overview on page 31](#)
- [Configuring a Basic Parameter Definition for QoS Administrators on page 75](#)
- [Configuring a Parameter Definition for IP Multicast Bandwidth Adjustment on page 77](#)
- [Example: QoS Parameter Configuration for Hierarchical Parameters on page 115](#)
- [qos-parameter-define on page 142](#)

Configuring a Parameter Definition for IP Multicast Bandwidth Adjustment

Before you configure a parameter definition for IP multicast bandwidth:

- Define a multicast bandwidth map and the QoS adjustment for a virtual router.

See *JunosE Multicast Routing Configuration Guide*.

To associate a parameter instance with the IP multicast bandwidth adjustment application:

1. Configure traffic classes.

```
host1(config)#traffic-class voice
host1(config-traffic-class)#exit
host1(config)#traffic-class best-effort
host1(config-traffic-class)#exit
```

2. Create a parameter definition.

- a. Configure the QoS parameter name and the application.

```
host1(config)#qos-parameter-define ipm application ip-multicast hierarchical
```

- b. Configure a controlled-interface type.

```
host1(config-qos-parameter-define)#controlled-interface-type vlan
host1(config-qos-parameter-define)#exit
```

3. Create a parameter instance that globally defines the value of the IP multicast adjustment as 0.

```
host1(config)#qos-parameter ipm 0
```

4. Reference the parameter within a scheduler profile parameter expression.

```
host1(config)#scheduler-profile vlan-subscriber
host1(config-scheduler-profile)#shared-shaping-rate 1000000 - ipm burst 50
milliseconds auto
host1(config-scheduler-profile)#exit
```

5. Add the scheduler profile to a QoS profile.

```
host1(config)#qos-profile vlan-subscriber
host1(config-qos-profile)#vlan queue traffic-class best-effort
host1(config-qos-profile)#vlan queue traffic-class voice scheduler-profile 192k
host1(config-qos-profile)#vlan node scheduler-profile vlan-subscriber
host1(config-qos-profile)#exit
```

6. Attach the parameter definition to a logical interface.

```
host1(config)#interface gigabitEthernet 7/0
host1(config-if)#encapsulation vlan
host1(config-if)#exit
host1(config)#interface gigabitEthernet 7/0.1
host1(config-if)#vlan id 200
host1(config-if)#qos-profile vlan-subscriber
host1(config-if)#ip address 1.1.1.1 255.255.255.0
```

After the QoS profile is attached to the interface, the IP multicast bandwidth adjustment application begins to adjust rates based on IGMP joins and leaves received on that interface.

- Related Documentation**
- [IP Multicast Bandwidth Adjustment for QoS Overview on page 31](#)
 - [Example: QoS Parameter Configuration for IP Multicast Bandwidth Adjustment on page 104](#)
 - [controlled-interface-type on page 129](#)
 - encapsulation vlan
 - interface gigabitEthernet
 - [node on page 139](#)
 - [qos-parameter-define on page 142](#)
 - [qos-profile on page 145](#)
 - [queue on page 147](#)
 - [scheduler-profile on page 149](#)
 - [shared-shaping-rate on page 153](#)
 - [traffic-class on page 155](#)
 - vlan id

Configuring a Parameter Definition to Shape Ethernet Traffic Using Cell Mode

To associate a parameter instance with the QoS cell mode application:

1. Configure traffic classes.

```
host1(config)#traffic-class voice
host1(config-traffic-class)#exit
host1(config)#traffic-class best-effort
host1(config-traffic-class)#exit
```

2. Create a parameter definition.

- a. Configure the QoS parameter name and the application.

```
host1(config)#qos-parameter-define shaping-mode application qos-cell-mode
```

- b. Configure a controlled-interface type.

```
host1(config-qos-parameter-define)#controlled-interface-type vlan
host1(config-qos-parameter-define)#controlled-interface-type ip
```

- c. Configure an instance-interface type.

```
host1(config-qos-parameter-define)#instance-interface-type vlan
```

3. Create the parameter instance and configure the shaping mode.

When you create the parameter instance and configure the shaping mode, the value of frame shaping mode is 0; the value for cell shaping mode is 1.

```
host1(config)#interface gigabitEthernet 6/0/2
host1(config-if)#encapsulation vlan
host1(config-if)#interface gigabitEthernet 6/0/2.1
```

```

host1(config-if)#vlan id 1
host1(config-if)#qos-parameter cell-mode 1
host1(config-if)#ip address 6.10.10.10 255.255.255.255
host1(config-if)#exit
host1(config)#interface gigabitEthernet 6/0/2
host1(config-if)#svlan 1 qos-parameter cell-mode 1
host1(config-if)#exit
host1(config)#interface gigabitEthernet 6/0/2
host1(config-if)#qos-parameter cell-mode 1

```

Related Documentation

- Example: QoS Parameter Configuration for QoS Cell Mode and Byte Adjustment for Cell Shaping
- Scheduler Profiles and Parameter Expressions for QoS Administrators
- QoS Downstream Rate Application Overview
- [controlled-interface-type on page 129](#)
- [instance-interface-type on page 136](#)
- interface gigabitEthernet
- ip address
- [qos-parameter on page 140](#)
- [qos-parameter-define on page 142](#)
- [scheduler-profile on page 149](#)
- svlan qos-parameter
- [traffic-class on page 155](#)
- vlan id

Configuring a Parameter Definition for QoS Downstream Rate

To associate a parameter instance with the QoS downstream rate application:

1. Configure traffic classes.

```

host1(config)#traffic-class voice
host1(config-traffic-class)#exit
host1(config)#traffic-class best-effort
host1(config-traffic-class)#exit

```

2. Create a parameter definition for the QoS downstream rate application.

- a. Configure the QoS parameter name and the application.

```

host1(config)#qos-parameter-define downstreamVLAN application
qos-downstream-rate

```

- b. Configure controlled-interface types.

```

host1(config-qos-parameter-define)#controlled-interface-type vlan
host1(config-qos-parameter-define)#controlled-interface-type ip

```


- c. Configure subscriber-interface types.

```
host1(config-qos-parameter-define)#subscriber-interface-type vlan
```

3. Do one of the following:

- For VLANs, configure the shaping mode by creating a parameter definition with the QoS cell mode application. Ensure that you specify a subscriber-interface type.

See [“Configuring a Parameter Definition to Shape Ethernet Traffic Using Cell Mode” on page 79](#).

- For ATM VCs, configure the shaping mode by issuing the **qos-shaping-mode** command.

See [Configuring the QoS Shaping Mode for ATM Interfaces](#).

4. Enable QoS adaptive mode for the system by issuing the **qos-adaptive-mode** command in L2C Configuration mode.

```
host1(config)#l2c
host1(config-l2c)#qos-adaptive-mode
```

5. Enable the QoS downstream rate application to use downstream rates obtained from the Actual-Data-Rate-Downstream [26-130] DSL Forum VSA.

```
host1(config)#aaa qos downstream-rate
```

6. Configure the scheduler profile for the shaping rate.

```
host1(config)#scheduler-profile vlan1
host1(config-scheduler-profile)#shared-shaping-rate downstreamVLAN * 5 auto
```

7. Configure the QoS profile for the shaping rate.

```
host1(config)#qos-profile vlan1
host1(config-qos-profile)#vlan node scheduler-profile vlan1
```

8. Attach the QoS profile to a logical Ethernet interface.

ANCP or AAA dynamically creates the parameter instances for the QoS downstream rate application, and if applicable, the QoS cell mode application; therefore, you do not need to specify them.

```
host1(config)#interface gigabitEthernet 6/0/2
host1(config-if)#encapsulation vlan
host1(config-if)#interface gigabitEthernet 6/0/2.1
host1(config-if)#vlan id 1
host1(config-if)#qos-profile vlan1
host1(config-if)#ip address 6.10.10.10 255.255.255.255
```

For more information about configuring ANCP (L2C) parameters, see *JunosE IP Services Configuration Guide*.

Related Documentation

- [Example: QoS Parameter Configuration for QoS Downstream Rate on page 110](#)
- DSL Forum VSAs
- `aaa qos downstream-rate`
- [controlled-interface-type on page 129](#)

- encapsulation vlan
- [instance-interface-type on page 136](#)
- ip address
- [node on page 139](#)
- [qos-parameter on page 140](#)
- qos-adaptive-mode
- [qos-parameter-define on page 142](#)
- [qos-profile on page 145](#)
- [queue on page 147](#)
- [shared-shaping-rate on page 153](#)
- subscriber-interface-type
- [traffic-class on page 155](#)
- vlan id

Creating Parameter Instances

You can create QoS parameter instances globally, for an interface, or for a subinterface.

Tasks to create parameter instances are:

- [Creating a Global Parameter Instance on page 82](#)
- [Creating a Parameter Instance for an Interface on page 82](#)
- [Creating a Parameter Instance for an ATM VP on page 83](#)
- [Creating a Parameter Instance for an S-VLAN on page 83](#)

Creating a Global Parameter Instance

To create a global parameter instance:

- Create a parameter instance in Global Configuration mode.
`host1(config)#qos-parameter max-subscriber-bandwidth 6000000`

Creating a Parameter Instance for an Interface

To create a parameter instance for an interface:

1. Specify an interface.
`host1(config)#interface atm 11/0.1`
`host1(config)#interface gigabitEthernet 2/0`
2. Specify the parameter name and the value.
`host1(config-subif)#qos-parameter max-subscriber-bandwidth 6000000`

Creating a Parameter Instance for an ATM VP

Use this procedure to attach a parameter instance to a VP on the interface. Optionally, use the **qos-profile** keyword to attach a parameter instance to a QoS profile.

To create a parameter instance for an ATM VP:

1. Configure the ATM VP.

```
host1(config)#interface atm 2/0
host1(config-if)#atm vp-tunnel 4
```

2. Do either of the following:

- Attach the parameter instance to an ATM VP on the interface.

```
host1(config-if)#atm-vp 4 qos-parameter max-subscriber-bandwidth 375000
```

- Attach the parameter instance and associate with the QoS profile.

```
host1(config-if)#atm-vp 4 qos-profile video qos-parameter
max-subscriber-bandwidth 375000
```

Creating a Parameter Instance for an S-VLAN

Use this procedure to attach a parameter instance to a specified S-VLAN ID on the interface. Optionally, use the **qos-profile** keyword to attach a parameter instance to a QoS profile.

To create a parameter instance for an S-VLAN:

1. Specify the Ethernet interface and create the VLAN.

```
host1(config)#interface gigabitEthernet 3/0
host1(config-if)#encapsulation vlan
host1(config-if)#interface gigabitEthernet 3/0.1
```

2. Specify the S-VLAN ID.

```
host1(config-if)#svlan id 1 202
```

3. Attach the parameter instance to an S-VLAN ID on the interface.

```
host1(config-if)#svlan 202 qos-parameter max-subscriber-bandwidth 6000000
```

Related Documentation

- [Parameter Instances for QoS Clients Overview](#)
- [JunosE Broadband Access Configuration Guide](#)
- [atm-vp qos-parameter](#)
- [atm vp-tunnel](#)
- [encapsulation vlan](#)
- [interface](#)
- [qos-parameter on page 140](#)
- [svlan id](#)

- `svlan qos-parameter`

Configuration Tasks for Statistics Collection

- [Configuring Statistic Profiles for QoS on page 85](#)
- [Configuring Rate Statistics on page 85](#)
- [Configuring Event Statistics on page 86](#)

Configuring Statistic Profiles for QoS

To begin to configure a statistics profile, enter Statistics Profile Configuration mode.

- Issue the **statistics-profile** command from Global Configuration mode:

```
host1(config)#statistics-profile statpro-1
host1(config-statistics-profile)#
```

The router supports up to 250 statistics profiles.

Related Documentation

- [Configuring Rate Statistics on page 85](#)
- [Configuring Event Statistics on page 86](#)
- [Monitoring QoS Statistics for Rates and Events](#)
- [statistics-profile on page 197](#)

Configuring Rate Statistics

To gather rate statistics:

1. Create the statistics profile.

```
host1(config)#statistics-profile statpro-5
```

2. Set the length of time during which statistics are counted.

```
host1(config-statistics-profile)#rate-period 45
```

Rate period range is 1–43200 seconds.

3. Reference the statistics profile by a QoS profile.

```
host1(config)#qos-profile qospro-3
host1(config-qos-profile)#ip queue traffic-class tc1 scheduler-profile sp1
statistics-profile statpro-5
```

4. Attach the QoS profile to the appropriate interface.

```
host1(config)#interface gigabitEthernet 1/0
host1(config-subif)#qos-profile qospro-3
host1(config-subif)#exit
```

5. (Optional) Display the rate statistics.

```
host1#show egress-queue rates interface gigabitEthernet 1/0
```

Related Documentation

- [Configuring Statistic Profiles for QoS on page 85](#)
- [Configuring a QoS Profile on page 61](#)
- [Monitoring QoS Statistics for Rates and Events](#)
- [interface](#)
- [qos-profile on page 145](#)
- [queue on page 147](#)
- [rate-period on page 196](#)
- [statistics-profile on page 197](#)

Configuring Event Statistics

To configure the router to count events on a queue, you configure the threshold above which forwarding or drop events are counted.

A forwarding rate event occurs each time the forwarding rate exceeds the threshold during the specified rate period.

A drop event occurs each time the number of packets dropped exceeds the threshold during the specified rate period.

To gather event statistics:

1. Create the statistics profile.

```
host1(config)#statistics-profile statpro-1
```

2. Set the length of time during which statistics are counted.

```
host1(config-statistics-profile)#rate-period 30
```

Rate period range is 1–43200 seconds.

3. (Optional) Set the threshold above which forwarding rate events are counted.

```
host1(config-statistics-profile)#forwarding-rate-threshold 10000000
```

Forwarding rate threshold range is 0–1073741824 bps; default is no threshold.

4. (Optional) Set a threshold for committed (green) packets.

```
host1(config-statistics-profile)#committed-drop-threshold 2000000
```

Drop rate threshold range is 0–1073741824 bps; default is no threshold.

5. (Optional) Set a threshold for conformed (yellow) packets.

```
host1(config-statistics-profile)#conformed-drop-threshold 4000000
```

Drop rate threshold range is 0–1073741824 bps; default is no threshold.

6. (Optional) Set a threshold for exceeded (red) packets.

```
host1(config-statistics-profile)#exceeded-drop-threshold 6000000
```

Drop rate threshold range is 0–1073741824 bps; default is no threshold.

7. Reference the statistics profile in a QoS profile.

```
host1(config)#qos-profile qospro-1
host1(config-qos-profile)#ip queue traffic-class tc1 scheduler-profile sp1
statistics-profile statpro-1
```

8. Attach the QoS profile to the appropriate interface.

```
host1(config)#interface gigabitEthernet 1/0
host1(config-subif)#qos-profile qospro-1
host1(config-subif)#exit
```

9. (Optional) Display the event statistics.

```
host1#show egress-queue events interface gigabitEthernet 1/0
```

Related Documentation

- [Configuring Statistic Profiles for QoS on page 85](#)
- [Configuring a QoS Profile on page 61](#)
- [Monitoring QoS Statistics for Rates and Events](#)
- [committed-drop-threshold on page 192](#)
- [conformed-drop-threshold on page 193](#)
- [exceeded-drop-threshold on page 194](#)
- [forwarding-rate-threshold on page 195](#)
- [qos-profile on page 145](#)
- [queue on page 147](#)
- [rate-period on page 196](#)
- [statistics-profile on page 197](#)

Examples


- [Example: QoS Parameter Configuration for Controlling Subscriber Bandwidth on page 89](#)
- [Example: QoS Parameter Configuration for IP Multicast Bandwidth Adjustment on page 104](#)
- [Example: QoS Parameter Configuration for QoS Downstream Rate on page 110](#)
- [Example: QoS Parameter Configuration for Hierarchical Parameters on page 115](#)

Example: QoS Parameter Configuration for Controlling Subscriber Bandwidth

The example in this section illustrates how to use parameters to control the minimum and maximum bandwidth of a subscriber. The example includes procedures for both QoS administrators and QoS clients.

Through QoS parameter definitions, the QoS administrator defines a QoS scheduler hierarchy that corresponds to the physical network topology shown in [Figure 15 on page 89](#).

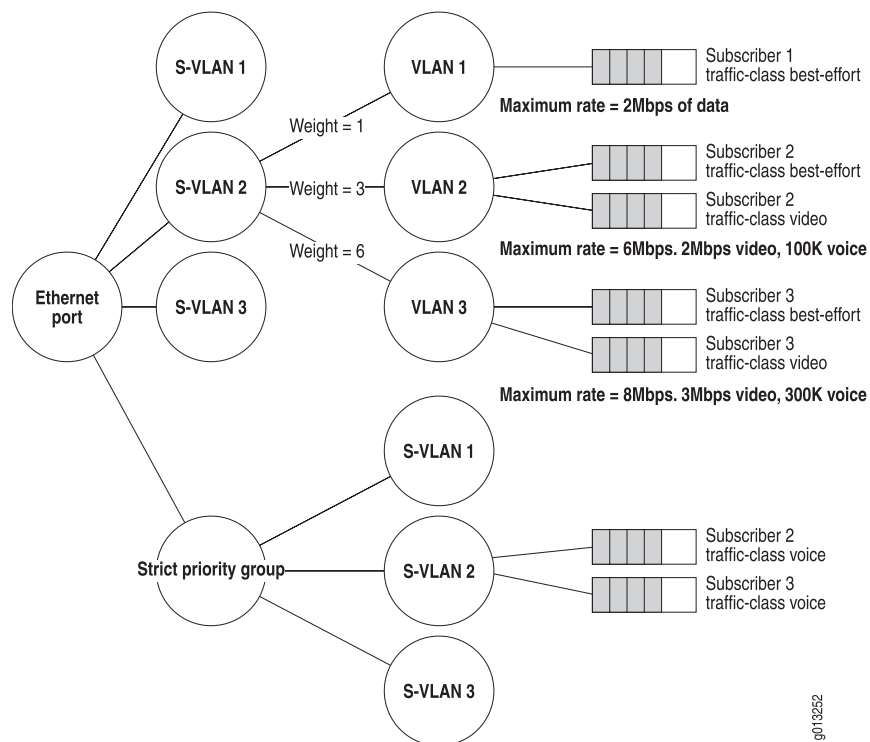
Figure 15: Physical Network Topology



The S-VLAN scheduler nodes correspond to the DSLAM in the physical network topology; the VLAN scheduler nodes correspond to the subscribers.

[Figure 16 on page 90](#) shows the QoS scheduler hierarchy that the QoS client creates when configuring a different service for each subscriber.

Figure 16: QoS Scheduler Hierarchy



For Subscriber 1, the QoS client configures a basic best-effort data service, with a maximum rate of 2 Mbps, and assigns a scheduler weight value of 1.

For Subscriber 2, the QoS client configures a basic triple-play service consisting of voice, video, and best-effort data services. This service enables the subscriber to transmit up to 6 Mbps of combined voice, video, and best-effort data traffic. The service limits video traffic to 2 Mbps and enables low-latency bandwidth for one 100 Kbps voice call. The QoS client then assigns this subscriber a scheduler weight value of 3, enabling this subscriber to claim up to three times the bandwidth than the basic data service configured for Subscriber 1.

For Subscriber 3, the QoS client configures an enhanced triple-play service consisting of voice, video and best-effort data services. This enhanced triple-play service enables the subscriber to transmit up to 8 Mbps of combined voice, video, and best-effort data traffic. This service limits video traffic to 3 Mbps and enables low-latency bandwidth for up to three 100 Kbps voice calls. The QoS client then assigns this subscriber a scheduler weight value of 6, enabling this subscriber to claim up to six times the bandwidth of the basic data service subscriber configured for Subscriber 1, and up to twice the bandwidth of the basic triple-play subscriber configured for Subscriber 2.

Procedure for QoS Administrators

This section describes the procedures to configure the scheduler hierarchy shown in Figure 16 on page 90 by using QoS parameters.

Configuring Traffic Classes and Traffic Class Groups

The QoS administrator configures traffic classes and traffic-class groups for best-effort data, video, and voice services.

1. Configure the traffic classes.
 - a. Configure the traffic class named best-effort.
 - b. Configure the traffic class named video.
 - c. Configure the traffic class named voice.
 - d. Enable the voice traffic class to provide a strict priority treatment throughout the fabric.

```
host1(config)#traffic-class best-effort
host1(config-traffic-class)#exit
```

```
host1(config)#traffic-class video
host1(config-traffic-class)#exit
```

```
host1(config)#traffic-class voice
host1(config-traffic-class)#fabric-strict-priority
host1(config-traffic-class)#exit
```

2. Configure a traffic-class group for low-latency expedited forwarding (EF) and add the voice traffic class into the traffic-class group EF.
 - a. Configure the EF traffic-class group with strict-priority scheduling.
 - b. Add the voice traffic class to the traffic-class group.

```
host1(config)#traffic-class-group EF auto-strict-priority
host1(config-traffic-class-group)#traffic-class voice
host1(config-traffic-class-group)#exit
```

The remaining traffic classes, best-effort and video, remain in the default traffic-class group.

Configuring the Parameter Definitions

After configuring the traffic classes and traffic-class groups, the QoS administrator configures the parameter definitions for Subscribers 1, 2, and 3.

1. Configure a parameter definition for the maximum subscriber bandwidth.
 - a. Configure the parameter definition named max-subscriber-bandwidth.
 - b. Enable the parameter to control VLANs.
 - c. Enable the parameter to have instances created on VLAN subinterfaces.
 - d. Specify the valid range of this parameter as 512 Kbps–8 Mbps.

```
host1(config)#qos-parameter-define max-subscriber-bandwidth
host1(config-qos-parameter-define)#controlled-interface-type vlan
host1(config-qos-parameter-define)#instance-interface-type vlan
host1(config-qos-parameter-define)#range 512000 8192000
```

```
host1(config-qos-parameter-define)#exit
```

2. Configure a parameter definition for a subscriber's weight in the hierarchical round-robin (HRR) scheduler. This parameter is used to provide different scheduler weights for each of the three service offerings.
 - a. Configure the parameter definition named subscriber-weight.
 - b. Enable the parameter to control VLANs.
 - c. Enable the parameter to have instances created on VLAN subinterfaces.
 - d. Specify the valid range of this parameter as 1–6.

```
host1(config)#qos-parameter-define subscriber-weight
host1(config-qos-parameter-define)#controlled-interface-type vlan
host1(config-qos-parameter-define)#instance-interface-type vlan
host1(config-qos-parameter-define)#range 1 6
host1(config-qos-parameter-define)#exit
```

3. Configure a parameter definition for the subscriber's maximum video bandwidth. By creating a parameter instance on S-VLANs, the QoS administrator can specify a subscriber's maximum video bandwidth for each DSLAM in the hierarchy.
 - a. Configure the parameter definition named max-subscriber-video-bandwidth.
 - b. Enable the parameter to control VLANs.
 - c. Enable the parameter to have instances created on both SVLAN and VLAN subinterfaces.
 - d. Specify the valid range of this parameter as 1 Mbps–5 Mbps.

```
host1(config)#qos-parameter-define max-subscriber-video-bandwidth
host1(config-qos-parameter-define)#controlled-interface-type vlan
host1(config-qos-parameter-define)#instance-interface-type vlan
host1(config-qos-parameter-define)#instance-interface-type svlan
host1(config-qos-parameter-define)#range 1000000 5000000
host1(config-qos-parameter-define)#exit
```

4. Configure a parameter definition for the maximum number of 100 Kbps voice calls supported for the subscriber.
 - a. Configure the parameter definition named max-100Kbps-voice-calls.
 - b. Enable the parameter to control VLANs.
 - c. Enable the parameter to have instances created on VLAN subinterfaces.
 - d. Specify the valid range of this parameter as 1–3.

```
host1(config)#qos-parameter-define max-100Kbps-voice-calls
host1(config-qos-parameter-define)#controlled-interface-type vlan
host1(config-qos-parameter-define)#instance-interface-type vlan
host1(config-qos-parameter-define)#range 1 3
host1(config-qos-parameter-define)#exit
```

Configuring the Scheduler Profiles

The QoS administrator can then reference the parameter definitions within a scheduler profile, which defines the shaping rates for the parameter.

1. Configure a scheduler profile to specify the maximum bandwidth of the subscriber's best-effort data.
 - a. Configure the scheduler profile named subscriber-best-effort.
 - b. Configure the shared-shaping rate by referencing the max-subscriber-bandwidth parameter and choosing automatic shared shaping.

```
host1(config)#scheduler-profile subscriber-best-effort
host1(config-scheduler-profile)#shared-shaping-rate max-subscriber-bandwidth
auto
host1(config-scheduler-profile)#exit
```

2. Configure a scheduler profile to specify the maximum bandwidth of the subscriber's video service.
 - a. Configure the scheduler profile named subscriber-video.
 - b. Configure the shaping rate by referencing the max-subscriber-video-bandwidth parameter.

```
host1(config)#scheduler-profile subscriber-video
host1(config-scheduler-profile)#shaping-rate max-subscriber-video-bandwidth
host1(config-scheduler-profile)#exit
```

3. Configure a scheduler profile for the subscriber's weight within the HRR scheduler.
 - a. Configure the scheduler profile named subscriber-weight.
 - b. Configure the weight using the default for the subscriber-weight parameter.

```
host1(config)#scheduler-profile subscriber-weight
host1(config-scheduler-profile)#weight subscriber-weight
host1(config-scheduler-profile)#exit
```

4. Configure a scheduler profile for the subscriber's voice service.
 - a. Configure the scheduler profile named subscriber-voice.
 - b. Configure the shaping rate by referencing the max-100Kbps-voice-calls parameter and multiplying it by 100 Kbps of voice calls.

```
host1(config)#scheduler-profile subscriber-voice
host1(config-scheduler-profile)#shaping-rate max-100Kbps-voice-calls * 100000
host1(config-scheduler-profile)#exit
```

Configuring the QoS Profiles

By referencing the scheduler profiles within QoS profiles, the QoS administrator creates the scheduler hierarchy. In this portion of the example, the QoS administrator configures QoS profiles for the best-effort data and triple-play service offerings.

1. Define a QoS profile for the best-effort data service.
 - a. Create the QoS profile named subscriber-data-service.
 - b. Create a node for S-VLAN subinterfaces.

- c. Specify a node for VLAN subinterfaces and reference the subscriber-weight scheduler profile.
- d. Specify a queue for VLAN subinterfaces, referencing the best-effort traffic class and the subscriber-best-effort scheduler-profile.

```
host1(config)#qos-profile subscriber-data-service
host1(config-qos-profile)#svlan node
host1(config-qos-profile)#vlan node scheduler-profile subscriber-weight
host1(config-qos-profile)#vlan queue traffic-class best-effort scheduler-profile
subscriber-best-effort
host1(config-qos-profile)#exit
```

The best-effort queue rule for VLAN subinterfaces refers to the subscriber-best-effort scheduler profile. The scheduler profile refers to the max-subscriber-bandwidth parameter that controls the maximum rate of this subscriber's best-effort queue.

- 2. Define a QoS profile for the triple-play service and specify S-VLAN nodes and VLAN nodes.
 - a. Create a QoS profile named subscriber-triple-play.
 - b. Specify a node for S-VLAN subinterfaces.
 - c. Specify a node for VLAN subinterfaces and reference the subscriber-weight scheduler profile.
 - d. Specify a node for S-VLAN subinterfaces and reference the EF traffic-class group.
 - e. Specify a queue for VLAN subinterfaces, referencing the best-effort traffic class and the subscriber-best-effort scheduler profile.
 - f. Specify a queue for VLAN subinterfaces, referencing the video traffic class and the subscriber-video scheduler profile.
 - g. Specify a queue for VLAN subinterfaces, referencing the voice traffic-class and the subscriber-voice scheduler profile.

```
host1(config)#qos-profile subscriber-triple-play
host1(config-qos-profile)#svlan node
host1(config-qos-profile)#vlan node scheduler-profile subscriber-weight
host1(config-qos-profile)#svlan node group EF
host1(config-qos-profile)#vlan queue traffic-class best-effort scheduler-profile
subscriber-best-effort
host1(config-qos-profile)#vlan queue traffic-class video scheduler-profile
subscriber-video
host1(config-qos-profile)#vlan queue traffic-class voice scheduler-profile
subscriber-voice
host1(config-qos-profile)#exit
```

VLAN queues are used for each service. The VLAN queue rules reference scheduler profiles that define the scheduler rates for the service.

- 3. Configure a QoS profile and attach to all Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces in the chassis.

- a. Create a QoS profile named ethernet-default.
- b. Remove the QoS profile rule for creating IP nodes.
- c. Remove the IP queue for the best-effort traffic-class.

```
host1(config)#qos-profile ethernet-default
host1(config-qos-profile)#no ip node
host1(config-qos-profile)#no ip queue traffic-class best-effort
host1(config-qos-profile)#exit
```

4. Configure the Fast Ethernet interface and VLAN subinterfaces.
 - a. Configure the Fast Ethernet interface in slot 9, port 0.
 - b. Configure the VLAN major interface.
 - c. Configure the VLAN subinterface at slot 9, port 0, subinterface 1.
 - d. Assign an S-VLAN ID of 2 and a VLAN ID of 1 to the VLAN subinterface.
 - e. Assign an IP address to the VLAN subinterface.
 - f. Repeat Steps a–e to configure VLAN subinterfaces in slot 9, port 0, subinterface 2 and in slot 9, port 0, subinterface 3.

```
host1(config)# interface fastEthernet 9/0
host1(config-if)#encapsulation vlan
host1(config-if)#exit
host1(config)#interface fastEthernet 9/0.1
host1(config-subif)#svlan id 2 1
host1(config-subif)#ip address 192.1.1.1 255.255.255.0
host1(config)#interface fastEthernet 9/0.2
host1(config-subif)#svlan id 2 2
host1(config-subif)#ip address 192.2.1.1 255.255.255.0
host1(config-subif)#exit
host1(config)#interface fastEthernet 9/0.3
host1(config-subif)#svlan id 2 3
host1(config-subif)#ip address 192.3.1.1 255.255.255.0
host1(config-subif)#exit
```

Procedure for QoS Clients

This section describes procedures to create parameter instances for Subscribers 1, 2, and 3.

Creating a Global Parameter Instance

The QoS client creates global parameter instances to provide a minimal level of default service for the router. In this portion of the example, the QoS client configures 2 Mbps of data traffic and configures a scheduler weight of 1 for Subscriber 1. For Subscribers 2 and 3, the QoS client then configures a maximum of 2 Mbps of video bandwidth and 1 voice call.

To create a global parameter instance:

1. Create a global parameter instance for max-subscriber-bandwidth with a value of 2000000.
2. Create a global parameter instance for subscriber-weight with a value of 1.
3. Create a global parameter instance for subscriber-video-bandwidth with a value of 2000000.
4. Create a global parameter instance for max-100Kbps-voice-calls with a value of 1.

```
host1(config)#qos-parameter max-subscriber-bandwidth 2000000
host1(config)#qos-parameter subscriber-weight 1
host1(config)#qos-parameter max-subscriber-video-bandwidth 2000000
host1(config)#qos-parameter max-100Kbps-voice-calls 1
```

Creating a Global Parameter Instance for Individual DSLAMs

Instead of creating global parameter instances, the QoS client can create different parameter instances for the DSLAMs that correspond to the S-VLAN nodes shown in [Figure 16 on page 90](#). In this portion of the example, the QoS client creates 1 Mbps video streams by default on DSLAM 1, rather than the 2Mbps global parameter instance.

1. Specify the Fast Ethernet interface in slot 9, port 0.
2. Attach the QoS parameter max-subscriber-video-bandwidth to S-VLAN 1.

```
host1(config)#interface fastEthernet 9/0
host1(config-if)#svlan 1 qos-parameter max-subscriber-video-bandwidth 1000000
host1(config-if)#exit
```

Creating Parameter Instances for Subscribers

The QoS client creates a parameter instance for Subscribers 1, 2, and 3.

1. Configure the basic-data service for Subscriber 1.
 - a. Specify the Fast Ethernet interface in slot 9, port 0.
 - b. Attach the QoS profile subscriber-data-service to the subscriber's Fast Ethernet interface.

```
host1(config)#interface fastEthernet 9/0.1
host1(config-subif)#qos-profile subscriber-data-service
host1(config-subif)#exit
```

This QoS profile references the scheduler profiles, which then reference the parameter instances max-subscriber-bandwidth and subscriber-weight. These global parameter instances are created with values 2 Mbps and 1.

2. Configure a basic triple-play service consisting of voice, video, and data services for Subscriber 2.
 - a. Specify the Fast Ethernet interface in slot 9, port 0.
 - b. Create a parameter instance for max-subscriber-bandwidth, enabling the subscriber to transmit up to 6 Mbps of combined voice, video, and data traffic.
 - c. Create a parameter instance for subscriber-weight with a value of 3. This value enables the subscriber to claim up to three times the bandwidth of Subscriber 1, with basic data service.

- d. Create a parameter instance for max-subscriber-video-bandwidth, limiting video traffic to 2 Mbps.
- e. Create a parameter instance for max-100Kbps-voice-calls, enabling bandwidth for one 100 Kbps voice call.
- f. Attach the QoS profile subscriber-triple-play to the subscriber's interface.

```
host1(config)#interface fastEthernet 9/0.2
host1(config-if)#qos-parameter max-subscriber-bandwidth 6000000
host1(config-if)#qos-parameter subscriber-weight 3
host1(config-if)#qos-parameter max-subscriber-video-bandwidth 2000000
host1(config-if)#qos-parameter max-100Kbps-voice-calls 1
host1(config-if)#qos-profile subscriber-triple-play
host1(config-if)#exit
```

3. Configure an enhanced triple-play service consisting of voice, video, and data services for Subscriber 3. Enable the subscriber to have twice as much bandwidth as Subscriber 2, with basic triple-play service.
 - a. Create a parameter instance for max-subscriber-bandwidth, enabling the subscriber to transmit up to 8 Mbps of combined voice, video, and data traffic.
 - b. Create a parameter instance for subscriber-weight with a value of 6, enabling the subscriber to claim up to six times the bandwidth of Subscriber 1, with basic data service.
 - c. Create a parameter instance for max-subscriber-video-bandwidth, limiting video traffic to 3 Mbps.
 - d. Create a parameter instance for max-100Kbps-voice-calls, enabling up to three 100 Kbps voice calls.
 - e. Attach the QoS profile subscriber-triple-play to the subscriber's interface.

```
host1(config)#interface fastEthernet 9/0.3
host1(config-if)#qos-parameter max-subscriber-bandwidth 8000000
host1(config-if)#qos-parameter subscriber-weight 6
host1(config-if)#qos-parameter max-subscriber-video-bandwidth 3000000
host1(config-if)#qos-parameter max-100Kbps-voice-calls 3
host1(config-if)#qos-profile subscriber-triple-play
host1(config-if)#exit
```

Monitoring the Subscriber Configuration

After completing the configuration, both the QoS administrator and the QoS client can monitor it by issuing **show** commands.

1. To display the traffic classes for best-effort, video, and voice, issue the **show traffic-class** command.

```
host1#show traffic-class
```

		fabric
traffic	fabric	strict
class	weight	priority

```

-----
best-effort      8      no
video           8      no
voice           8      yes

```

- To display the traffic-class group EF, issue the **show traffic-class-group** command.

```

host1#show traffic-class-group

traffic-class-group EF auto-strict-priority
traffic-class voice

```

- To display the settings for all four QoS parameter definitions (max-subscriber-bandwidth, subscriber-weight, max-subscriber-video-bandwidth, and max-100Kbps-voice-calls), issue the **show qos-parameter-define** command.

```

host1#show qos-parameter-define

```

parameter name	controlled interface types	instance interface types	subscriber interface types
max-subscriber-bandwidth	vlan	vlan	<none>
subscriber-weight	vlan	vlan	<none>
max-subscriber-video-bandwidth	vlan	vlan, svlan	<none>
max-100Kbps-voice-calls	vlan	vlan	<none>

parameter name	value range	properties
max-subscriber-bandwidth	512000 - 8192000	<none>
subscriber-weight	1 - 10	<none>
max-subscriber-video-bandwidth	1000000 - 5000000	<none>
max-100Kbps-voice-calls	1 - 3	<none>

- To display the shaping rates and burst for the four scheduler profiles (subscriber-best-effort, subscriber-video, subscriber-weight, and subscriber-voice), issue the **show scheduler-profile** command.

```

host1#show scheduler-profile

```

scheduler	shaping rate	shaping burst
default	<none>	<none>
subscriber-best-effort	<none>	<none>
subscriber-video	max-subscriber-video-bandwidth	default
subscriber-weight	<none>	<none>
subscriber-voice	max-100Kbps-voice-calls * 100000	default

scheduler	weight	strict priority	assured rate
default	8	no	<none>
subscriber-best-effort	8	no	<none>
subscriber-video	8	no	<none>
subscriber-weight	subscriber-weight	no	<none>
subscriber-voice	8	no	<none>

scheduler	shared shaping rate	shared shaping burst	shared shaping constituent
default	<none>	<none>	<none>
subscriber-best-effort	max-subscriber-bandwidth	default	<none>

subscriber-video	<none>	<none>	<none>
subscriber-weight	<none>	<none>	<none>
subscriber-voice	<none>	<none>	<none>

scheduler	shared shaping mode
default	<none>
subscriber-best-effort	auto implicit
subscriber-video	<none>
subscriber-weight	<none>
subscriber-voice	<none>

5. To display the settings for the QoS profile subscriber-triple-play, issue the **show qos-profile** command.

host1#show qos-profile subscriber-triple-play

```
qos-profile subscriber-triple-play:
t-class interface rule traffic
group type type class scheduler profile queue drop statistics
profile profile profile
-----
```

	vlan	node		subscriber-weight			
	svlan	node		default			
	vlan	queue	best-effort	subscriber-best-effort	default	default	default
	vlan	queue	video	subscriber-video	default	default	default
EF	svlan	node		default			
EF	vlan	queue	voice	subscriber-voice	default	default	default

6. To display the attachments on all QoS profiles, issue the **show qos-profile references** command.

```
host1#show qos-profile references
qos profile attachment
-----
```

atm-default	(qos-port-type-profile)
serial-default	(qos-port-type-profile)
ethernet-default	(qos-port-type-profile)
server-default	(qos-port-type-profile)
subscriber-data-service	vlan FastEthernet9/0.1
subscriber-triple-play	vlan FastEthernet9/0.2
subscriber-triple-play	vlan FastEthernet9/0.3

Port attachments: 4
Interface attachments: 3
Not attached: 0

7. To display global and interface attachments on all of the QoS parameter instances, issue the **show qos-parameter references** command.

```
host1#show qos-parameter references
```

interface	parameter name	value
global	max-subscriber-bandwidth	2000000
global	subscriber-weight	1
global	max-subscriber-video-bandwidth	2000000
global	max-100Kbps-voice-calls	1
FastEthernet9/0.2	max-subscriber-bandwidth	6000000
FastEthernet9/0.2	subscriber-weight	3

```

FastEthernet9/0.2      max-subscriber-video-bandwidth 2000000
FastEthernet9/0.2      max-100Kbps-voice-calls      1
FastEthernet9/0.3      max-subscriber-bandwidth      8000000
FastEthernet9/0.3      subscriber-weight      6
FastEthernet9/0.3      max-subscriber-video-bandwidth 3000000
FastEthernet9/0.3      max-100Kbps-voice-calls      3
FastEthernet9/0  svlan 1 max-subscriber-video-bandwidth 1000000

```

```

Global parameter instances: 4
Parameter instances reported: 13

```

8. To display the queue forwarding rates for the VLANs on the Fast Ethernet interface in slot 9, port 0, issue the **show egress-queue rates** command.

```
host1#show egress-queue rates full interface fastEthernet 9/0
```

	traffic	forwarded	aggregate	minimum	maximum
interface	class	rate	drop rate	rate	rate
ethernet FastEthernet9/0	best-effort	*	*	0	100000000
vlan FastEthernet9/0.1	best-effort	*	*	0	2000000
vlan FastEthernet9/0.2	best-effort	*	*	0	6000000
	video	*	*	0	2000000
	voice	*	*	100000	100000
vlan FastEthernet9/0.3	best-effort	*	*	0	8000000
	video	*	*	0	3000000
	voice	*	*	300000	300000

```

Queues reported: 0
Queues filtered (under threshold): 0
* Queues disabled (no rate period): 8
**Queues disabled (no resources): 0
Total queues: 8

```

9. To display the shared-shaper settings for the VLANs on the Fast Ethernet interface in slot 9, port 0, issue the **show qos shared-shaper** command.

```
host1#show qos shared-shaper interface fastEthernet 9/0
```

interface	resource	shared shaping rate	shaping rate	other rate
vlan Eth9/0.1	vlan node			
	A vlan queue best-effort	2000000		2000000
vlan Eth9/0.2	vlan node			
	A vlan queue best-effort	6000000		6000000
	A vlan queue video		2000000	
	A vlan queue EF voice		100000	
vlan Eth9/0.3	vlan node			
	A vlan queue best-effort	8000000		8000000
	A vlan queue video		3000000	
	A vlan queue EF voice		300000	
Total shared shapers:		3		
Total constituents:		10		

Total shared shaper failovers: 0
Compound shared shapers are not supported.

10. To display the scheduler hierarchy for the Fast Ethernet interface in slot 9, port 0, issue the **show qos scheduler-hierarchy** command.

```
host1# show qos scheduler-hierarchy interface fastEthernet 9/0
Scheduler hierarchy for the default traffic-class group
```

interface	resource	shaping rate	shared shaping rate	assured rate or weight
-----	-----	-----	-----	-----
ethernet Eth9/0	ethernet port			wgt 8
ethernet Eth9/0	ethernet queue			wgt 8
svlan Eth9/0 svlan 2	svlan node			wgt 8
vlan Eth9/0.1	vlan node			wgt 1
vlan Eth9/0.1	vlan queue best-effort		2000000	wgt 8
vlan Eth9/0.2	vlan node			wgt 3
vlan Eth9/0.2	vlan queue video	2000000		wgt 8
vlan Eth9/0.2	vlan queue best-effort		6000000	wgt 8
vlan Eth9/0.3	vlan node			wgt 6
vlan Eth9/0.3	vlan queue video	3000000		wgt 8
vlan Eth9/0.3	vlan queue best-effort		8000000	wgt 8
Scheduler hierarchy for traffic-class group EF				

interface	resource	shaping rate	shared shaping rate	assured rate or weight
-----	-----	-----	-----	-----
ethernet Eth9/0	ethernet group node EF			wgt 8
svlan Eth9/0 svlan 2	svlan node EF			wgt 8
vlan Eth9/0.2	vlan queue EF voice	100000		wgt 8
vlan Eth9/0.3	vlan queue EF voice	300000		wgt 8

Complete Configuration Example

You can use the complete configuration examples provided for each of the configurations in your own network. To customize the configuration example for your needs, copy the text into a text editor, and modify it.

To use the example for immediate use, copy it to the local console or Telnet session from which you access the router.

You can also save the example as a script (.scr) file that executes the commands as though they were entered at the terminal. For information about executing .scr files, see *JunosE System Basics Configuration Guide*.

QoS Administrator Configuration

From Global Configuration mode:

```
! Configure traffic classes and traffic-class groups.
traffic-class best-effort
exit
traffic-class video
exit
traffic-class voice
```

```

fabric-strict-priority
exit
traffic-class-group EF auto-strict-priority
traffic-class voice
exit
!Configure the max-subscriber-bandwidth parameter definition.
qos-parameter-define max-subscriber-bandwidth
controlled-interface-type vlan
instance-interface-type vlan
range 512000 8192000
exit
!Configure the subscriber-weight parameter definition.
qos-parameter-define subscriber-weight
controlled-interface-type vlan
instance-interface-type vlan
range 1 6
exit
!Configure the max-subscriber-video parameter definition.
qos-parameter-define max-subscriber-video-bandwidth
controlled-interface-type vlan
instance-interface-type vlan
instance-interface-type svlan
range 1000000 5000000
exit
!Configure the max-100Kbps-voice-calls parameter definition.
qos-parameter-define max-100Kbps-voice-calls
controlled-interface-type vlan
instance-interface-type vlan
range 1 3
exit
! Configure the subscriber-best-effort scheduler profile.
scheduler-profile subscriber-best-effort
shared-shaping-rate max-subscriber-bandwidth auto
exit
! Configure the subscriber-video scheduler profile.
scheduler-profile subscriber-video
shaping-rate max-subscriber-video-bandwidth
exit
! Configure the subscriber-weight scheduler profile.
scheduler-profile subscriber-weight
weight subscriber-weight
exit
! Configure the subscriber-voice scheduler profile.
scheduler-profile subscriber-voice
shaping-rate max-100Kbps-voice-calls * 100000
exit
! Configure the subscriber-data-service QoS profile.
qos-profile subscriber-data-service
svlan node
vlan node scheduler-profile subscriber-weight
vlan queue traffic-class best-effort scheduler-profile subscriber-best-effort
exit
! Configure the subscriber-triple-play QoS profile.
qos-profile subscriber-triple-play
svlan node
vlan node scheduler-profile subscriber-weight

```

```

svlan node group EF
vlan queue traffic-class best-effort scheduler-profile subscriber-best-effort
vlan queue traffic-class video scheduler-profile subscriber-video
vlan queue traffic-class voice scheduler-profile subscriber-voice
exit
! Configure the ethernet-default QoS profile.
qos-profile ethernet-default
no ip node
no ip queue traffic-class best-effort
exit
! Attach the QoS profile to the VLAN and S-VLAN subinterfaces.
interface fastEthernet 9/0
encapsulation vlan
exit
interface fastEthernet 9/0.1
svlan id 2 1
ip address 192.1.1.1 255.255.255.0
interface fastEthernet 9/0.2
svlan id 2 2
ip address 192.2.1.1 255.255.255.0
exit
interface fastEthernet 9/0.3
svlan id 2 3
ip address 192.3.1.1 255.255.255.0
exit

```

QoS Client Configuration

From Global Configuration mode:

```

! Configure the max-subscriber-bandwidth, subscriber-weight,
max-subscriber-video-bandwidth, and max-100Kbps-voice-calls global parameter
instances.
qos-parameter max-subscriber-bandwidth 2000000
qos-parameter subscriber-weight 1
qos-parameter max-subscriber-video-bandwidth 2000000
qos-parameter max-100Kbps-voice-calls 1
! Configure a global parameter instance for individual DSLAMs.
interface fastEthernet 9/0
svlan 1 qos-parameter max-subscriber-video-bandwidth 1000000
exit
! Configure the basic-data service for Subscriber 1.
interface fastEthernet 9/0.1
qos-profile subscriber-data-service
exit
! Configure the basic triple-play service for Subscriber 2.
interface fastEthernet 9/0.2
qos-parameter max-subscriber-bandwidth 6000000
qos-parameter subscriber-weight 3
qos-parameter max-subscriber-video-bandwidth 2000000
qos-parameter max-100Kbps-voice-calls 1
qos-profile subscriber-triple-play
exit
! Configure the enhanced triple-play service for Subscriber 3.
interface fastEthernet 9/0.3
qos-parameter max-subscriber-bandwidth 8000000

```

```

qos-parameter subscriber-weight 6
qos-parameter max-subscriber-video-bandwidth 3000000
qos-parameter max-100Kbps-voice-calls 3
qos-profile subscriber-triple-play
exit

```

Example: QoS Parameter Configuration for IP Multicast Bandwidth Adjustment

In this example, a QoS administrator configures a QoS parameter definition to associate with the IP multicast bandwidth adjustment application.

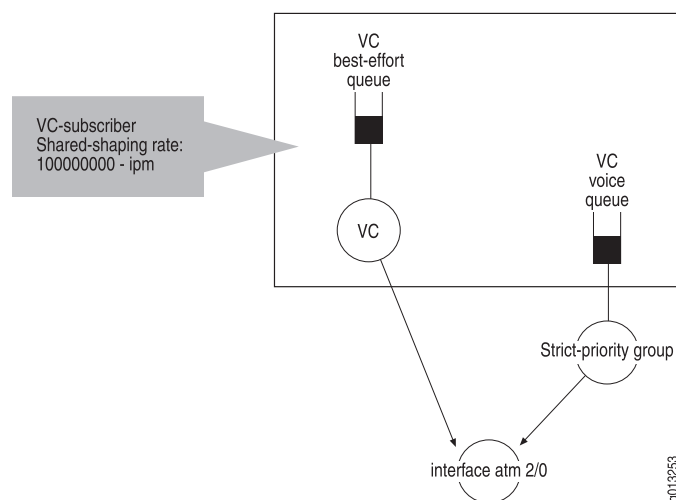
The QoS administrator configures the parameter definition to perform the QoS adjustment on an ATM VC subscriber. By specifying the **ip-multicast** keyword with the **qos-parameter-define** command, the IP parameter instances are created when the Internet Group Management Protocol (IGMP) joins and leaves.

When you specify a controlled-interface type for atm-vc, the system explicitly creates a parameter instance at the ATM VC with a value that is equal to the sum of the IP adjustments above this interface. This parameter value is referred by a scheduler profile and a QoS profile to create the QoS scheduler hierarchy that adjusts the shared-shaping rate when IGMP joins and leaves.

This subscriber has data, voice, and video service with total subscriber bandwidth of 10 Mbps. Voice traffic is shaped at 192 Kbps and belongs to the strict priority group. Video traffic is provided by the IP multicast bandwidth adjustment application and its rate is configured in the bandwidth map.

Figure 17 on page 104 shows the scheduler hierarchy built in this configuration.

Figure 17: Scheduler Hierarchy with QoS Adjustment for IP Multicast



Configuring Traffic Classes and Traffic-Class Groups

The QoS administrator configures the traffic classes and traffic-class groups for best-effort data and voice services. The QoS administrator does not need to configure a traffic class for the video service because it is transmitted through the IP multicast connection.

1. Configure the traffic classes.
 - a. Configure the traffic class named best-effort.
 - b. Configure the traffic class named voice.

```
host1(config)#traffic-class voice
host1(config-traffic-class)#exit
```

```
host1(config)#traffic-class best-effort
host1(config-traffic-class)#exit
```

2. Configure a traffic-class group for low-latency expedited forwarding (EF) and add the traffic class for voice service into the traffic-class group EF.
 - a. Configure the EF traffic-class group with strict-priority scheduling.
 - b. Add the traffic class voice to the traffic-class group.

```
host1(config)#traffic-class-group EF auto-strict-priority
host1(config-traffic-class-group)#traffic-class voice
host1(config-traffic-class-group)#exit
```

The remaining traffic class, best-effort, remains in the default traffic-class group.

Configuring the QoS Parameter Definition and Global Parameter Instance

The QoS administrator configures the QoS parameter definition and specifies the IP multicast bandwidth adjustment application. The QoS administrator must configure the parameter as hierarchical.

The QoS scheduler hierarchy is constructed when the subscriber logs on. However, because the parameter instance has not yet been created, the shared-shaping rate is undefined (that is, there is no shaping rate).

Therefore, the QoS administrator creates a global parameter instance to shape the subscriber to the desired bandwidth. The initial value is determined based on the application; in this example, the QoS administrator specifies 0 as the default.

1. Configure the QoS parameter definition ipm, associate it with the ip-multicast application, and assign it as a hierarchical parameter.
2. Configure a controlled-interface type of atm-vc.
3. Configure the global parameter instance.

```
host1(config)#qos-parameter-define ipm application ip-multicast hierarchical
host1(config-qos-parameter-define)#controlled-interface-type atm-vc
host1(config-qos-parameter-define)#exit
host1(config)#qos-parameter ipm 0
```

Therefore, the initial shared-shaping rate is 10 Mbps (10 Mbps - ipm value of 0).

Reference the Parameter Definition Within a Scheduler Profile

The QoS administrator configures the scheduler profile for the ATM VC subscriber and configures the shared-shaping rate. When a scheduler profile references the parameter instance, it enables the IP multicast bandwidth adjustment application to adjust the subscriber bandwidth to account for the video traffic.

The QoS administrator then configures the scheduler profile to shape voice traffic.

1. Configure the scheduler profile for the ATM VC subscriber.
 - a. Configure the scheduler profile named vc-subscriber.
 - b. Configure the shared-shaping rate by referencing an expression to limit the subscriber bandwidth to 10 Mbps.

```
host1(config)#scheduler-profile vc-subscriber
host1(config-scheduler-profile)#shared-shaping-rate 10000000 - ipm burst 50
milliseconds auto
host1(config-scheduler-profile)#exit
```

2. Configure the scheduler profile for shaping voice traffic.
 - a. Configure the scheduler profile named 192K.
 - b. Configure the shaping rate at 1920000.

```
host1(config)#scheduler-profile 192K
host1(config-scheduler-profile)#shaping rate 192000
host1(config-scheduler-profile)#exit
```

Adding the Scheduler Profiles to a QoS Profile

The IP multicast adjustment application is initialized when IGMP joins or leaves. The QoS administrator specifies the scheduler hierarchy by using a QoS profile rule that refers to a scheduler profile with a parameter expression.

1. Create the QoS profile named ipm-adjusted.
2. Configure a queue for ATM VC subinterfaces with the best-effort traffic class.
3. Configure a queue for ATM VC subinterfaces with the voice traffic class and reference the 192K scheduler profile.
4. Configure a node for ATM VC subinterfaces and reference the scheduler profile vc-subscriber.

```
host1(config)#qos-profile ipm-adjusted
host1(config-qos-profile)#atm-vc queue traffic-class best-effort
host1(config-qos-profile)#atm-vc queue traffic-class voice scheduler-profile 192k
host1(config-qos-profile)#atm-vc node scheduler-profile vc-subscriber
host1(config-qos-profile)#exit
```

Attaching the Parameter Definition to an Interface

The QoS administrator creates a logical interface and attaches the parameter definition. The scheduler hierarchy is created when the QoS administrator creates the interface.

1. Configure the ATM interface in slot 2, port 0 as a point-to-point ATM interface.
2. Configure the ATM PVC with aal5snap encapsulation.
3. Attach the QoS profile vc-subscriber to the subinterface.
4. Configure the IP address for the ATM subinterface.

```
host1(config)#interface atm 2/0
host1(config-if)#interface atm 2/0.1 point-to-point
host1(config-subif)#atm pvc 100 0 100 aal5snap
```

```
host1(config-subif)#qos-profile ipm-adjusted
host1(config-subif)#ip address 1.1.1.255.255.255.0
```

IP Multicast Bandwidth Adjustment

When an IGMP join occurs, the IP multicast bandwidth adjustment application creates the parameter instance ipm for the IP interface and the ATM VC subinterface. Because the shared-shaping rate of the ATM VC references the ipm parameter, the rate is recalculated. If the ipm parameter has a value of 2 Mbps, the resulting shared-shaping rate is 8 Mbps (10 Mbps - 2 = 8 Mbps).

When another IGMP join occurs, the IP multicast bandwidth adjustment application recalculates the value for parameter ipm and configures it to another value (for example, 7 Mbps). The system readjusts the ipm at the ATM VC and readjusts the shared-shaping rate. If the voice traffic is 100 Kbps, then the best-effort traffic is 2.9 Mbps.

When an IGMP leave occurs, the IP multicast bandwidth adjustment application configures the ipm parameter instance with a new value and readjusts the shared-shaping rate.

Monitoring the Configuration

After completing the configuration, the QoS administrator can monitor it by issuing **show** commands.

1. To display the traffic classes for best-effort and voice, issue the **show traffic-class** command.

```
host1#show traffic-class
```

traffic class	fabric weight	fabric strict priority
best-effort	8	no
voice	8	no

2. To display the traffic-class group, issue the **show traffic-class-group** command.

```
host1#show traffic-class-group
traffic-class-group EF auto-strict-priority
traffic-class voice
```

3. To display the scheduler profile settings for vc-subscriber and 192K, issue the **show scheduler-profile** command.

```
host1#show scheduler-profile
```

scheduler	shaping rate	shaping burst	weight	strict priority	assured rate
default	<none>	<none>	8	no	<none>
vc-subscriber	<none>	<none>	8	no	<none>
192k	192000	default	8	no	<none>

scheduler	shared shaping rate	shared shaping burst	shared shaping constituent	shared shaping mode
default	<none>	<none>	<none>	<none>
vc-subscriber	10000000 - ipm	50 bytes	<none>	simple implicit
192k	<none>	<none>	<none>	<none>

4. To display the attachments on all QoS profiles, including ipm-adjust, issue the **show qos-profile references** command.

```

host1#show qos-profile references
      qos profile                                attachment
-----
atm-default                                (qos-port-type-profile)
serial-default                            (qos-port-type-profile)
ethernet-default                          (qos-port-type-profile)
server-default                            (qos-port-type-profile)
ipm-adjust                                atm-vc ATM2/0.1

Port attachments:      4
Interface attachments: 1
Not attached:          0

```

5. To display the settings for the ipm-adjust QoS profile, issue the **show qos-profile** command.

```

host1#show qos-profile ipm-adjust
qos-profile ipm-adjust:
t-class interface rule traffic scheduler queue drop statistics
group  type      type class  profile  profile profile profile
-----
      atm-vc    node      vc-subscriber
      atm-vc    queue best-effort default default default default
EF     atm-vc    queue voice    192k    default default default

```

6. To display the settings for the ipm QoS parameter definition, issue the **show qos-parameter-define** command.

```

host1#show qos-parameter-define
      controlled instance subscriber
parameter interface interface interface value
name      types      types      types      range
-----
ipm       atm-vc      <none>    <none>    <none>

parameter
name      properties
-----
ipm       ip-multicast-adjustment, hierarchical

```

7. To display global and interface attachments on the ipm QoS parameter instance, issue the **show qos-parameter references** command.

```

host1#show qos-parameter references
      parameter
interface name      value
-----
global    ipm        0

Global parameter instances: 1
Parameter instances reported: 1

host1#show qos-parameter references interface atm 1/0.1
      parameter instance
interface name      value      Type
-----

```

```

atm-vc ATM1/0.1 ipm          200 hierarchical
ip ATM1/0.1 ipm              200 ip-multicast

```

```

Explicit parameter instances: 0
Heirarchical parameter instances: 1
IP multicast parameter instances: 1
Parameter instances reported: 2

```

8. To display the queue forwarding rates for the ATM VC and IP interfaces on the ATM interface in slot 2, port 0, issue the **show egress-queue rates** command.

```
host1#show egress-queue rates interface atm 2/0.1
```

interface	traffic class	forwarded rate	aggregate drop rate	minimum rate	maximum rate
atm-vc ATM2/0.1	voice	0	0	192000	192000
ip ATM2/0.1	best-effort	0	0	0	10000000

```

Queues reported: 2
Queues filtered (under threshold): 0
* Queues disabled (no rate period): 0
**Queues disabled (no resources): 0
Total queues: 2

```

9. To display the shared shaper settings for the ATM VC on the ATM interface in slot 2, port 0, issue the **show qos shared-shaper** command.

```
host1#show qos shared-shaper interface atm 2/0.1
```

interface	resource	shared shaping rate	shaping rate	other rate
atm-vc ATM2/0.1	A atm-vc node	10000000		10000000
	A atm-vc queue EF voice		192000	

```

Total shared shapers: 1
Total constituents: 2
Total shared shaper failovers: 0
Compound shared shapers are supported.

```

Complete Configuration Example

You can use the complete configuration examples provided for each of the configurations in your own network. To customize the configuration example for your needs, copy the text into a text editor, and modify it.

To use the example for immediate use, copy it to the local console or Telnet session from which you access the router.

You can also save the example as a script (.scr) file that executes the commands as though they were entered at the terminal. For information about executing .scr files, see *JunosE System Basics Configuration Guide*.

From Global Configuration mode:

```

! Create the voice traffic class.
traffic-class voice
exit

```

```

traffic-class best-effort
exit
traffic-class-group EF auto-strict-priority
traffic-class best-effort
exit
! Create the ipm QoS parameter definition.
qos-parameter-define ipm application ip-multicast hierarchical
controlled-interface-type atm-vc
exit
! Create a global parameter instance of the ipm QoS parameter.
qos-parameter ipm 0
! Configure the vc-subscriber and 192K scheduler profiles.
scheduler-profile vc-subscriber
shared-shaping-rate 10000000 - ipm burst 50 milliseconds auto
exit
scheduler-profile 192K
shaping-rate 192000
exit
! Add the scheduler profiles to the ipm-adjusted QoS profile.
qos-profile ipm-adjusted
atm-vc queue traffic-class best-effort
atm-vc queue traffic-class voice scheduler-profile 192k
atm-vc node scheduler-profile vc-subscriber
exit
! Attach the parameter definition to an interface.
interface atm 2/0.1 point-to-point
atm pvc 100 0 100 aal5snap
qos-profile ipm-adjusted
ip address 1.1.1.1 255.255.255

```

Related Documentation

- [IP Multicast Bandwidth Adjustment for QoS Overview on page 31](#)

Example: QoS Parameter Configuration for QoS Downstream Rate

This example illustrates how to use parameters to control the downstream rate obtained from ANCP.

In this example, the subscribers on the 0.1 access loop are configured on VLAN1. They subscribe to voice, video, and data traffic with a bandwidth of 10 Mbps. Subscribers on the 1.1 access loop are configured on VLAN2, and subscribe to 1 Mbps of data traffic.

[Table 12 on page 110](#) lists the shaping mode and shaping rate information received by the QoS downstream rate application upon access loop synchronization. The parameter instances are created with these values.

Table 12: Shaping Rate and Shaping Mode

	VLAN1	VLAN2
Shaping mode	Cell	Cell
Shaping rate	10000000 bps	1000000 bps

Configuring Traffic Classes	<p>The QoS administrator configures the traffic classes for voice and video services.</p> <ol style="list-style-type: none"> 1. Configure the traffic class named voice. <pre data-bbox="529 373 878 426">host1(config)#traffic-class voice host1(config-traffic-class)#exit</pre> 2. Configure the traffic class named video. <pre data-bbox="529 499 878 552">host1(config)#traffic-class video host1(config-traffic-class)#exit</pre>
Configuring the QoS Parameter Definition for QoS Downstream Rate	<p>The QoS administrator configures a parameter definition for the QoS downstream rate application. Using subscriber-interface types, the QoS administrator then enables ANCP to create parameter instances of the QoS downstream rate application.</p> <ol style="list-style-type: none"> 1. Configure a parameter definition named ancpVlan. <pre data-bbox="529 772 1401 793">host1(config)#qos-parameter-define ancpVlan application qos-downstream-rate</pre> 2. Define the controlled-interface types for vlan and ip to adjust the shaping rate for the VLAN and IP queues. <ol style="list-style-type: none"> a. Configure the controlled-interface type for VLAN. b. Configure the controlled-interface type for IP. <pre data-bbox="565 997 1287 1050">host1(config-qos-parameter-define)#controlled-interface-type vlan host1(config-qos-parameter-define)#controlled-interface-type ip</pre> 3. Define the subscriber-interface types for vlan and ethernet. <pre data-bbox="529 1129 1304 1207">host1(config-qos-parameter-define)#subscriber-interface-type vlan host1(config-qos-parameter-define)#subscriber-interface-type ethernet host1(config-qos-parameter-define)#exit</pre>
Configuring the QoS Parameter Definition for QoS Cell Mode	<p>The QoS administrator then configures the QoS shaping mode using the QoS cell mode application. Using subscriber-interface types, the QoS administrator then enables ANCP to create parameter instances using the QoS cell mode application.</p> <ol style="list-style-type: none"> 1. Configure a parameter definition named cellmodeVlan. <pre data-bbox="529 1430 1369 1451">host1(config)#qos-parameter-define cellmodeVlan application qos-cell-mode</pre> 2. Define the controlled-interface types for vlan and ip for the shaping mode. <pre data-bbox="529 1524 1255 1602">host1(config-qos-parameter-define)#controlled-interface-type vlan host1(config-qos-parameter-define)#controlled-interface-type ip host1(config-qos-parameter-define)#exit</pre> 3. Define the subscriber-interface types for vlan and ethernet. <pre data-bbox="529 1675 1304 1759">host1(config-qos-parameter-define)#subscriber-interface-type vlan host1(config-qos-parameter-define)#subscriber-interface-type ethernet host1(config-qos-parameter-define)#exit</pre>
Enabling QoS Adaptive Mode	<p>The QoS administrator enables QoS adaptive mode for ANCP.</p> <ol style="list-style-type: none"> 1. Enter Layer 2 Control Configuration mode.

	<pre>host1(config)#l2c</pre>
	<ol style="list-style-type: none"> 2. Enable QoS adaptive mode for the system. <pre>host1(config-l2c)#qos-adaptive-mode</pre>
Reference the Parameter Definition Within a Scheduler Profile	<p>The QoS administrator configures the shaping rate and the shared-shaping rate within scheduler profiles for the VLAN1 and VLAN2 subscribers.</p> <ol style="list-style-type: none"> 1. Configure the scheduler profile for the subscriber vlan1. <ol style="list-style-type: none"> a. Configure the scheduler profile named vlan1. b. Configure the shared-shaping rate by referencing the ancpVlan parameter with a burst of 10 milliseconds. <pre>host1(config)#scheduler-profile vlan1 host1(config-scheduler-profile)#shared-shaping-rate ancpVlan burst 10 milliseconds auto host1(config-scheduler-profile)#exit</pre> 2. Configure the scheduler profile for the voice service. <ol style="list-style-type: none"> a. Configure the scheduler profile named voice. b. Configure the shaping rate of 100000 with a burst of 10 milliseconds. <pre>host1(config)#scheduler-profile voice host1(config-scheduler-profile)#shaping-rate 100000 burst 10 milliseconds host1(config-scheduler-profile)#exit</pre> 3. Configure the scheduler profile for the video service. <ol style="list-style-type: none"> a. Configure the scheduler profile named video. b. Configure the shaping rate of 8000000 with a burst of 10 milliseconds. <pre>host1(config)#scheduler-profile video host1(config-scheduler-profile)#shaping-rate 8000000 burst 10 milliseconds host1(config-scheduler-profile)#exit</pre> 4. Configure the scheduler profile for the subscriber vlan2. <ol style="list-style-type: none"> a. Configure the scheduler profile named vlan2. b. Configure the shaping rate by referencing the ancpVlan parameter with a burst of 10 milliseconds. <pre>host1(config)#scheduler-profile vlan2 host1(config-scheduler-profile)#shaping-rate ancpVlan burst 10 milliseconds host1(config-scheduler-profile)#exit</pre>
Adding the Scheduler Profiles to a QoS Profile	<p>After configuring the scheduler profiles, the QoS administrator then configures QoS profiles for the VLAN1 and VLAN2 subscribers.</p> <ol style="list-style-type: none"> 1. Configure the vlan1 QoS profile with a shared-shaping rate that matches the downstream rate.

- a. Configure the QoS profile named vlan1.
- b. Configure the vlan node and reference the scheduler profile vlan1.
- c. Configure the vlan queue and reference the voice traffic class and the voice scheduler profile.
- d. Configure the vlan queue and reference the video traffic class and the video scheduler profile.

```
host1(config)#qos-profile vlan1
host1(config-qos-profile)#vlan node scheduler-profile vlan1
host1(config-qos-profile)#vlan queue traffic-class voice scheduler-profile voice
host1(config-qos-profile)#vlan queue traffic-class video scheduler-profile video
host1(config-qos-profile)#exit
```

2. Configure the vlan2 QoS profile with a shaping rate of 1 Mbps.

- a. Configure the QoS profile named vlan2.
- b. Configure the vlan node and reference the scheduler profile vlan2.

```
host1(config)#qos-profile vlan2
host1(config-qos-profile)#vlan node scheduler-profile vlan2
host1(config-qos-profile)#exit
```

Attaching the QoS Profile to an Interface

The QoS administrator creates logical interfaces for VLAN1 and VLAN2 and attaches the QoS profiles to them. As the subscribers log in, ANCP creates the parameter instances for cellmodeVlan and ancpVlan using RADIUS VSAs.

1. Attach the vlan1 QoS profile to VLAN1.
 - a. Configure the Gigabit Ethernet interface in slot 6, adapter 0, port 0.
 - b. Configure the VLAN major interface.
 - c. Configure the Gigabit Ethernet interface in slot 6, adapter 0, port 0, subinterface 1.
 - d. Assign VLAN ID of 1.
 - e. Attach the QoS profile vc1 to the interface.

```
host1(config)#interface gigabitEthernet 6/0/0
host1(config-if)#encapsulation vlan
host1(config-if)#interface gigabitEthernet 6/0/0.1
host1(config-if)#vlan id 1
host1(config-if)#qos-profile vlan1
host1(config-if)#exit
```

2. Attach the vlan2 QoS profile to VLAN2.
 - a. Specify the Gigabit Ethernet interface in slot 6, adapter 0, port 1.
 - b. Assign a VLAN ID of 2.
 - c. Attach the QoS profile vlan2 to the interface.

```
host1(config-if)#interface gigabitEthernet 6/0/1.1
host1(config-if)#vlan id 2
host1(config-if)#qos-profile vlan2
host1(config-if)#exit
```

Complete Configuration Example

You can use the complete configuration examples provided for each of the configurations in your own network. To customize the configuration example for your needs, copy the text into a text editor, and modify it.

To use the example for immediate use, copy it to the local console or Telnet session from which you access the router.

You can also save the example as a script (.scr) file that executes the commands as though they were entered at the terminal. For information about executing .scr files, see *JunosE System Basics Configuration Guide*.

From Global Configuration mode:

```
! Configure the traffic-classes for video and voice.
traffic-class voice
exit
traffic-class video
exit
! Create the ancpVlan QoS parameter definition.
qos-parameter-define ancpVlan application qos-downstream-rate
controlled-interface-type vlan
controlled-interface-type ip
instance-interface-type vlan
instance-interface-type ethernet
exit
! Create the cellmodeVlan QoS parameter definition.
qos-parameter-define cellmodeVlan application qos-cell-mode
controlled-interface-type vlan
controlled-interface-type ip
instance-interface-type vlan
instance-interface-type ethernet
exit
! Enable QoS adaptive mode for ANCP.
l2c
qos-adaptive-mode
exit
! Configure the vlan1 and vlan2 scheduler profiles.
scheduler-profile vlan1
shared-shaping-rate ancpVlan burst 10 milliseconds auto
exit
scheduler-profile voice
shaping-rate 100000 burst 10 milliseconds
exit
scheduler-profile video
shaping-rate 8000000 burst 10 milliseconds
exit
scheduler-profile vlan2
```

```

shaping-rate ancpVlan burst 10 milliseconds
exit
! Add the scheduler profiles to the vlan1 and vlan2 QoS profiles.
qos-profile vlan1
vlan node scheduler-profile vlan1
vlan queue traffic-class voice scheduler-profile voice
vlan queue traffic-class video scheduler-profile video
exit
qos-profile vlan2
vlan node scheduler-profile vlan2
exit
! Configure the QoS downstream rate adjustment for VLAN1 and VLAN2.
interface gigabitEthernet 6/0/0
encapsulation vlan
interface gigabitEthernet 6/0/1.1
vlan id 1
qos-profile vlan1
exit
interface gigabitEthernet 6/0/1.1
vlan id 2
qos-profile vlan2
exit

```

Related Documentation

- [QoS Downstream Rate Application Overview](#)


Example: QoS Parameter Configuration for Hierarchical Parameters

The example in this section illustrates how to configure hierarchical parameters for VLANs and S-VLANs.

[Figure 18 on page 116](#) shows the QoS scheduler hierarchy that the QoS client creates for the VLANs and S-VLANs in the interface stack. The QoS client creates explicit parameter instances using the parameter definition `max-sub-bw` to shape rates at the VLAN subinterfaces 100 and 101.

An S-VLAN node is located below the two VLAN nodes in the interface stack. The QoS client creates an implicit parameter instance by applying a shaper to the S-VLAN subinterface 10 that equals the total rate at the VLANs (3072000).

Figure 18: Hierarchical Parameters Scheduler Hierarchy



Procedure for QoS Administrators

This section describes the procedures to configure the scheduler hierarchy shown in [Figure 18 on page 116](#) by using QoS parameters.

Configuring the Parameter Definition

The QoS administrator configures the parameter definition for the maximum subscriber bandwidth.

To configure a parameter definition for the maximum subscriber bandwidth:

1. Configure the parameter definition named max-sub-bw.
2. Enable the parameter to control S-VLANs.
3. Enable the parameter to control VLANs.
4. Enable the parameter to have instances created on VLAN subinterfaces.
5. Specify that the QoS client can create the parameter instance for VLANs, which represent subscribers.

```

host1(config)#qos-parameter-define max-sub-bw hierarchical
host1(config-qos-parameter-define)#controlled-interface-type svlan
host1(config-qos-parameter-define)#controlled-interface-type vlan
host1(config-qos-parameter-define)#instance-interface-type vlan
host1(config-qos-parameter-define)#subscriber-interface-type vlan
host1(config-qos-parameter-define)#exit
    
```

Configuring the Scheduler Profiles

The QoS administrator can then reference the parameter definition within a scheduler profile, which defines the shaping rates for the parameter.

1. Configure a scheduler profile to shape the throughput the explicit QoS parameters for VLANs.
 - a. Configure the scheduler profile named sp-shape-cvlan.
 - b. Configure the shaping rate by referencing the parameter max-sub-bw.

```
host1(config)#scheduler-profile sp-shape-cvlan
host1(config-scheduler-profile)#shaping-rate max-sub-bw
host1(config-scheduler-profile)#exit
```

2. Configure a scheduler profile to shape the S-VLAN throughput.
 - a. Configure the scheduler profile named sp-shape-svlan.
 - b. Configure the shaping rate by referencing the parameter max-sub-bw.

```
host1(config)#scheduler-profile sp-shape-svlan
host1(config-scheduler-profile)#shaping-rate max-sub-bw
host1(config-scheduler-profile)#exit
```

Configuring the QoS Profiles

By referencing the scheduler profiles within QoS profiles, the QoS administrator creates the scheduler hierarchy. In this portion of the example, the QoS administrator configures QoS profiles for the VLAN and the S-VLAN.

1. Configure the QoS profile for the VLAN interfaces.
 - a. Configure the QoS profile named qp-shape-cvlan.
 - b. Configure the VLAN queue and reference the best-effort traffic class.
 - c. Configure the VLAN node and reference the scheduler profile for shaping VLANs.

```
host1(config)#qos-profile qp-shape-cvlan
host1(config-qos-profile)#vlan queue traffic-class best-effort
host1(config-qos-profile)#vlan node scheduler-profile sp-shape-cvlan
host1(config-qos-profile)#exit
```

2. Configure the QoS profile for the S-VLAN interface.
 - a. Configure the QoS profile named qp-shape-svlan.
 - b. Configure the S-VLAN node and reference the scheduler profile sp-shape-svlan.

```
host1(config)#qos-profile qp-shape-svlan
host1(config-qos-profile)#svlan node scheduler-profile sp-shape-svlan
host1(config-qos-profile)#exit
```

Procedure for QoS Clients

This section describes procedures to create parameter instances at VLAN subinterface 100 and VLAN subinterface 101.

1. Create an explicit parameter instance at VLAN subinterface 100.
 - a. Specify the Gigabit Ethernet interface in slot 2, port 0.
 - b. Configure the VLAN major interface.
 - c. Configure the VLAN subinterface at slot 2, port 0, subinterface 100.
 - d. Assign an S-VLAN ID of 10 and a VLAN ID of 100 to the VLAN subinterface.

- e. Attach the max-sub-bw QoS parameter to the subinterface with a value of 1024000.

- f. Attach the qp-shape-cvlan QoS profile to the subinterface.

```
host1(config)#interface gigabitEthernet 2/0
host1(config-if)#encapsulation vlan
host1(config)#interface gigabitEthernet 2/0.100
host1(config-if)#svlan id 10 100
host1(config-if)#qos-parameter max-sub-bw 1024000
host1(config-if)#qos-profile qp-shape-cvlan
host1(config-if)#exit
```

2. Create an explicit parameter instance at VLAN subinterface 101.

- a. Specify the VLAN subinterface 101 in slot 2, port 0.
- b. Assign an S-VLAN ID of 10 and a VLAN ID of 101 to the VLAN subinterface.
- c. Attach the max-sub-bw QoS parameter to the subinterface with a value of 2048000.
- d. Attach the qp-shape-cvlan QoS profile to the subinterface.

```
host1(config-if)#interface gigabitEthernet 2/0.101
host1(config-if)#svlan id 10 101
host1(config-if)#qos-parameter max-sub-bw 2048000
host1(config-if)#qos-profile qp-shape-cvlan
host1(config-if)#exit
```

3. Create an implicit parameter instance at S-VLAN subinterface 10.

- a. Specify the Gigabit Ethernet interface at slot 2, port 0.
- b. Attach the qp-shape-svlan QoS profile to the node at S-VLAN subinterface 10.

```
host1(config)#interface gigabitEthernet 2/0
host1(config-if)#svlan 10 qos-profile qp-shape-svlan
```

Monitoring Hierarchical QoS Parameters

After completing the configuration, both the QoS administrator and the QoS client can monitor it by issuing the **show qos-parameter references** command. To display the information about hierarchical parameter instances, you must specify the Gigabit Ethernet interface.

```
host1#show qos-parameter max-sub-bw references interface gigabitEthernet 2/0
```

interface	parameter name	value	instance Type
GigabitEthernet2/0 svlan 10	max-sub-bw	3072000	hierarchical
GigabitEthernet2/0.100	max-sub-bw	1024000	explicit
GigabitEthernet2/0.101	max-sub-bw	2048000	explicit
Explicit parameter instances:		2	
Hierarchical parameter instances:		1	
IP multicast parameter instances:		0	
Parameter instances reported:		3	

Complete Configuration Example

You can use the complete configuration examples provided for each of the configurations in your own network. To customize the configuration example for your needs, copy the text into a text editor, and modify it.

To use the example for immediate use, copy it to the local console or Telnet session from which you access the router.

You can also save the example as a script (.scr) file that executes the commands as though they were entered at the terminal. For information about executing .scr files, see *JunosE System Basics Configuration Guide*.

QoS Administrator Configuration

From Global Configuration mode:

```
! Configure the max-sub-bw QoS parameter definition.
qos-parameter-define max-sub-bw hierarchical
controlled-interface-type svlan
controlled-interface-type vlan
instance-interface-type vlan
subscriber-interface-type vlan
exit
! Configure the sp-shape-cvlan and sp-shape-svlan scheduler profiles.
scheduler-profile sp-shape-cvlan
shaping-rate max-sub-bw
exit
scheduler-profile sp-shape-svlan
shaping-rate max-sub-bw
exit
! Configure the qp-shape-cvlan and qp-shape-svlan QoS profiles.
qos-profile qp-shape-cvlan
vlan queue traffic-class best-effort
vlan node scheduler-profile sp-shape-cvlan
exit
qos-profile qp-shape-svlan
svlan node scheduler-profile sp-shape-svlan
exit
```

QoS Client Configuration

From Global Configuration mode:

```
! Configure the QoS parameter max-sub-bw for VLAN subinterface 100.
interface gigabitEthernet 2/0
encapsulation vlan
interface gigabitEthernet 2/0.100
svlan id 10 100
qos-parameter max-sub-bw 1024000
qos-profile qp-shape-cvlan
exit
! Configure the QoS parameter max-sub-bw for VLAN subinterface 101.
interface gigabitEthernet 2/0.101
svlan id 10 101
```

```
qos-parameter max-sub-bw 2048000
qos-profile qp-shape-cvlan
! Attach the QoS profile to the S-VLAN subinterface 10.
interface gigabitEthernet 2/0
svlan 10 qos-profile qp-shape-svlan
```

Related Documentation

- [Hierarchical QoS Parameters Overview on page 31](#)

CHAPTER 15

Configuration Commands

average-length-exponent

Syntax `average-length-exponent exponent`
 `no average-length-exponent`

Release Information Command introduced before JunosE Release 7.1.0.

Description Specifies the exponent used to weight the average queue length over time, controlling WRED responsiveness. The **no** version negates the average-length-exponent.

Options • *exponent*—Total average queue length (TAQL) coefficient

Mode Drop Profile Configuration

Related Documentation • Configuring RED
 • [Configuring WRED on page 54](#)

buffer-weight

Syntax	<code>buffer-weight <i>bufferWeight</i></code> <code>no buffer-weight</code>
Release Information	Command introduced before JunosE Release 7.1.0.
Description	Sets the buffer weight of the queue. The no version returns the queue to its default buffer weight.
Options	<ul style="list-style-type: none">• <i>bufferWeight</i>—Range 1–63; default value is 8
Mode	Queue Profile Configuration
Related Documentation	<ul style="list-style-type: none">• Configuring Queue Profiles to Manage Buffers and Thresholds on page 53

committed-length

Syntax committed-length *minimumCommittedLength* [*maximumCommittedLength*]
no committed-length

Release Information Command introduced before JunosE Release 7.1.0.

Description Sets minimum and maximum constraints for the queue's committed lengths. The **no** version removes constraints on the queue's committed length.

- Options**
- *minimumCommittedLength*—Range 0–1073741824
 - *maximumCommittedLength*—Range 0–1073741824

Mode Queue Profile Configuration

Related Documentation

- [Configuring Queue Profiles to Manage Buffers and Thresholds on page 53](#)

committed-threshold

Syntax	<code>committed-threshold { percent <i>MinThresholdPercent</i> <i>MaxThresholdPercent</i> <i>MinThresholdBytes</i> <i>MaxThresholdBytes</i> } <i>MaxDropProbability</i></code> <code>no committed-threshold</code>
Release Information	Command introduced before JunosE Release 7.1.0.
Description	Specifies the committed queue thresholds and maximum drop probability. The no version removes committed threshold.
Options	<ul style="list-style-type: none">• <code>percent</code>—Specifies committed queue thresholds as percentages• <i>MinThresholdPercent</i>—Minimum queue threshold as a percentage of queue length• <i>MaxThresholdPercent</i>—Maximum queue threshold as a percentage of queue length• <i>MinThresholdBytes</i>—Minimum queue threshold in bytes• <i>MaxThresholdBytes</i>—Maximum queue threshold in bytes• <i>MaxDropProbability</i>—Maximum drop probability
Mode	Drop Profile Configuration
Related Documentation	<ul style="list-style-type: none">• Configuring RED• Configuring WRED on page 54

conformed-fraction

Syntax `conformed-fraction conformedFraction`
 `no conformed-fraction`

Release Information Command introduced before JunosE Release 7.1.0.

Description Sets the percentage of the total queue that can be occupied before dropping conformed packets. The **no** version returns the conformed fraction to its default setting.

Options • *conformedFraction*—Percentage in the range 0–100; default value is 50

Mode Queue Profile Configuration

Related Documentation • [Configuring Queue Profiles to Manage Buffers and Thresholds on page 53](#)

conformed-length

Syntax conformed-length *minimumConformedLength* [*maximumConformedLength*]
no conformed-length

Release Information Command introduced before JunosE Release 7.1.0.

Description Sets minimum and maximum constraints for the queue's conformed lengths. The **no** version removes constraints on the queue's conformed length.

- Options**
- *minimumConformedLength*—Number in the range 0–1073741824
 - *maximumConformedLength*—Number in the range 0–1073741824

Mode Queue Profile Configuration

Related Documentation

- [Configuring Queue Profiles to Manage Buffers and Thresholds on page 53](#)

conformed-threshold

Syntax	<p>conformed-threshold { percent <i>MinThresholdPercent</i> <i>MaxThresholdPercent</i> <i>MinThresholdBytes</i> <i>MaxThresholdBytes</i> } <i>MaxDropProbability</i></p> <p>no conformed-threshold</p>
Release Information	Command introduced before JunosE Release 7.1.0.
Description	Specifies the conformed queue thresholds and maximum drop probability. The no version removes conformed threshold.
Options	<ul style="list-style-type: none"> • percent—Specifies conformed queue thresholds as percentages • <i>MinThresholdPercent</i>—Minimum queue threshold as a percentage of queue length • <i>MaxThresholdPercent</i>—Maximum queue threshold as a percentage of queue length • <i>MinThresholdBytes</i>—Minimum queue threshold in bytes • <i>MaxThresholdBytes</i>—Maximum queue threshold in bytes • <i>MaxDropProbability</i>—Maximum drop probability
Mode	Drop Profile Configuration
Related Documentation	<ul style="list-style-type: none"> • Configuring RED • Configuring WRED on page 54

controlled-interface-type

Syntax	<pre>controlled-interface-type { <i>controlledInterfaceType</i> set superset }</pre> <pre>no controlled-interface-type { <i>controlledInterfaceType</i> all }</pre>
Release Information	<p>Command introduced in JunosE Release 7.1.0.</p> <p>set and superset keywords added in JunosE Release 9.2.0.</p>
Description	<p>Assigns controlled-interface types to a QoS parameter definition. Controlled-interface types specify the types of logical interfaces whose queues and nodes can be controlled by instances of the parameter definition. You can specify up to four controlled-interface types for each parameter definition. The no version removes the controlled-interface type from the parameter definition.</p>
Options	<ul style="list-style-type: none"> • <i>controlledInterfaceType</i>—One of the following controlled-interface types: atm, atm-vc, atm-vp, bridge, ethernet, fr-vc, ip, ip-tunnel, ipv6, l2tp-session, l2tp-tunnel, lsp, serial, svlan, server-port, vlan • set—Specifies an interface set as a controlled-interface type • superset—Specifies an interface superset as a controlled-interface type • all—Removes all controlled-interface types
Mode	QoS Parameter Definition
Related Documentation	<ul style="list-style-type: none"> • Configuring a Basic Parameter Definition for QoS Administrators on page 75 • Creating a QoS Parameter on an Interface Superset or Interface Set

drop-profile

Syntax [no] drop-profile *dropProfileName*

Release Information Command introduced before JunosE Release 7.1.0.

Description Creates a drop profile. The **no** version removes the drop profile.

Options • *dropProfileName*—Name for the drop profile

Mode Global Configuration

Related Documentation • Configuring RED
• [Configuring WRED on page 54](#)

exceeded-fraction

Syntax	<code>exceeded-fraction <i>exceededFraction</i></code> <code>no exceeded-fraction</code>
Release Information	Command introduced before JunosE Release 7.1.0.
Description	Sets the percentage of the total queue length that can be occupied before dropping exceeded packets. The no version returns the exceeded fraction to its default setting.
Options	<ul style="list-style-type: none">• <i>exceededFraction</i>—Percentage range 0–100; default value is 25
Mode	Queue Profile Configuration
Related Documentation	<ul style="list-style-type: none">• Configuring Queue Profiles to Manage Buffers and Thresholds on page 53

exceeded-length

Syntax `exceeded-length minimumExceededLength [maximumExceededLength]`
 `no exceeded-length`

Release Information Command introduced before JunosE Release 7.1.0.

Description Sets minimum and maximum constraints for the queue's exceeded lengths. The **no** version removes constraints on the queue's exceeded length.

Options • *minimumExceededLength*—Range 0–1073741824
 • *maximumExceededLength*—Range 0–1073741824

Mode Queue Profile Configuration

Related Documentation • [Configuring Queue Profiles to Manage Buffers and Thresholds on page 53](#)

exceeded-threshold

Syntax	<code>exceeded-threshold { percent <i>MinThresholdPercent</i> <i>MaxThresholdPercent</i> <i>MinThresholdBytes</i> <i>MaxThresholdBytes</i> } <i>MaxDropProbability</i></code> <code>no exceeded-threshold</code>
Release Information	Command introduced before JunosE Release 7.1.0.
Description	Specifies the exceeded queue thresholds and maximum drop probability. The no version removes exceeded threshold.
Options	<ul style="list-style-type: none">• <i>percent</i>—Specifies <i>exceeded</i> queue thresholds as percentages• <i>MinThresholdPercent</i>—Minimum queue threshold as a percentage of queue length• <i>MaxThresholdPercent</i>—Maximum queue threshold as a percentage of queue length• <i>MinThresholdBytes</i>—Minimum queue threshold in bytes• <i>MaxThresholdBytes</i>—Maximum queue threshold in bytes• <i>MaxDropProbability</i>—Maximum drop probability
Mode	Drop Profile Configuration
Related Documentation	<ul style="list-style-type: none">• Configuring RED• Configuring WRED on page 54

fabric-strict-priority

Syntax [no] fabric-strict-priority

Release Information Command introduced before JunosE Release 7.1.0.

Description Specifies strict priority scheduling for queues in the traffic class in the fabric. The **no** version deletes the strict priority setting.

Mode Traffic Class Configuration

fabric-weight

Syntax `fabric-weight weight`
 `no fabric-weight`

Release Information Command introduced before JunosE Release 7.1.0.

Description Specifies the relative weight for queues in the traffic class in the fabric. The **no** version sets the fabric weight to the default value.

Options • *weight*—Range 1–63; default value is 8

Mode Traffic Class Configuration

instance-interface-type

Syntax	<p>instance-interface-type {<i>instanceInterfaceType</i> set superset }</p> <p>no instance-interface-type { <i>instanceInterfaceType</i> all }</p>
Release Information	<p>Command introduced in JunosE Release 7.1.0.</p> <p>lag keyword added in JunosE Release 8.1.0.</p> <p>set and superset keywords added in JunosE Release 9.2.0.</p>
Description	<p>Assigns an instance-interface type to a QoS parameter definition. Instance-interface types indicate the interfaces for which QoS clients can assign QoS parameter instances. You can specify up to eight instance-interface types for each parameter definition. The no version removes the specified instance-interface type from the parameter definition.</p>
Options	<ul style="list-style-type: none"> • <i>instanceInterfaceType</i>—One of the following instance-interface types: atm, atm-vc, atm-vp, bridge, ethernet, fr-vc, ip, ip-tunnel, ipv6, lag, l2tp-session, l2tp-tunnel, lsp, pppoe, serial, server-port, svlan, vlan • set—Specifies an interface set as an instance-interface type • superset—Specifies an interface superset as an instance-interface type • all—Removes all instance-interface types
Mode	QoS Parameter Definition
Related Documentation	<ul style="list-style-type: none"> • Configuring a Basic Parameter Definition for QoS Administrators on page 75 • Creating a QoS Parameter on an Interface Superset or Interface Set

load-rebalance

Syntax	[no] load-rebalance [period <i>rebalancePeriod</i> start-threshold <i>rebalanceStartThreshold</i> [percent subscribers] stop-threshold <i>rebalanceStopThreshold</i> [percent subscribers] maximum-improvement <i>rebalanceMaximumImprovement</i> [percent subscribers]]
Release Information	Command introduced in JunosE Release 8.1.0.
Description	Configures the QoS algorithm for rebalancing the links in an 802.3ad link aggregation group (LAG). To configure the algorithm to dynamically rebalance the LAG using existing parameters, issue the load-rebalance command without any keywords. The no version restores the default parameters.
Options	<ul style="list-style-type: none"> • <i>rebalancePeriod</i>—Time period for rebalancing in seconds; in the range 0–86400; the default is 60 seconds • <i>rebalanceStartThreshold</i>—Amount of imbalance in the LAG that triggers the algorithm to start rebalancing; the default is 0 percent <ul style="list-style-type: none"> • percent—Specifies that the threshold is measured as a percentage of the average load per link; in the range 0–100 • subscribers—Specifies that the threshold is measured by the number of subscribers away from the average subscriber count per link in the LAG; in the range 0–100000 • <i>rebalanceStopThreshold</i>—Amount of imbalance in the LAG that triggers the algorithm to stop rebalancing; the default is 0 percent <ul style="list-style-type: none"> • percent—Specifies that the amount of imbalance is measured as a percentage of the average load per link; in the range 0–100 • subscribers—Specifies that the threshold is measured by the number of subscribers away from the average subscriber count link in the LAG; in the range 0–100000 • <i>rebalanceMaximumImprovement</i>—Maximum number of links in the LAG to rebalance; the default is 100 percent <ul style="list-style-type: none"> • percent—Specifies that the maximum number of links is measured as a percentage of the total links; in the range 0–100 • subscribers—Specifies that the maximum number of links is measured by the number of subscribers; in the range 0–100000
Mode	Interface Configuration
Related Documentation	<ul style="list-style-type: none"> • Configuring Load Rebalancing for 802.3ad Link Aggregation Groups on page 71

member-interface

Syntax [no] member-interface *interfaceType interfaceSpecifier*

Release Information Command introduced before JunosE Release 7.1.0.

Description Adds a member interface, also known as a bundle member, to an MLPPP bundle, an MLFR bundle, or an IEEE 802.3ad link aggregation group (LAG) bundle. The **no** version removes the specified member interface from the bundle.

- Options**
- *interfaceType*—One of the following interface types listed in Interface Types and Specifiers
 - serial (MLFR bundle or MLPPP bundle)
 - pos (MLFR bundle only)
 - fastEthernet (IEEE 802.3ad LAG bundle only)
 - gigabitEthernet (IEEE 802.3ad LAG bundle only)
 - *interfaceSpecifier*—Particular interface; format varies according to interface type; see Interface Types and Specifiers

Mode Interface Configuration, Subinterface Configuration

node

Syntax	[no] { <i>typeOfInterface</i> set superset } node [group <i>trafficClassGroup</i> scheduler-profile <i>schedulerProfileName</i>]*
Release Information	Command introduced before JunosE Release 7.1.0. svlan keyword added in JunosE Release 7.1.0. set and superset keywords added in JunosE Release 9.2.0.
Description	Specifies that a scheduler node be configured for each interface of the given interface type. The no version removes this rule from the QoS profile.
Options	<ul style="list-style-type: none"> • <i>typeOfInterface</i>—Interface types for scheduler nodes to be configured: atm, atm-vc, atm-vp, bridge, ethernet, fr-vc, ip, ip-tunnel, ipv6, l2tp-session, l2tp-tunnel, lsp, serial, server-port, svlan, vlan • set—Configures the node for an interface set • superset—Configures the node for an interface superset • <i>trafficClassGroup</i>—Name of the traffic class group • <i>schedulerProfileName</i>—Name of the scheduler profile • *—Indicates that one or more parameters can be repeated multiple times in a list in the command line
Mode	QoS Profile Configuration
Related Documentation	<ul style="list-style-type: none"> • Configuring a QoS Profile on page 61 • Configuring Shadow Nodes • Configuring a Basic Parameter Definition for QoS Administrators on page 75 • Attaching a QoS Profile to an Interface Superset or an Interface Set

qos-parameter

Syntax In Global Configuration, Interface Configuration, QoS Interface Set Configuration, and QoS Interface Superset Configuration modes:

```
qos-parameter qosParameterInstanceName qosParameterInstanceValue
no qos-parameter
```

In Profile Configuration mode:

```
qos-parameter qosParameterInstanceName [ qosParameterInstanceValue |
add qosParameterAddValue [ initial-value qosParameterInitialValue ] ]
no qos-parameter qosParameterInstanceName [ add ]
```

Release Information Command introduced in JunosE Release 7.1.0.
add and **initial-value** keywords added in JunosE Release 7.2.0.
 Profile Configuration mode added in JunosE Release 7.2.0.
 QoS Interface Set Configuration and QoS Interface Superset Configuration modes added in JunosE Release 9.2.0.

Description In Global Configuration mode, creates a QoS parameter instance and assigns a value to the parameter. A global parameter instance is typically used to provide a global default for a parameter value. The **no** version deletes the parameter instance.

In Interface Configuration mode, creates a parameter instance, assigns a value to the parameter, and attaches the parameter instance to the interface. The **no** version detaches the parameter instance from the interface.

In Profile Configuration mode, creates a parameter instance command in a profile for use with Service Manager. When the service is activated, the parameter instances are created for the subscriber interface. The **no** version removes the parameter instance command from the profile.

In QoS Interface Set Configuration mode, creates a parameter instance, assigns a value to the parameter, and attaches the parameter instance to a QoS interface set. The **no** version detaches the parameter instance from the interface set.

In QoS Interface Superset Configuration mode, creates a parameter instance, assigns a value to the parameter, and attaches the parameter instance to a QoS interface superset. The **no** version detaches the parameter instance from the interface superset.

- Options**
- *qosParameterInstanceName*—Name of the QoS parameter instance
 - *qosParameterInstanceValue*—Number of the rate for the parameter instance; the default value is the minimum value defined in the parameter definition

- *qosParameterAddValue*—Number of the rate that is added to an existing parameter instance
- *qosParameterInitialValue*—Number of the initial rate of a newly created parameter instance

Mode Global Configuration, Interface Configuration, Profile Configuration, QoS Interface Set Configuration, QoS Interface Superset Configuration

- Related Documentation**
- [Creating Parameter Instances on page 82](#)
 - Creating a QoS Parameter on an Interface Superset or Interface Set

qos-parameter-define

Syntax	[no] qos-parameter-define <i>qosParameterDefinitionName</i> [application <i>applicationName</i>] [hierarchical]
Release Information	Command introduced in JunosE Release 7.1.0.
Description	Specifies a QoS parameter name and accesses QoS Parameter Definition mode. The no version deletes the QoS parameter definition.
Options	<ul style="list-style-type: none"> • <i>qosParameterDefinitionName</i>—Name of the parameter definition • <i>applicationName</i>—Name of the application that you want to associate with the parameter definition: <ul style="list-style-type: none"> • ip-multicast—Specifies the IP multicast bandwidth adjustment application. You must also specify the hierarchical keyword when you specify this application. • qos-byte-adjustment—Specifies the cell byte-adjustment application, which enables you to adjust the shaping rate to account for different layer 2 encapsulations in ADSL configurations. If you have configured the QoS shaping mode as cell, the system adjusts the shaping rate to account for the ATM cell pad, header, and trailer. • qos-frame-byte-adjustment—Specifies the frame byte-adjustment application, which enables you to shape traffic based on frames for VDSL configurations. If you have configured the QoS shaping mode as frame, the system adjusts the shaping rate based on bytes within frames. • qos-cell-mode—Specifies the QoS cell mode application, which enables you to configure the operational shaping mode (frame or cell) for ATM, Gigabit Ethernet, or 10-Gigabit Ethernet interfaces. • qos-downstream-rate—Specifies the QoS downstream rate application, which enables you to adjust the downstream rate of VLANs and ATM VCs based on parameter instances that are created dynamically by ANCP or AAA. The values of the parameter instances track the bandwidth of the local loop that are communicated by ANCP. • hierarchical—Specifies that the parameter instance is hierarchical. Hierarchical parameters have explicit instances that are associated with the logical interfaces of instance-interface types, as well as implicit instances that are associated with the logical interfaces of controlled-interface types. The system computes the values of an implicit instance as the sum of the values of the explicit instances stacked above the implicit instance.
Mode	Global Configuration
Related Documentation	<ul style="list-style-type: none"> • Configuring a Basic Parameter Definition for QoS Administrators on page 75 • Configuring a Parameter Definition to Calculate Hierarchical Instances on page 77

- [Configuring a Parameter Definition for IP Multicast Bandwidth Adjustment on page 77](#)
- [Configuring a Parameter Definition to Shape Ethernet Traffic Using Cell Mode on page 79](#)
- [Configuring a Parameter Definition to Adjust Cell Shaping Rates for ADSL Traffic](#)
- [Configuring a Parameter Definition to Adjust Frame Shaping Rates for VDSL Traffic](#)
- [Configuring a Parameter Definition for QoS Downstream Rate on page 80](#)

qos-port-type-profile

Syntax	<code>qos-port-type-profile <i>typeOfInterface</i> qos-profile <i>qosProfileName</i></code>
Release Information	Command introduced before JunosE Release 7.1.0. lag keyword added in JunosE Release 8.1.0.
Description	Associates a QoS profile with all the ports of a given interface type. There is no no version.
Options	<ul style="list-style-type: none"> • <i>typeOfInterface</i>—One of the following interface types to be associated with the QoS port-type profile; atm, ethernet, lag, serial, server-port • <i>qosProfileName</i>—Name of the QoS profile
Mode	Global Configuration
Related Documentation	<ul style="list-style-type: none"> • Attaching a QoS Profile to an Interface on page 62 • Enabling Default Subscriber Load Balancing for 802.3ad Link Aggregation Groups on page 70

qos-profile

Syntax [no] qos-profile *qosProfileName*

Release Information Command introduced before JunosE Release 7.1.0.
 Profile Configuration mode added in JunosE Release 7.2.0.
 QoS Interface Set Configuration and QoS Interface Superset Configuration modes added in JunosE Release 9.2.0.

Description In Global Configuration mode, creates a QoS profile on the router and enters QoS Profile Configuration mode. The **no** version deletes the QoS profile.

In Interface Configuration mode, attaches a QoS profile to an interface. The **no** version detaches the QoS profile from the interface.

In Profile Configuration mode, adds a QoS profile command for use with Service Manager. When the service is activated, the QoS profile is created and attached to the subscriber interface. The **no** version removes the QoS profile from the profile.

In QoS Interface Set Configuration mode, attaches a QoS profile to the QoS interface set. The **no** version detaches the QoS profile from the interface set.

In QoS Interface Superset Configuration mode, attaches a QoS profile to the QoS interface superset. The **no** version detaches the QoS profile from the interface superset.

Options • *qosProfileName*—Name of the QoS profile

Mode Global Configuration, Interface Configuration, Profile Configuration, QoS Interface Set Configuration, QoS Interface Superset Configuration

Related Documentation

- [Configuring a QoS Profile on page 61](#)
- [Attaching a QoS Profile to an Interface on page 62](#)
- [Configuring Shadow Nodes](#)
- [Configuring a Basic Parameter Definition for QoS Administrators on page 75](#)
- [Creating Parameter Instances on page 82](#)
- [Attaching a QoS Profile to an Interface Superset or an Interface Set](#)
- [Creating a QoS Parameter on an Interface Superset or Interface Set](#)

qos-shaping-mode

Syntax [no] qos-shaping-mode [cell | frame]

Release Information Command introduced before JunosE Release 7.1.0.

Description Specifies either cell-based or frame-based traffic shaping for ATM, Gigabit Ethernet, or 10-Gigabit Ethernet interfaces. The shaping mode is configured for a major interface and affects scheduling for all nodes and queues stacked above the interface. In cell shaping mode, queues and nodes are scheduled as if they were ATM cells. All newly configured ports use the shaping mode from port 0; frame is the default shaping mode for port 0. If you do not specify an option, the command restores the default, frame. The **no** version restores the default, frame.

- Options**
- cell—Shapes traffic based on the number of bytes in a cell, and accounts for ATM cell encapsulation and padding overhead
 - frame—Shapes traffic based on the number of bytes in a frame, without considering cell encapsulation and padding overhead; the default shaping mode for port 0

Mode Interface Configuration

- Related Documentation**
- Configuring the QoS Shaping Mode for ATM Interfaces
 - Configuring the QoS Shaping Mode for Ethernet Interfaces

queue

Syntax	[no] { <i>typeOfInterface</i> set superset } queue traffic-class <i>trafficClassName</i> [queue-profile <i>queueProfileName</i> [scheduler-profile <i>schedProfileName</i>] scheduler-profile <i>schedProfileName</i> [queue-profile <i>queueProfileName</i>]] [drop-profile <i>dropProfileName</i>] [statistics-profile <i>statisticsProfileName</i>]
Release Information	Command introduced before JunosE Release 7.1.0. svlan keyword added in JunosE Release 7.1.0. set and superset keywords added in JunosE Release 9.2.0.
Description	Specifies that a queue traffic class be configured for the selected interface type. The no version removes this rule from the QoS profile.
Options	<ul style="list-style-type: none"> • <i>typeOfInterface</i>—Interface type for queue traffic classes to be configured; atm, atm-vp, atm-vc, bridge, ethernet, fr-vc, ip, ipv6, ip-tunnel, l2tp-session, l2tp-tunnel, lsp, serial, server-port, svlan, vlan • set—Configures a queue for an interface set • superset—Configures a queue for an interface superset • <i>trafficClassName</i>—Name of the traffic class • <i>queueProfileName</i>—Name of the queue profile • <i>schedProfileName</i>—Name of the scheduler profile • <i>dropProfileName</i>—Name of the drop profile • <i>statisticsProfileName</i>—Name of the statistics profile
Mode	QoS Profile Configuration
Related Documentation	<ul style="list-style-type: none"> • Configuring a QoS Profile on page 61 • Configuring Shadow Nodes • Configuring a Basic Parameter Definition for QoS Administrators on page 75 • Configuring Rate Statistics on page 85 • Configuring Event Statistics on page 86 • Attaching a QoS Profile to an Interface Superset or an Interface Set

queue-profile

Syntax [no] queue-profile *queueProfileName*

Release Information Command introduced before JunosE Release 7.1.0.

Description Configures a queue profile. The **no** version removes the named queue profile.

Options • *queueProfileName*—Name of the queue profile

Mode Global Configuration

Related Documentation • [Configuring Queue Profiles to Manage Buffers and Thresholds on page 53](#)

scheduler-profile

Syntax [no] scheduler-profile *schedulerProfileName*

Release Information Command introduced before JunosE Release 7.1.0.

Description Configures a scheduler profile. The router supports up to 1000 scheduler profiles. The **no** version deletes the scheduler profile.

Options • *schedulerProfileName*—Name of the scheduler profile

Mode Global Configuration

Related Documentation

- [Configuring a Scheduler Hierarchy on page 57](#)
- [Configuring a Scheduler Profile for a Scheduler Node or Queue](#)
- [Configuring a Basic Parameter Definition for QoS Administrators on page 75](#)

shaping-rate

Syntax `shaping-rate shapingRate [operator operandValue]* [bps | kbps]`
`[burst burstSize [milliseconds | bytes]]`

`shaping-rate operandValue [operator operandValue]* [bps | kbps]`

`no shaping-rate`

Release Information Command introduced before JunosE Release 7.1.0.
milliseconds and **bytes** keywords added in JunosE Release 7.1.0.
bps and **kbps** keywords added in JunosE Release 8.0.0.

Description Sets the shaping rate and burst size. The **no** version deletes the shaping rate.

- Options**
- *shapingRate*—Constant shaping rate in bits per second or kilobits per second; in the range 1–10000000000 bps/kbps. You can set the shaping rate to vary from 1 bps to 1000 gbps (which is denoted by entering 10000000000 kbps in the CLI for this command).
 - *operator*—Mathematical function
 - *operandValue*—Input for the operator; can be a QoS parameter definition name or an integer
 - *—Indicates that one or more parameters can be repeated multiple times in a list in the command line
 - bps—Specifies shaping rate in bits per second
 - kbps—Specifies shaping rate in kilobits per second



NOTE: The lower and higher limits for the shaping rate range apply to both the **bps** and **kbps** keywords available with this command. For example, if you want to set the shaping rate to 1 mbps, you can either enter 1000 as the value for the *shapingRate* argument and suffix it with the **kbps** keyword, or enter 1000000 as the value for the *shapingRate* argument and suffix it with the **bps** keyword. Both the methods of configurations result in the same shaping rate value to be set.

- *burstSize*—Number, in the range 0–522240 (0–510 KB); 0 enables the router to select an applicable default value
- milliseconds—Specifies burst size in milliseconds
- bytes—Specifies burst size in bytes

Mode Scheduler Profile Configuration

**Related
Documentation**

- Configuring Rate Shaping for a Scheduler Node or Queue
- Configuring Port Shaping
- [Configuring a Basic Parameter Definition for QoS Administrators on page 75](#)

shared-shaping-constituent

Syntax shared-shaping-constituent [priority [*priorityValue*] | weight [*weightValue*]]
no shared-shaping-constituent

Release Information Command introduced before JunosE Release 7.1.0.

Description Sets the attributes of implicit and explicit shared-shaping constituents and specifies explicit shared-shaping constituents. Constituents default to priority with a priority value of 8. Priority constituents are ordered before weighted constituents. The **no** version deletes the attributes or explicit constituent.

- Options**
- *priorityValue*—Value, in the range 1–8, that specifies the order in which the constituent can claim bandwidth from among all priority constituents; a lower value correlates to a higher claim; 8 is the default value
 - *weightValue*—Value, in the range 1–31, that specifies the order in which the constituent can claim bandwidth from among all weighted constituents; a lower value correlates to a higher claim; 8 is the default value

Mode Scheduler Profile Configuration

- Related Documentation**
- Configuring Implicit Constituents for Simple or Compound Shared Shaping
 - Configuring Explicit Constituents for Simple or Compound Shared Shaping

shared-shaping-rate

Syntax `shared-shaping-rate sharedShapingRate [operator operandValue]* [bps | kbps]`
`[burst burstSize [milliseconds | bytes]] { simple | compound | auto }`
`[explicit-constituents]`

`shared-shaping-rate operandValue [operator operandValue]* [bps | kbps]`
`[burst burstSize [milliseconds | bytes]] { simple | compound | auto }`
`[explicit-constituents]`

`no shared-shaping-rate`

Release Information Command introduced before JunosE Release 7.1.0.
operator and *operandValue* variables added in JunosE Release 7.1.0.
milliseconds and **bytes** keywords added in JunosE Release 7.1.0.
bps and **kbps** keywords added in JunosE Release 8.0.0.

Description Sets the shared-shaping rate and burst size for the logical interface. This command must appear in the scheduler profile for either the best-effort queue or the best-effort scheduler node. The **no** version deletes the shared-shaping rate.

- Options**
- *sharedShapingRate*—Constant shared-shaping rate in bits per second or kilobits per second; in the range 1–1000000000 bps/kbps. You can set the shaping rate to vary from 1 bps to 1000 gbps (which is denoted by entering 1000000000 kbps in the CLI for this command).
 - *operator*—Mathematical function
 - *operandValue*—Input for the operator; can be a QoS parameter definition name or an integer
 - *—Indicates that one or more parameters can be repeated multiple times in a list in the command line
 - bps—Specifies shared-shaping rate in bits per second
 - kbps—Specifies shared-shaping rate in kilobits per second



NOTE: The lower and higher limits for the shaping rate range apply to both the **bps** and **kbps** keywords available with this command. For example, if you want to set the shaping rate to 1 mbps, you can either enter 1000 as the value for the *shapingRate* argument and suffix it with the **kbps** keyword, or enter 1000000 as the value for the *shapingRate* argument and suffix it with the **bps** keyword. Both the methods of configurations result in the same shaping rate value to be set.

- *burstSize*—Number, in the range 0–522240 (0–510 KB); 0 enables the router to select an applicable default value
- milliseconds—Specifies burst size in milliseconds

- bytes—Specifies burst size in bytes
- simple—Specifies the simple form of shared shaping, which does not manage voice and video traffic, but shapes data queue rates to the value of the shared rate minus the combined voice and video traffic rate
- auto—Specifies that the router automatically selects the type of shared shaping depending on the module; compound is selected only for line modules that support it, and simple is selected for all other line modules; this is the default mode
- compound—Specifies the compound form of shared shaping, which actively shapes voice and video traffic so that the shared rate cannot be exceeded, and shapes data queue rates to the value of the shared rate minus the combined voice and video traffic rate; requires special hardware
- explicit-constituents—Overrides automatic selection of compound shared-shaping constituents and enables you to explicitly specify constituents and bandwidth allocation; generates an error message and has no effect when applied to modules that do not support compound shared shaping

Mode Scheduler Profile Configuration

- Related Documentation**
- Configuring Simple Shared Shaping
 - Configuring Implicit Constituents for Simple or Compound Shared Shaping
 - Configuring Explicit Constituents for Simple or Compound Shared Shaping
 - [Configuring a Basic Parameter Definition for QoS Administrators on page 75](#)

traffic-class

Syntax In Classifier Group Configuration mode:

[no] [suspend] traffic-class *trafficClassName1*

In Global Configuration and Traffic Class Group Configuration modes:

[no] traffic-class *trafficClassName2*

Release Information Command introduced before JunosE Release 7.1.0.

Description In Classifier Group Configuration mode, specifies a traffic class in a policy list for policy management. The **no** version removes a traffic class from a policy list; the **suspend** version temporarily suspends the policy rule; the **no suspend** version resumes application of a suspended rule.



NOTE: This command replaces the Policy List Configuration version of the **traffic-class** command, which may be removed completely in a future release.

In Global Configuration mode, configures a traffic class in the E Series router. In Traffic Class Group Configuration mode, specifies a traffic class that belongs to the traffic-class group. The **no** version deletes the traffic class.

- Options**
- *trafficClassName1*—Name of the traffic class; up to 40 characters
 - *trafficClassName2*—Name of the traffic class; up to 31 characters

Mode Classifier Group Configuration, Global Configuration, Traffic Class Group Configuration

- Related Documentation**
- [Configuring Traffic Classes That Define Service Levels on page 51](#)
 - [Configuring Traffic-Class Groups That Define Service Levels on page 52](#)
 - Policy Rule Precedence

traffic-class-group

Syntax	[no] traffic-class-group <i>trafficClassGroupName</i> [slot <i>slotNumber</i> auto-strict-priority extended]
Release Information	Command introduced before JunosE Release 7.1.0.
Description	Configures a traffic-class group. The no version deletes the selected traffic-class group. You must remove all local (slot-based) instances of a traffic-class group before you can remove the global group.
Options	<ul style="list-style-type: none"> • <i>trafficClassGroupName</i>—Name of the traffic class group; up to 31 characters • <i>slotNumber</i>—Number of the slot associated with the group, in the range 0–17 • auto-strict-priority—Specifies strict-priority scheduling for the group, regardless of whether the scheduler profile associated with the group node specifies strict-priority scheduling. Only one auto-strict-priority group can exist; this is the default behavior for a group. • extended—Specifies that strict-priority scheduling for the group is determined by the scheduler profile associated with the group node; scheduling is either hierarchical round-robin or strict priority, but if a strict-priority traffic-class group already exists, this group must be scheduled via HRR
Mode	Global Configuration
Related Documentation	<ul style="list-style-type: none"> • Configuring Traffic Classes That Define Service Levels on page 51 • Configuring Traffic-Class Groups That Define Service Levels on page 52 • Configuring QoS for an L2TP Session

PART 3

Administration

- [Monitoring QoS Statistics on page 159](#)
- [Monitoring QoS on Interfaces on page 161](#)
- [Monitoring Traffic Classification and Queuing on page 165](#)
- [Monitoring Scheduling on page 171](#)
- [Monitoring QoS Parameters on page 185](#)
- [Statistics-Gathering Commands on page 189](#)
- [Monitoring Commands on page 199](#)

Monitoring QoS Statistics

- [Monitoring Queue Statistics for the Fabric on page 159](#)
- [Clearing QoS Statistics on the Egress Queue on page 160](#)
- [Clearing QoS Statistics on the Fabric Queue on page 160](#)

Monitoring Queue Statistics for the Fabric

Purpose Display forwarded and dropped statistics for the fabric.

Action To display general information about the fabric queue:

```
host1#show fabric-queue
```

traffic class	egress slot	type	forwarded packets	forwarded bytes	dropped packets	dropped bytes
best-effort	all	committed	0	0	0	0
best-effort	all	conformed	0	0	0	0
best-effort	all	exceeded	0	0	0	0

To display detailed information about the fabric queue in a specific traffic class:

```
host1#show fabric-queue traffic-class video detail
```

To display information about the fabric queue on the egress slot:

```
host1#show fabric-queue egress-slot 0
```

Meaning [Table 13 on page 159](#) lists the **show fabric-queue** command output fields.

Table 13: show fabric-queue Output Fields

Field Name	Field Description
traffic class	Name of the traffic class
egress slot	Egress slot for which statistics are being displayed
type	Type of packet
forwarded packets	Number of forwarded packet
forwarded bytes	Number of forwarded bytes

Table 13: show fabric-queue Output Fields (*continued*)

Field Name	Field Description
dropped packets	Number of dropped packets
dropped bytes	Number of dropped bytes

- Related Documentation**
- [Configuring Rate Statistics on page 85](#)
 - [Configuring Event Statistics on page 86](#)
 - [show fabric-queue](#)

Clearing QoS Statistics on the Egress Queue

To clear statistics from the egress queue for the specified interface and traffic class:

- Issue the **clear egress-queue** command.

```
host1#clear egress-queue atm 3/0 explicit traffic-class class15
```

Use the **explicit** keyword to clear queues only on the specified interface and not queues stacked above the interface.

- Related Documentation**
- [Monitoring QoS Statistics for Rates and Events](#)
 - [clear egress-queue on page 190](#)

Clearing QoS Statistics on the Fabric Queue

To clear statistics from the fabric queue for the specified traffic class and egress slot:

- Issue the **clear fabric-queue** command.

```
host1#clear fabric-queue traffic-class class15 egress-slot 3
```

By default, statistics for all traffic classes and all slots are cleared.

- Related Documentation**
- [Monitoring QoS Statistics for Rates and Events](#)
 - [clear fabric-queue on page 191](#)

Monitoring QoS on Interfaces

- [Monitoring the QoS Configuration of Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet Interfaces on page 161](#)
- [Monitoring the QoS Configuration of IEEE 802.3ad Link Aggregation Group Bundles on page 163](#)

Monitoring the QoS Configuration of Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet Interfaces

Purpose Display information about the QoS configuration for a specific Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet interface.

Action To display the QoS configuration for a Fast Ethernet interface:

```
host1#show interfaces fastEthernet 6/0
GigEthernet6/0 is Up, Administrative status is Up
Hardware is Intel 21440, address is 0090.1a40.5508
MAU is 100BASE-TX
MTU: Operational 1522, Administrative 1522
Duplex Mode: Operational Full Duplex, Administrative Auto Negotiate
Speed: Operational 100 Mbps, Administrative Auto Negotiate
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
In: Bytes 0, Unicast 0
Multicast 0, Broadcast 0
Errors 0, Discards 0, Mac Errors 0, Alignment 0
CRC 0, Too Longs 0, Symbol Errors 0
Out: Bytes 64, Unicast 0
Multicast 0, Broadcast 1
Errors 0, Discards 0, Mac Errors 0, Deferred 0, No Carrier 0
Collisions: Single 0, Multiple 0, Late 0, Excessive 0
Administrative qos-shaping-mode: cell
Operational qos-shaping-mode: cell
Attached QoS profile: ss
```

To display the QoS configuration for a Gigabit Ethernet interface:

```
host1#show interfaces gigabitEthernet 2/0
```

To display the QoS configuration for a 10-Gigabit Ethernet interface:

```
host1#show interfaces tenGigabitEthernet 5/0/0
```

Meaning [Table 14 on page 162](#) lists the related **show interfaces** command output fields.

Table 14: show interfaces Output Fields

Field Name	Field Description
Administrative qos-shaping-mode	<p>Configured shaping mode for the interface:</p> <ul style="list-style-type: none"> disabled—Shaping mode is configured but disabled. frame—Default shaping mode for shaping and policing rates. Reports QoS statistics such as transmitted bytes and dropped bytes based on bytes within frames. cell—Shaping mode for shaping and policing rates is cell-based; resulting traffic stream conforms exactly to the policing rates configured in downstream devices. Reports statistics in bytes within cells and also accounts for cell encapsulation and padding overhead. none—Shaping mode is not configured.
Operational qos-shaping-mode	<p>Actual shaping mode for the interface. The router determines the operational shaping mode based on the value configured using the qos-shaping-mode command. For more information, see “QoS Shaping Mode for Ethernet Interfaces Overview” on page 37.</p> <ul style="list-style-type: none"> disabled—Shaping mode is configured but disabled. frame—Default shaping mode for shaping and policing rates. Reports QoS statistics such as transmitted bytes and dropped bytes based on bytes within frames. cell—Shaping mode for shaping and policing rates is cell-based; resulting traffic stream conforms exactly to the policing rates configured in downstream devices. Reports statistics in bytes within cells and also accounts for cell encapsulation and padding overhead. none—Shaping mode is not configured.
Attached QoS profile	<p>QoS profile attachment at or below the displayed interface. For example, if the interface being displayed is a VLAN subinterface, and the attachment is at the Gigabit Ethernet interface, the Gigabit Ethernet attachment is displayed.</p>

Related Documentation

- [Configuring the QoS Shaping Mode for Ethernet Interfaces](#)
- [Creating Parameter Instances on page 82](#)
- [Monitoring the Status of Fast Ethernet Interfaces](#)
- [show interfaces on page 202](#)

Monitoring the QoS Configuration of IEEE 802.3ad Link Aggregation Group Bundles

Purpose Display information about the QoS configuration for Ethernet member links in all IEEE 802.3ad link aggregation group (LAG) bundles configured on the router, or about the member links in a specified IEEE 802.3ad LAG bundle.

Action To display the QoS configuration for a specific LAG bundle:

```
host1#show interfaces lag lg0 members
Lag lg0 is Up, Administrative status is Up
MAC Address is 0090.1a40.01be
MTU: Operational 1526
Duplex Mode: Operational Full Duplex
Speed: Operational 100 Mbps
System Priority 32768 System MAC Address is 0090.1a00.00e0 key 8
Partner System Priority 0 System MAC Address is 0000.0000.0000 key 0
QoS parameter: vlan 1500000
Attached QoS profile: eth1
Member-interface FastEthernet11/2 is Up
(LACP disabled, state collecting/distributing)
Member-interface FastEthernet11/3 is Down
(LACP disabled, state waiting)
Member-interface FastEthernet11/4 is Up
(LACP disabled, state collecting/distributing)
```

Meaning [Table 15 on page 163](#) lists the related **show interfaces lag members** command output fields.

Table 15: show interfaces lag members Output Fields

Field Name	Field Description
Lag	Name of the LAG bundle
QoS parameter	QoS parameter instance at the displayed interface
Attached QoS profile	QoS profile attachment at the displayed interface

- Related Documentation**
- [Configuring the Scheduler Hierarchy for Hashed Load Balancing in 802.3ad Link Aggregation Groups on page 70](#)
 - [Configuring the Scheduler Hierarchy for Subscriber Load Balancing in 802.3ad Link Aggregation Groups on page 71](#)
 - [Creating Parameter Instances on page 82](#)
 - *JunosE Physical Layer Configuration Guide*
 - **show interfaces lag members**

Monitoring Traffic Classification and Queuing

- [Monitoring Service Levels with Traffic Classes on page 165](#)
- [Monitoring Service Levels with Traffic-Class Groups on page 166](#)
- [Monitoring Queues and Buffers on page 167](#)
- [Monitoring Queue Thresholds on page 167](#)

Monitoring Service Levels with Traffic Classes

Purpose Display information about traffic classes.

Action To display information about all traffic classes:

```
host1#show traffic-class
```

traffic class	fabric weight	fabric strict priority
-----	-----	-----
best-effort	8	no
best-effort	8	no
tc1	8	no
tc2	8	no
tc3	8	no
tcs4	8	yes
tcs5	8	yes

To display the number of times that a QoS profile references the traffic class:

```
host1#show traffic-class brief
```

```
traffic-class best-effort referenced 17 times in qos-profiles
```

To display a list of QoS profiles and traffic-class groups that reference the traffic class:

```
host1#show traffic-class references
```

```
traffic-class best-effort
```

```
Referenced by QoS profiles:
```

```
atm-default
```

```
serial-default
```

```
ethernet-default
```

```
server-default
```

```
Referenced by traffic class groups:
```

```
None
```

Meaning [Table 16 on page 166](#) lists the **show traffic-class** command output fields.

Table 16: show traffic-class Output Fields

Field Name	Field Description
traffic class	Name of the traffic class
fabric weight	Weight of the queue in the fabric
fabric strict priority	Setting strict-priority queues in the fabric
Referenced by QoS profiles	QoS profiles that reference this traffic class
Referenced by traffic class groups	Traffic-class groups that reference this traffic class

Related Documentation

- [Configuring Traffic Classes That Define Service Levels on page 51](#)
- [show traffic-class on page 215](#)

Monitoring Service Levels with Traffic-Class Groups

Purpose Display the name of a traffic-class group and the classes in the group.

Action To display the traffic classes in a traffic-class group:

```
host1#show traffic-class-group
traffic-class-group assured-fwd
    traffic-class video

traffic-class-group assured-fwd slot 11
    traffic-class video
    traffic-class voice
```

To display the number of times each traffic-class group is referenced by a profile:

```
host1#show traffic-class-group brief
traffic-class-group g2 referenced 1 time in qos-profiles
traffic-class-group g3 referenced 1 time in qos-profiles
traffic-class-group g4 referenced 0 times in qos-profiles
traffic-class-group g1 referenced 0 times in qos-profiles
```

To display a list of profiles and QoS profiles that reference the traffic-class group:

```
host1#show traffic-class-group references
traffic-class-group g2
    Referenced by QoS profiles:
        profile1

traffic-class-group g3
    Referenced by QoS profiles:
        None
```

Meaning [Table 17 on page 167](#) lists the **show traffic-class-group** command output fields.

Table 17: show traffic-class-group Output Fields

Field Name	Field Description
traffic-class group	Name of the traffic-class group
traffic-class	Name of the traffic class
Referenced in qos-profiles	Number of times group is referenced by QoS profiles
Referenced by QoS profiles	QoS profiles that reference this traffic class

- Related Documentation**
- [Configuring Traffic-Class Groups That Define Service Levels on page 52](#)
 - [show traffic-class-group on page 216](#)

Monitoring Queues and Buffers

To monitor queues and buffers, see:

- [Monitoring Queue Thresholds on page 167](#)
- [Monitoring Queue Profiles](#)

Monitoring Queue Thresholds

Purpose Display the color-based thresholds for queues on each egress slot.

Showing queue thresholds by queue profile shows buffer memory information for each queue profile and, within that profile, shows the thresholds for each region.

In addition, showing queue thresholds by region organizes the buffer memory information by queue region and, within each region, shows the buffer allocations for each queue profile.

Action To display the color-based queue thresholds for each of the 2000 video queues when 8000 total queues are configured:

```
host1#show qos queue-thresholds egress-slot 9 queue-profile video
```

```
queue-profile video 2000 queues
```

region	egress memory	exceeded length	conformed length	committed length	total committed memory
0	0MB - 4MB	34944	69888	139648	279296000
1	4MB - 8MB	24448	48896	97792	195584000
2	8MB - 12MB	14080	28032	55936	111872000
3	12MB - 16MB	7040	14080	28032	56064000
4	16MB - 20MB	5248	10496	20992	41984000
5	20MB - 24MB	1280	2560	5120	10240000
6	24MB - 28MB	1152	2176	4224	8448000
7	28MB - 32MB	896	1792	3456	6912000

As shown, when all of the egress memory in use is between 0 MB and 4 MB, each video queue can queue 139,648 bytes of committed traffic. Because the default conformed fraction is 50 percent and the default exceeded fraction is 25 percent, half of the committed length, or 69,888 bytes, can be queued before conformed traffic is dropped, and one quarter of the committed length, or 34,944 bytes, can be queued before exceeded traffic is dropped. While memory fills, the video queues are given progressively smaller amounts of memory. For example, when 28 to 32 MB of buffer memory is in use, each video queue is limited to 3456 bytes. While memory fills beyond the last region, all frames are dropped except control traffic, until the queues are drained and memory usage falls back into one of the regions.

To display the router's memory management:

```
host1#show qos queue-thresholds egress-slot 9 region 0
region 0 (0MB - 4MB) oversubscription 3330%
```

queue-profile	exceeded length	conformed length	committed length	queue count	total committed memory
default	34944	69888	139648	2000	279296000
video	34944	69888	139648	2000	279296000
multicast	34944	69888	139648	2000	279296000
internet	34944	69888	139648	2000	279296000

Static and dynamic oversubscription determines that when 8000 queues are configured and 0–4 MB of egress buffer memory is in use, memory is oversubscribed by 3330 percent. If significantly fewer queues are configured, there is less oversubscription. This example illustrates static oversubscription.

Because all of the queues in Example 2 use default queue profiles, all queues have the same lengths. Each queue is allocated 139,648 bytes of committed buffer memory when operating within this region. This allocation allows active queues to burst traffic by using memory that is unused by quiescent queues. This example illustrates dynamic oversubscription, which is based on the assumption that when a large number of queues is configured, only a fraction of the queues is active at a given time. While more queues become active, memory fills and spills into another region. When this occurs, queues are given progressively smaller queue limits.

In memory regions 1 through 5, queue limits are progressively reduced. In region 6, memory is strictly partitioned among queues.

To display oversubscription in region 6:

```
host1#show qos queue-thresholds egress-slot 9 region 6
region 6 (24MB - 28MB) oversubscription 100%
```

queue-profile	exceeded length	conformed length	committed length	queue count	total committed memory
default	1152	2176	4224	2000	8448000
video	1152	2176	4224	2000	8448000
multicast	1152	2176	4224	2000	8448000
internet	1152	2176	4224	2000	8448000

Oversubscription is 100 percent. When 24–28 MB of the memory is in use, there is no oversubscription of egress buffer memory; 32 MB of the 32-MB memory is allocated. In

Example 3, each of the 8000 egress queues is given a queue of 4224 bytes, for a total of 16 MB.

If memory continues to fill into region 7, egress buffer memory is undersubscribed, allowing control traffic to flow within the router. As shown in Example 4, when operating in region 7, only 80 percent of the 32-MB memory is allocated.

To display oversubscription in region 7:

```
host1#show qos queue-thresholds egress-slot 9 region 7
region 7 (28MB - 32MB) oversubscription 80%
```

queue-profile	exceeded length	conformed length	committed length	queue count	total committed memory
default	896	1792	3456	2000	6912000
video	896	1792	3456	2000	6912000
multicast	896	1792	3456	2000	6912000
internet	896	1792	3456	2000	6912000

Region 7 has 2000 IP users, each with four queues. Each of the four queues use default queue profiles.

To display the queue thresholds in the multicast queue profile:

```
host1#show qos queue-thresholds egress-slot 9 queue-profile multicast
queue-profile multicast 2000 queues
```

region	egress memory	exceeded length	conformed length	committed length	total committed memory
0	0MB - 4MB	5120	10112	20096	40192000
1	4MB - 8MB	5120	10112	20096	40192000
2	8MB - 12MB	5120	10112	20096	40192000
3	12MB - 16MB	5120	10112	20096	40192000
4	16MB - 20MB	5120	10112	20096	40192000
5	20MB - 24MB	1280	2560	10112	20224000
6	24MB - 28MB	1152	2176	4224	8448000
7	28MB - 32MB	896	1792	3456	6912000

The multicast queue profile is configured with a committed length of 10,000 minimum and 20,000 maximum. When in regions 0–4, these queues would normally get more memory than the 20,000 byte maximum requested. In this case, the queue is limited to the maximum, and any excess memory is redistributed to other queues. Region 5 does not have enough memory to honor the 20,000-byte maximum requested.

Although a 20,000 byte maximum was requested, the router provisions memory in 128 byte blocks, rounded up or down per each request; 20,096 bytes is 157 blocks of 128 bytes.

In region 6, memory is strictly partitioned, and neither the minimum nor maximum request is honored. Instead, each multicast queue is given a fair share of the queue length so that aggressive bandwidth consumers cannot starve out moderate traffic consumers.

In region 7, memory is underprovisioned to allow queues to drain and to avoid starvation that occurs when egress buffer memory fills completely.

To display the queue thresholds for video queues:

```
host1#show qos queue-thresholds egress-slot 9 region 0
region 0 (OMB - 4MB) oversubscription 3330%
```

queue-profile	exceeded length	conformed length	committed length	queue count	total committed memory
default	33664	67328	134656	2000	269312000
video	67328	134656	269184	2000	538368000
multicast	5120	10112	20096	2000	40192000
internet	33664	67328	134656	2000	269312000

You can configure video queues with a buffer weight of 16 and Internet and multicast queues with a buffer weight of 8 to ensure that video queues get to queue twice as much traffic as Internet and multicast queues.

Meaning [Table 18 on page 170](#) lists the **show qos queue-thresholds** command output fields.

Table 18: show qos queue-thresholds Output Fields

Field Name	Field Description
queue profile	Name of the queue profile
region	Egress buffer memory region
egress memory	Amount of memory in each region
exceeded length	Amount of exceeded traffic that can be queued at this egress memory usage
conformed length	Amount of conformed traffic that can be queued at this egress memory usage
committed length	Amount of committed traffic that can be queued at this egress memory usage
total committed memory	Amount of committed memory allocated to the queue

- Related Documentation**
- [Configuring Queue Profiles to Manage Buffers and Thresholds on page 53](#)
 - [show qos queue-thresholds on page 208](#)

Monitoring Scheduling

- [Monitoring the Configuration of Scheduler Profiles on page 171](#)
- [Monitoring Shared Shapers on page 172](#)
- [Monitoring the QoS Scheduler Hierarchy on page 174](#)
- [Monitoring the Configuration of QoS Profiles on page 180](#)
- [Monitoring the QoS Profiles Attached to an Interface on page 182](#)
- [Monitoring a Scheduler Hierarchy on an Interface with QoS Profiles on page 184](#)

Monitoring the Configuration of Scheduler Profiles

Purpose Display information about scheduler profiles. If you do not specify the scheduler profile name, data for all scheduler profiles is displayed.

You can display the values that you configured using a QoS parameter for assured rate, shaping rate, and shared-shaping rate.

Action To display information about all scheduler profiles:

```
host1#show scheduler-profile
```

	shaping			strict	
scheduler	rate	burst	weight	priority	assured rate
-----	-----	-----	-----	-----	-----
default	<none>	32767	8	no	<none>
wf100	128000	32767	20	no	75000
spSV25	5000000	32767	40	no	64000
videoHar	<none>	32767	8	no	hierarchical

To display the number of times that a QoS profile references the scheduler profile:

```
host1#show scheduler-profile brief
scheduler-profile default referenced 39 times in qos-profiles
scheduler-profile wf100 referenced 1 time in qos-profiles
scheduler-profile spSV25 referenced 2 times in qos-profiles
```

To display a list of QoS profiles that reference the scheduler profile:

```
host1#show scheduler-profile references
scheduler-profile default
  Referenced by QoS profiles:
    atm-default
    serial-default
    ethernet-default
    server-default
```

```

scheduler-profile wf100
  Referenced by QoS profiles:
    ipV610

scheduler-profile spSV25
  Referenced by QoS profiles:
    qospro25

```

Meaning [Table 19 on page 172](#) lists the **show scheduler-profile** command output fields.

Table 19: show scheduler-profile Output Fields

Field Name	Field Description
scheduler	Name of the scheduler profile
shaping rate	Maximum bandwidth, in bits per second, provided to a node or queue
burst	Catch-up number associated with the shaper
weight	HRR weight of a node or queue
strict priority	Status of strict priority, yes or no
assured rate	Desired bandwidth, in bits per second, provided to a node or queue, or the keyword, hierarchical, to indicate that HAR is used
Referenced by QoS profiles	QoS profiles that reference this profile

- Related Documentation**
- [Configuring a Scheduler Hierarchy on page 57](#)
 - Configuring Simple Shared Shaping
 - Configuring Compound Shared Shaping
 - [show scheduler-profile on page 213](#)

Monitoring Shared Shapers

Purpose Display information about the configured shared shapers.

The best-effort queue is listed as the first resource for shared shapers that are queue controlled. The best-effort scheduler node is listed as the first resource for shared shapers that are node controlled.

Action To display information about configured shared shapers for a specific interface:

```
host1#show qos shared-shaper interface atm 11/0
```

interface	resource	shared shaping rate	shaping rate	other rate
-----	-----	-----	-----	-----

```

atm-vc ATM11/0.10 A atm-vc node          500000          500000
                   atm-vc queue best-effort
                   atm-vc node EF
                   A atm-vc queue EF voice    100000
                   atm-vc node AF
                   A atm-vc queue AF video    200000
atm-vc ATM11/0.11 A atm-vc node          500000          500000
                   atm-vc queue best-effort
                   atm-vc node EF
A atm-vc queue EF voice    100000
                   atm-vc node AF
                   A atm-vc queue AF video    200000

Total shared shapers:      2
Total constituents:        12
Total shared shaper failovers: 0
Compound shared shapers are not supported

```

To display information about configured shared shapers for a specific L2TP session:

```
host1#show qos shared-shaper l2tp-session session1
```

To display information about the interface at the root of the scheduler hierarchy located on the tunnel-service interface or at the same hierarchy for LNS GRE tunnel traffic:

```
host1#show qos shared-shaper tunnel-server 6/0
```

To display information about the shared shapers for an interface set:

```
host1#show qos shared-shaper qos-interface-set gigEbusiness
```

To display information about the shared shapers for an interface superset:

```
host1#show qos shared-shaper qos-interface-superset allservices
```

Meaning [Table 20 on page 173](#) lists the **show qos shared-shaper** command output fields.

Table 20: show qos shared-shaper Output Fields

Field Name	Field Description
interface	Type of interface
resource	Traffic resource associated with the logical interface
shared shaping rate	Configured shared-shaping rate in bits per second
shaping rate	Individual shaping rate of a traffic resource in bits per second
other rate	Actual current shaping rate in bits per second
Total shared shapers	Total number of shared shapers
Total constituents	Total number of resource constituents for all shared shapers

Table 20: show qos shared-shaper Output Fields (*continued*)

Field Name	Field Description
Total number of shared shapers that are disabled (in failover mode) due to lack of resources	Total number of shared shapers that are disabled (in failover mode) due to lack of resources
Compound shared shapers are [not] supported	Indication of whether compound shared shapers are supported; determined by installed hardware

- Related Documentation**
- [Configuring a Scheduler Hierarchy on page 57](#)
 - Configuring Simple Shared Shaping
 - Configuring Compound Shared Shaping
 - Configuring Interface Sets for QoS
 - Configuring Interface Supersets for QoS
 - [show qos shared-shaper on page 211](#)

Monitoring the QoS Scheduler Hierarchy

Purpose Display information about the QoS scheduler hierarchy, including interfaces, resources, and shaping rates on a particular interface. Phantom nodes are not displayed in the output for this command.

If you do not specify the **traffic-class-group** keyword, the output displays information for the default traffic-class group.

Action To display the scheduler hierarchy for a particular interface:

```
host1# show qos scheduler-hierarchy interface fastEthernet 9/0
Scheduler hierarchy for the default traffic-class group
```

interface	resource	shaping rate	shared shaping rate	assured rate or weight
ethernet Eth9/0	ethernet port			wgt 8
ethernet Eth9/0	ethernet queue			wgt 8
svlan Eth9/0 svlan 2	svlan node			wgt 8
vlan Eth9/0.1	vlan node			wgt 1
vlan Eth9/0.1	vlan queue best-effort		2000000	wgt 8
vlan Eth9/0.2	vlan node			wgt 3
vlan Eth9/0.2	vlan queue video	2000000		wgt 8
vlan Eth9/0.2	vlan queue best-effort		6000000	wgt 8
vlan Eth9/0.3	vlan node			wgt 6
vlan Eth9/0.3	vlan queue video	3000000		wgt 8
vlan Eth9/0.3	vlan queue best-effort		8000000	wgt 8

Scheduler hierarchy for traffic-class group EF

interface	resource	shaping rate	shared shaping rate	assured rate or weight
-----------	----------	--------------	---------------------	------------------------

```

-----
ethernet Eth9/0      ethernet group node EF          wgt 8
svlan Eth9/0 svlan 2  svlan node EF          wgt 8
vlan Eth9/0.2        vlan queue EF voice 100000      wgt 8
vlan Eth9/0.3        vlan queue EF voice 300000      wgt 8

```

To display the scheduler hierarchy from the specified interface down to the port, then up from the specified interface:

```
host1#show qos scheduler-hierarchy interface fastEthernet 9/0.2 level 0
```

Scheduler hierarchy for the default traffic-class group

interface	resource	shaping rate	shared shaping rate	assured rate or weight
ethernet Eth9/0	ethernet port			wgt 8
svlan Eth9/0 svlan 2	svlan node			wgt 8
vlan Eth9/0.2	vlan node			wgt 3
vlan Eth9/0.2	vlan queue video	2000000		wgt 8
vlan Eth9/0.2	vlan queue best-effort		6000000	wgt 8

Scheduler hierarchy for the default traffic-class group

interface	resource	shaping rate	shared shaping rate	assured rate or weight
ethernet Eth9/0	ethernet port			wgt 8
ethernet Eth9/0	ethernet group node EF			wgt 8
svlan Eth9/0 svlan 2	svlan node EF			wgt 8
vlan Eth9/0.2	vlan queue EF voice	100000		wgt 8

To display the QoS scheduler hierarchy for a specified interface rather than those stacked above the interface:

```
host1#show qos scheduler-hierarchy interface fastEthernet 9/0.2 explicit
```

Scheduler hierarchy for the default traffic-class group

interface	resource	shaping rate	shared shaping rate	assured rate or weight
vlan Eth9/0.2	vlan node			wgt 3
vlan Eth9/0.2	vlan queue video	2000000		wgt 8
vlan Eth9/0.2	vlan queue best-effort		6000000	wgt 8

Scheduler hierarchy for traffic-class group EF

interface	resource	shaping rate	shared shaping rate	assured rate or weight
vlan Eth9/0.2	vlan queue EF voice	100000		wgt 8

To display the scheduler hierarchy of a specific traffic-class group or the default traffic-class group:

```
host1#show qos scheduler-hierarchy interface fastEthernet 9/0
traffic-class-group EF
```

Scheduler hierarchy for traffic-class group EF

interface	resource	shaping rate	shared shaping rate	assured rate or weight
-----	-----	-----	-----	-----
ethernet Eth9/0	ethernet group node EF			wgt 8
svlan Eth9/0 svlan 2	svlan node EF			wgt 8
vlan Eth9/0.2	vlan queue EF voice	100000		wgt 8
vlan Eth9/0.3	vlan queue EF voice	300000		wgt 8

To display a summary of the scheduler profiles stacked above the specified interface:

```
host1#show qos scheduler-hierarchy interface fastEthernet 9/0 summary
```

```
Total number of nodes: 7
  Level 0 nodes:      1
  Level 1 nodes:      2
  Level 2 nodes:      4
  Level 3 nodes:      0
Total number of queues: 8
  Level 0 queues:     0
  Level 1 queues:     1
  Level 2 queues:     0
  Level 3 queues:     7
```

To display information about a specified interface in condensed format:

```
host1#show qos scheduler-hierarchy interface fastEthernet 9/0 brief
```

Scheduler hierarchy for the default traffic-class group

interface	resource
-----	-----
ethernet Eth9/0	ethernet port
ethernet Eth9/0	ethernet queue
svlan Eth9/0 svlan 2	svlan node
vlan Eth9/0.1	vlan node
vlan Eth9/0.1	vlan queue best-effort
vlan Eth9/0.2	vlan node
vlan Eth9/0.2	vlan queue video
vlan Eth9/0.2	vlan queue best-effort
vlan Eth9/0.3	vlan node
vlan Eth9/0.3	vlan queue video
vlan Eth9/0.3	vlan queue best-effort

Scheduler hierarchy for traffic-class group EF

interface	resource
-----	-----
ethernet Eth9/0	ethernet group node EF
svlan Eth9/0 svlan 2	svlan node EF
vlan Eth9/0.2	vlan queue EF voice
vlan Eth9/0.3	vlan queue EF voice

To display the scheduler level, scheduler profile that controls QoS behavior of the scheduler nodes and queues, and the burst associated with shaping rates:


```
host1#show qos scheduler-hierarchy interface fastEthernet 9/0 full | include
subscriber-best-effort
```

```
vlan Eth9/0.1      subscriber-best-effort      2000000 default
vlan Eth9/0.2      subscriber-best-effort      6000000 default
vlan Eth9/0.3      subscriber-best-effort      8000000 default
```

To display the QoS scheduler hierarchy using a filter as an alternative to using the **level** keyword:

```
host1#show qos scheduler-hierarchy interface fastEthernet 9/0 full | include
level 2
```

```
vlan Eth9/0.1      vlan node      level 2
vlan Eth9/0.2      vlan node      level 2
vlan Eth9/0.3      vlan node      level 2
svlan Eth9/0 svlan 2  svlan node EF  level 2
```

To display the QoS scheduler hierarchy for an interface set:

```
host1#show qos scheduler—hierarchy qos-interface-set vlanset1
Scheduler hierarchy for the default traffic-class group
```

interface	resource	assured		
		shared		rate
		shaping	shaping	or
rate	rate	weight		
-----	-----	-----	-----	-----
ethernet Eth1/0/0	ethernet port			wgt 8
superset cluster	superset node		800000000	wgt 8
set vlanset1	set node		300000000	wgt 8
vlan Eth1/0/0.1	vlan queue best-effort			wgt 8
vlan Eth1/0/0.2	vlan queue best-effort			wgt 8

Scheduler hierarchy for traffic-class group EF

interface	resource	assured		
		shared		rate
		shaping	shaping	or
rate	rate	weight		
-----	-----	-----	-----	-----
ethernet Eth1/0/0	ethernet port			wgt 8
ethernet Eth1/0/0	ethernet group node EF			wgt 8
superset cluster	superset node EF	100000000		wgt 8

```
set vlnaset1          set queue EF EF          wgt 8
```

Scheduler hierarchy for traffic-class group AF

interface	resource	rate	shaping	shaping	or	weight
ethernet Eth1/0/0	ethernet port					wgt 8
superset cluster	superset node AF	100000000				wgt 8
set vlnaset1	set node AF					wgt 8
vlan Eth1/0/0.1	vlan queue AF AF					wgt 8
vlan Eth1/0/0.2	vlan queue AF AF					wgt 8

To display the QoS scheduler hierarchy for an interface superset:

```
host1#show qos scheduler—hierarchy qos-interface-superset cluster
Scheduler hierarchy for the default traffic-class group
```

interface	resource	rate	shaping	shaping	or	weight
ethernet Eth1/0/0	ethernet port					wgt 8
superset cluster	superset node			800000000		wgt 8
set vlnaset1	set node			300000000		wgt 8
vlan Eth1/0/0.1	vlan queue best-effort					wgt 8
vlan Eth1/0/0.2	vlan queue best-effort					wgt 8
set vlnaset2	set node					wgt 8
vlan Eth1/0/0.3	vlan queue best-effort					wgt 8

Scheduler hierarchy for traffic-class group EF

assured

interface	resource	shared rate		
		shaping rate	shaping rate	or weight
-----	-----	-----	-----	-----
ethernet Eth1/0/0	ethernet port			wgt 8
ethernet Eth1/0/0	ethernet group node EF			wgt 8
superset cluster	superset node EF	100000000		wgt 8
set vlanet1	set queue EF EF			wgt 8
set vlanet2	set queue EF EF			wgt 8

Scheduler hierarchy for traffic-class group AF

interface	resource	assured		
		shared rate	shared rate	or weight
-----	-----	-----	-----	-----
ethernet Eth1/0/0	ethernet port			wgt 8
superset cluster	superset node AF	100000000		wgt 8
set vlanet1	set node AF			wgt 8
vlan Eth1/0/0.1	vlan queue AF AF			wgt 8
vlan Eth1/0/0.2	vlan queue AF AF			wgt 8
set vlanet2	set node AF			wgt 8
vlan Eth1/0/0.3	vlan queue AF AF			wgt 8

Meaning [Table 21 on page 179](#) lists the **show qos scheduler-hierarchy** command output fields.

Table 21: show qos scheduler-hierarchy Output Fields

Field Name	Field Description
interface	Type of interface
resource	Traffic resource associated with the logical interface
shaping rate	Individual shaping rate of a traffic resource in bits per second

Table 21: show qos scheduler-hierarchy Output Fields (*continued*)

Field Name	Field Description
shared shaping rate	Configured shared-shaping rate in bits per second
assured rate or weight	Configured assured rate in bits per second or configured weight

Related Documentation

- [Configuring a Scheduler Hierarchy on page 57](#)
- Configuring Simple Shared Shaping
- Configuring Compound Shared Shaping
- Configuring Interface Sets for QoS
- Configuring Interface Supersets for QoS
- [show qos scheduler-hierarchy on page 209](#)

Monitoring the Configuration of QoS Profiles

Purpose Display information about QoS profiles, including attachments to interfaces or port types.

This command displays groups, nodes, and queues, in that order, according to the following sequence:

- not members of a traffic-class group
- members of the strict-priority traffic-class group
- members of an extended traffic-class group in the order of configuration

Action To display information about a specific QoS profile:

```
host1# show qos-profile qpDiffServExample1
```

```
qos-profile qpDiffServExample1:
```

t-class group	interface type	rule type	traffic class	scheduler profile	queue profile	drop profile	statistics profile
	ip	queue	tc3	best-effort	default	default	default
	ip	queue	tc4	best-effort	default	default	default
	ip	queue	tc5	best-effort	default	default	default
expedited-forwarding	ethernet	group		expeditedGroup			
expedited-forwarding	ip	node		default			
expedited-forwarding	ip	queue	voice	voice	default	default	default
best-effort	ethernet	group		bestEffortGroup			
best-effort	ip	node		default			
best-effort	ip	queue	best-effort	best-effort	default	default	default
assured-forwarding	ethernet	group		assuredGroup			
assured-forwarding	ip	node		default			
assured-forwarding	ip	queue	video	video	default	default	default

To display information about the QoS profiles attached to an interface or port type:

```

host1# show qos-profile references interface fastEthernet 9/0 202
          qos profile                               attachment
-----
atm-default      (qos-port-type-profile)
serial-default   (qos-port-type-profile)
ethernet-default (qos-port-type-profile)
server-default   (qos-port-type-profile)
lag-default      (qos-port-type-profile)
subscriber-data-service  vlan FastEthernet9/0.1
subscriber-triple-play  vlan FastEthernet9/0.2
subscriber-triple-play  vlan FastEthernet9/0.3
  Port attachments:      4
  Interface attachments: 3
  DCM Profile attachments: 0
  Not attached:         0

```

To display the number of times the QoS profile is referenced by an interface or protocol profile:

```

host1#show qos-profile brief
qos-profile atm-default referenced by 1 attachment
qos-profile serial-default referenced by 1 attachment
qos-profile ethernet-default referenced by 1 attachment
qos-profile server-default referenced by 1 attachment
qos-profile lag-default referenced by 1 attachment

```

To display information about the QoS profiles attached to a specific tunnel interface, specify the interface at the root of the scheduler hierarchy located on the tunnel-service interface or at the same hierarchy for LNS GRE tunnel traffic:

```
host1#show qos-profile references tunnel-server 6/0
```

To display information about the QoS profiles attached to a specific L2TP session:

```
host1#show qos-profile references l2tp-session session1
```

To display attachments for QoS profiles only on the specified interface and not QoS profiles stacked above the interface:

```
host1#show qos-profile references interface gigabitEthernet 6/0 explicit
```

Meaning [Table 22 on page 181](#) lists the **show qos-profile** command output fields.

Table 22: show qos-profile Output Fields

Field Name	Field Description
qos-profile	Name of QoS profile
t-class group	Name of the traffic-class group associated with the interface
interface type	Type of interface
rule type	Whether the rule is a group node, scheduler node, queue, or shadow node
traffic class	Name of the traffic class associated with the interface

Table 22: show qos-profile Output Fields (*continued*)

Field Name	Field Description
scheduler profile	Name of the scheduler profile associated with the interface
queue profile	Name of the queue profile associated with the interface
drop profile	Name of the drop profile associated with the interface
statistics profile	Name of the statistics profile associated with the interface
attachment	Type of interface or port type to which the QoS profile is attached
Port attachments	Number of QoS profiles attached to port types
DCM Profile attachments	Number of QoS profiles attached to profiles for Service Manager
Interface attachments	Number of QoS profiles attached to interfaces
Not attached	Number of QoS profiles that are unattached

Related Documentation

- [Configuring a QoS Profile on page 61](#)
- [Attaching a QoS Profile to an Interface on page 62](#)
- [Creating Parameter Instances on page 82](#)
- [show qos-profile on page 207](#)

Monitoring the QoS Profiles Attached to an Interface

Purpose Display the QoS profiles in effect for and stacked above the specified interface. If no QoS profiles are attached to the interface or above the interface, the router displays the QoS profile that is in effect down the interface stack toward the port interface.

Action To display the interface hierarchy for a specific interface:

```
host1#show qos interface-hierarchy interface atm 11/0.1
```

```
attachment@ atm-vc ATM11/0.1:
```

qos profile	t-class group	interface type	rule type	traffic class	scheduler profile	queue profile
qp2@ATM11/0.1		atm-vc	node		default	default
qp2@ATM11/0.1		atm-vp	node		default	default
qp2@ATM11/0.1		atm-vc	queue	best-effort	default	default
qp2@ATM11/0.1		atm-vc	queue	tc5	default	default
qp2@ATM11/0.1		atm-vc	queue	tc6	default	default
qp2@ATM11/0.1	g1	atm	group		strictShaper	default

```

qp2@ATM11/0.1  g1      atm-vc  node      default  default
qp2@ATM11/0.1  g1      atm-vp  node      default  default
qp2@ATM11/0.1  g1      atm-vc  queue tc1 default  default
qp2@ATM11/0.1  g1      atm-vc  queue tc2 default  default
qp2@ATM11/0.1  g2      atm-vp  node      default  default
qp2@ATM11/0.1  g2      atm-vc  queue tc3 default  default
qp2@ATM11/0.1  g2      atm-vc  queue tc4 default  default

```

To display the interface hierarchy using an L2TP session:

```
host1#show qos interface-hierarchy l2tp-session session1
```

To display the interface hierarchy for a tunnel interface, specify the interface at the root of the scheduler hierarchy located on the tunnel-service interface or at the same hierarchy for LNS GRE tunnel traffic:

```
host1#show qos interface-hierarchy tunnel-server 6/0
```

To display the interface hierarchy for an interface set:

```
host1#show qos interface-hierarchy qos-interface-set gigEbusiness
```

To display the interface hierarchy for an interface superset:

```
host1#show qos interface-hierarchy qos-interface-superset allservices
```

Meaning [Table 23 on page 183](#) lists the **show qos interface-hierarchy** command output fields.

Table 23: show qos interface-hierarchy Output Fields

Field Name	Field Description
attachment@	Interface for which the hierarchy is being displayed
qos profile	Name of the QoS profile and its attachment point
t-class group	Traffic-class groups associated with the interface
interface type	Type of interface to which the profile is attached
rule type	Queue, node, group, or shadow node
traffic class	Name of the traffic class associated with the queue
scheduler profile	Scheduler profiles associated with the interface
queue profile	Queue profiles associated with the interface

Related Documentation

- [Configuring a QoS Profile on page 61](#)
- [Attaching a QoS Profile to an Interface on page 62](#)
- [Creating Parameter Instances on page 82](#)
- [Configuring Interface Sets for QoS](#)
- [Configuring Interface Supersets for QoS](#)

- [show qos interface-hierarchy on page 203](#)

Monitoring a Scheduler Hierarchy on an Interface with QoS Profiles

To monitor a scheduler hierarchy on an interface, see:

- [Monitoring the QoS Profiles Attached to an Interface on page 182](#)
- [Monitoring the Configuration of QoS Port-Type Profiles](#)
- [Monitoring the Configuration of QoS Profiles on page 180](#)
- [Monitoring the Configuration of Scheduler Profiles on page 171](#)
- [Monitoring QoS Parameter Instances on page 186](#)

Monitoring QoS Parameters

- [Monitoring QoS Parameter Definitions on page 185](#)
- [Monitoring QoS Parameter Instances on page 186](#)

Monitoring QoS Parameter Definitions

Purpose Display the QoS parameter definition settings for QoS administrators.

Action To display the settings for a specific QoS parameter definition:

```
host1#show qos-parameter-define ip-multicast
      controlled instance subscriber
parameter interface interface interface value
  name      types    types    types    range
-----
ip-multicast ip          ip, ipv6 <none> <none>
parameter
  name                        properties
-----
ip-multicast ip-multicast-adjustment, hierarchical
```

To display information about QoS parameter definitions in condensed format:

```
host1#show qos-parameter-define voice1 brief
```

To display references to all QoS parameter definitions:

```
host1#show qos-parameter-define references
```

Meaning [Table 24 on page 185](#) lists the **show qos-parameter-define** command output fields.

Table 24: show qos-parameter-define Output Fields

Field Name	Field Description
parameter name	Name of the parameter definition
controlled interface types	Types of controlled-interface types that are available for the parameter definition
instance interface types	Types of instance-interface types that are available for the parameter definition

Table 24: show qos-parameter-define Output Fields (*continued*)

Field Name	Field Description
subscriber interface types	Types of subscriber-interface types that are available for the parameter definition
value range	Range assigned to the parameter definition
properties	Applications and hierarchical settings assigned to the parameter definition

- Related Documentation**
- [Configuring a Basic Parameter Definition for QoS Administrators on page 75](#)
 - [show qos-parameter-define on page 206](#)

Monitoring QoS Parameter Instances

Purpose Display the QoS parameter instances for QoS clients.

Action To display information about the QoS parameters attached to a specific interface or port type:

```
host1#show qos-parameter max-subscriber-bw references
interface parameter name value
-----
global max-subscriber-bw 5000000
ATM11/0.1 max-subscriber-bw 6000000
```

```
Global parameter instances: 1
Parameter instances reported: 2
```

To display a list of all QoS parameters attached to all interfaces:

```
host1#show qos-parameter references
interface parameter name value
-----
global max-subscriber-bandwidth 2000000
global subscriber-weight 1
global max-subscriber-video-bandwidth 2000000
global max-100Kbps-voice-calls 1
FastEthernet9/0.2 max-subscriber-bandwidth 6000000
subscriber-weight 3
max-subscriber-video-bandwidth 2000000
max-100Kbps-voice-calls 1
FastEthernet9/0.3 max-subscriber-bandwidth 8000000
subscriber-weight 6
max-subscriber-video-bandwidth 3000000
max-100Kbps-voice-calls 3
FastEthernet9/0 svlan 1 max-subscriber-video-bandwidth 1000000
```

```
Global parameter instances: 4
Parameter instances reported: 13
```

To display the QoS profile name and attachment data for a specific interface:

```
host1#show qos-parameter references interface fastEthernet 9/0.3
```

interface	parameter name	value	instance Type
FastEthernet9/0.3	max-subscriber-bandwidth	8000000	explicit
	subscriber-weight	6	explicit
	max-subscriber-video-bandwidth	3000000	explicit
	max-100Kbps-voice-calls	3	explicit
Explicit parameter instances:		4	
Hierarchical parameter instances:		0	
IP multicast parameter instances:		0	
Parameter instances reported:		4	

To display information in expanded format, including Service Manager references:

```
host1#show qos-parameter video references full
```

interface	parameter name	value	source	service manager refs	persistence
GigabitEthernet6/0	video	50	default	none	persistent
Global parameter instances:		0			
Parameter instances reported:		1			

To display information about global parameter instance attachments in condensed format:

```
host1#show qos-parameter references global brief
```

To display information about the parameter instances attached to a specific tunnel interface, specify the interface at the root of the scheduler hierarchy located on the tunnel-service interface or at the same hierarchy for LNS GRE tunnel traffic:

```
host1#show qos-parameter references tunnel-server 6/0
```

To display information about the parameter instances attached to a specific L2TP session:

```
host1#show qos-parameter references l2tp-session session1
```

To display parameter instances only on the specified interface and not QoS parameters stacked above the interface:

```
host1#show qos-parameter references gigabitEthernet 6/0 explicit
```

Meaning [Table 25 on page 187](#) lists the **show qos-parameter** command output fields.

Table 25: show qos-parameter Output Fields

Field Name	Field Description
interface	Location of the interface to which the parameter instance is assigned; global indicates that the parameter is assigned to the chassis
parameter name	Name of the parameter instance
value	Value assigned to the parameter instance

Table 25: show qos-parameter Output Fields (*continued*)

Field Name	Field Description
source	Source of the parameter instance: <ul style="list-style-type: none"> • dcm—Parameter instance was created in a profile • radius—Parameter instance was created through RADIUS • service manager—Parameter instance was created through Service Manager • default—Parameter instance was created through the CLI or SNMP
service manager refs	Number of references of this parameter instance created through Service Manager
persistence	Status of the persistence of a parameter instance in the system: <ul style="list-style-type: none"> • persistent—Parameter instance is stored in NVS and is restored after a chassis reset • non-persistent—Parameter instance is not stored in NVS and are deleted after a chassis reset
Global parameter instances	Number of parameter instances assigned to the chassis
Parameter instances reported	Total number of parameter instances assigned
Explicit parameter instances	Total number of explicit parameter instances assigned
Hierarchical parameter instances	Total number of hierarchical parameter instances assigned
IP multicast parameter instances	Total number of parameter instances associated with the IP multicast bandwidth adjustment application

- Related Documentation**
- [Creating Parameter Instances on page 82](#)
 - [show qos-parameter on page 204](#)

CHAPTER 21

Statistics-Gathering Commands

clear egress-queue

Syntax clear egress-queue *interfaceType interfaceSpecifier* [*explicit*]
[*traffic-class trafficClassName*]

Release Information Command introduced before JunosE Release 7.1.0.

Description Clears egress queue statistics for the all queues bound to the specified interface for queues stacked at and above the interface, or only for the specified traffic class. There is no **no** version.

- Options**
- *interfaceType*—Interface type; see Interface Types and Specifiers
 - *interfaceSpecifier*—Particular interface; format varies according to interface type; see Interface Types and Specifiers
 - *explicit*—Clears queues only on the specified interface and not queues stacked above the interface
 - *trafficClassName*—Name of a traffic class for which egress queue statistics are cleared

Mode Privileged Exec

clear fabric-queue

Syntax clear fabric-queue [traffic-class *trafficClassName*] [egress-slot *egressSlot*]

Release Information Command introduced before JunosE Release 7.1.0.

Description Clears statistics for all fabric queues or for the specified traffic-class, egress-slot, or both. There is no **no** version.

- Options**
- *trafficClassName*—Name of a traffic class for which fabric-queue statistics will be cleared
 - *egressSlot*—Number of an egress slot for which fabric-queue statistics will be cleared

Mode Privileged Exec

committed-drop-threshold

Syntax	<code>committed-drop-threshold <i>committedDropThreshold</i></code> <code>no committed-drop-threshold</code>
Release Information	Command introduced before JunosE Release 7.1.0.
Description	Configures the threshold above which committed-drop-events are logged. The no version removes the threshold.
Options	<ul style="list-style-type: none">• <i>committedDropThreshold</i>—Bits per second in the range 0–1073741824
Mode	Statistics Profile Configuration
Related Documentation	<ul style="list-style-type: none">• Configuring Event Statistics on page 86

conformed-drop-threshold

Syntax	<code>conformed-drop-threshold <i>conformedDropThreshold</i></code> <code>no conformed-drop-threshold</code>
Release Information	Command introduced before JunosE Release 7.1.0. conditional , unconditional , and final keywords added in JunosE Release 7.2.0.
Description	Configures the threshold above which conformed-drop-events are logged. The no version removes the threshold.
Options	<ul style="list-style-type: none">• <i>conformedDropThreshold</i>—Bits per second in the range 0–1073741824
Mode	Statistics Profile Configuration
Related Documentation	<ul style="list-style-type: none">• Configuring Event Statistics on page 86

exceeded-drop-threshold

Syntax `exceeded-drop-threshold exceededDropThreshold`
 `no exceeded-drop-threshold`

Release Information Command introduced before JunosE Release 7.1.0.

Description Configures the threshold above which exceeded-drop-events are logged. The **no** version removes the threshold.

Options • *exceededDropThreshold*—Bits per second in the range 0–1073741824

Mode Statistics Profile Configuration

Related Documentation • [Configuring Event Statistics on page 86](#)

forwarding-rate-threshold

Syntax forwarding-rate-threshold *forwardingRateThreshold*
no forwarding-rate-threshold

Release Information Command introduced before JunosE Release 7.1.0.

Description Configures the threshold above which forwarded-rate-exceeded events are logged. The **no** version removes the threshold.

Options • *forwardingRateThreshold*—Bits per second in the range 0–1073741824

Mode Statistics Profile Configuration

Related Documentation • [Configuring Event Statistics on page 86](#)

rate-period

Syntax `rate-period ratePeriod`

`no rate-period`

Release Information Command introduced before JunosE Release 7.1.0.

Description Configures the length of time during which statistics are logged. The **no** version deletes the rate period and results in no statistics being gathered.

Options • *ratePeriod*—Number of seconds in the range 1–43200

Mode Statistics Profile Configuration

Related Documentation • [Configuring Rate Statistics on page 85](#)
• [Configuring Event Statistics on page 86](#)

statistics-profile

Syntax	[no] statistics-profile <i>statisticsProfileName</i>
Release Information	Command introduced before JunosE Release 7.1.0.
Description	Configures a statistics profile. The no version removes the named statistics profile.
Options	<ul style="list-style-type: none">• <i>statisticsProfileName</i>—Name of the statistics profile
Mode	Global Configuration
Related Documentation	<ul style="list-style-type: none">• Configuring Statistic Profiles for QoS on page 85• Configuring Rate Statistics on page 85• Configuring Event Statistics on page 86

CHAPTER 22

Monitoring Commands

show aaa qos downstream-rate

Syntax show aaa qos downstream-rate

Release Information Command introduced in JunosE Release 8.1.0.

Description Displays whether the QoS downstream rate application is enabled to use downstream rates obtained from the Actual-Data-Rate-Downstream [26-30] DSL Forum VSA.

Mode Privileged Exec

Related Documentation

- [Monitoring the AAA Downstream Rate for QoS](#)

show drop-profile

Syntax show drop-profile [*dropProfileName*] [brief | references] [*filter*]

Release Information Command introduced before JunosE Release 7.1.0.

Description Displays information about the drop profile.

- Options**
- *dropProfileName*—Name for the drop profile
 - *brief* —Displays information in a condensed format
 - *references*—Displays QoS profiles which reference the drop profile
 - *filter*—See Filtering show Commands

Mode Privileged Exec

Related Documentation • Monitoring Drop Profiles for RED and WRED

show interfaces

Syntax	show interfaces <i>interfaceType interfaceSpecifier</i> [<i>delta</i>] [<i>brief</i>] [<i>filter</i>]
Release Information	Command introduced before JunosE Release 7.1.0.
Description	Displays the current state of the interface you specify.
Options	<ul style="list-style-type: none"> • <i>interfaceType</i>—Interface type; see Interface Types and Specifiers • <i>interfaceSpecifier</i>—Particular interface; format varies according to interface type; see Interface Types and Specifiers; for ATM, subinterfaces are not supported by the syntax • <i>delta</i>—Displays baselined statistics • <i>brief</i>—Displays a brief summary of the interface • <i>filter</i>—See Filtering show Commands
Mode	Privileged Exec, User Exec
Related Documentation	<ul style="list-style-type: none"> • Monitoring the QoS Configuration of ATM Interfaces • Monitoring the QoS Configuration of Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet Interfaces on page 161 • Monitoring Interfaces and Policy Lists

show qos interface-hierarchy

Syntax To display the interface hierarchy for interfaces:

```
show qos interface-hierarchy interface interfaceType interfaceSpecifier [ atmVpi |
s-vlanIdValue ] [ filter ]
```

To display the interface hierarchy for L2TP sessions:

```
show qos interface-hierarchy l2tp session sessionName [ filter ]
```

To display the interface hierarchy for tunnel-service interfaces:

```
show qos interface-hierarchy tunnel-server interfaceSpecifier [ filter ]
```

To display the interface hierarchy for an interface set:

```
show qos interface-hierarchy qos-interface-set interfaceSetName [ filter ]
```

To display the interface hierarchy for an interface superset:

```
show qos interface-hierarchy qos-interface-superset interfaceSupersetName [ filter ]
```

Release Information Command introduced before JunosE Release 7.1.0.
atmVpi and *s-vlanIdValue* variables added in JunosE Release 7.1.0.
qos-interface-set keyword and *interfaceSetName* variable added in JunosE Release 9.2.0.
qos-interface-superset keyword and *interfaceSupersetName* variable added in JunosE Release 9.2.0.

Description Displays information about the router's QoS interface hierarchy.

- Options**
- *interfaceType*—Interface type; see Interface Types and Specifiers
 - *interfaceSpecifier*—Particular interface; format varies according to interface type; see Interface Types and Specifiers
 - *atmVpi*—Virtual path identifier of this PVC; number in the range 0–255
 - *s-vlanIdValue*—S-VLAN ID number in the range 0–4095
 - *sessionName* —Name of the L2TP session
 - *interfaceSupersetName*—Name of the interface superset
 - *interfaceSetName*—Name of the interface set
 - *filter*—See Filtering show Commands

Mode Privileged Exec, User Exec

Related Documentation

- [Monitoring the QoS Profiles Attached to an Interface on page 182](#)

show qos-parameter

Syntax To display settings for a specific parameter instance with references:

```
show qos-parameter [ qosParameterInstanceName ] references [ brief | full ] [ filter ]
```

To display references globally:

```
show qos-parameter [ qosParameterInstanceName ] references global  
[ qosParameterInstanceName ] [ brief | full ] [ filter ]
```

To display references for interfaces:

```
show qos-parameter [ qosParameterInstanceName ] references [ interface interfaceType  
interfaceSpecifier [ atmVpi | s-vlanIdValue ] ] [ explicit ] [ brief | full ]  
[ filter ]
```

To display references for L2TP sessions:

```
show qos-parameter [ qosParameterInstanceName ] references lt2p session sessionName  
[ explicit ] [ filter ]
```

Release Information Command introduced in JunosE Release 7.1.0.
full keyword added in JunosE Release 7.2.0.

Description Displays QoS parameter instance settings for QoS clients.

- Options**
- *qosParameterInstanceName*—Name of the parameter instance
 - *references*—Displays interfaces that reference this parameter instance
 - *global*—Displays information about global parameter instances
 - *interfaceType*—Interface type; see Interface Types and Specifiers
 - *interfaceSpecifier*—Particular interface; format varies according to interface type; see Interface Types and Specifiers
 - *atmVpi*—Virtual path identifier of this PVC; number in the range 0–255
 - *s-vlanIdValue*—S-VLAN ID number in the range 0–4095
 - *sessionName* —Name of the L2TP session
 - *brief*—Displays information in a condensed format
 - *full*—Displays information in expanded format
 - *explicit*—Displays parameter instances only on the specified interface and not parameter instances stacked above the interface
 - *filter*—See Filtering show Commands

Mode Privileged Exec, User Exec

Related Documentation • [Monitoring QoS Parameter Instances on page 186](#)

show qos-parameter-define

Syntax	show qos-parameter-define [<i>qosParameterDefinitionName</i>] [brief references] [<i>filter</i>]
Release Information	Command introduced in JunosE Release 7.1.0.
Description	Displays QoS parameter definition settings for QoS administrators.
Options	<ul style="list-style-type: none">• <i>qosParameterDefinitionName</i>—Name of the parameter definition• brief—Displays information in a condensed format• references—Display references to this parameter definition• <i>filter</i>—See Filtering show Commands
Mode	Privileged Exec, User Exec
Related Documentation	<ul style="list-style-type: none">• Monitoring QoS Parameter Definitions on page 185

show qos-profile

Syntax To display information about all QoS profiles or a specific QoS profile:

```
show qos-profile [ qosProfileName ] [ brief ] [ filter ]
```

To display information about the QoS profiles attached to a specific interface:

```
show qos-profile references interface interfaceType interfaceSpecifier  
[ atmVpi | s-vlanIdValue ] [ explicit ] [ brief ] [ filter ]
```

To display information about the QoS profiles attached to a specific L2TP session:

```
show qos-profile references lt2p session sessionName [ explicit ] [ brief ] [ filter ]
```

To display information about the QoS profiles attached to a specific tunnel-service interface:

```
show qos-profile references tunnel-server interfaceType [ explicit ] [ brief ] [ filter ]
```

Release Information Command introduced before JunosE Release 7.1.0.
references keyword added in JunosE Release 7.1.0.

Description Displays information about QoS profiles configured on the router. Use the **references** keyword to display information about QoS profiles attached to an interface, L2TP session, or tunnel-service interface.

- Options**
- *qosProfileName*—Name of the QoS profile
 - *references*—Displays interface profiles that reference this profile
 - *interfaceType*—Interface type; see Interface Types and Specifiers
 - *interfaceSpecifier*—Particular interface; format varies according to interface type; see Interface Types and Specifiers
 - *atmVpi*—Virtual path identifier of this PVC; number in the range 0–255
 - *s-vlanIdValue*—S-VLAN ID number in the range 0–4095
 - *sessionName* —Name of the L2TP session
 - *explicit*—Displays attachments for QoS profiles only on the specified interface and not QoS profiles stacked above the interface
 - *brief*—Displays information in a condensed format
 - *filter*—See Filtering show Commands

Mode Privileged Exec, User Exec

Related Documentation

- [Monitoring the Configuration of QoS Profiles on page 180](#)

show qos queue-thresholds

Syntax	show qos queue-thresholds egress-slot <i>egressSlot</i> [queue-profile [<i>queueProfileName</i>] region [<i>regionNumber</i>]] [<i>filter</i>]
Release Information	Command introduced before JunosE Release 7.1.0.
Description	<p>Displays the color-based thresholds for queues on each egress slot.</p> <p>Displaying queue thresholds by queue profile shows buffer memory information by queue profile, and, within that profile, shows the thresholds for each region.</p> <p>Displaying queue thresholds by region organizes the buffer memory information by queue region, and, within each region, shows the buffer allocations for each queue profile.</p>
Options	<ul style="list-style-type: none"> • queue-thresholds—Displays color-based thresholds for queues on an egress slot • <i>egressSlot</i>—Displays color-based thresholds for an egress slot • queue-profile—Displays thresholds for each region of the queue profile • <i>queueProfileName</i>—Name of the queue profile • region—Displays egress memory or buffer region oversubscription • <i>regionNumber</i>—Number identifying the egress memory or buffer region on the line module • <i>filter</i>—See Filtering show Commands
Mode	Privileged Exec, User Exec
Related Documentation	<ul style="list-style-type: none"> • Monitoring Queue Thresholds on page 167

show qos scheduler-hierarchy

Syntax To display information about the scheduler hierarchy on a specified interface:

```
show qos scheduler-hierarchy interfaceType interfaceSpecifier
[ atmVpi | s-vlanIdValue ] [ explicit | level levelNumber ]
[ traffic-class-group { trafficClassGroupName | default } ] [ brief | full | summary ]
[ filter ]
```

To display information about the scheduler hierarchy on a specified tunnel-service interface:

```
show qos scheduler-hierarchy tunnel-server interfaceType [ explicit |
level levelNumber ] [ traffic-class-group { trafficClassGroupName | default } ] [ brief | full
| summary ] [ filter ]
```

To display information about the scheduler hierarchy on a specified L2TP session:

```
show qos scheduler-hierarchy lt2p session sessionName [ explicit | level levelNumber ]
[ traffic-class-group { trafficClassGroupName | default } ] [ brief | full | summary ] [ filter
]
```

To display information about the scheduler hierarchy for an interface set:

```
show qos scheduler-hierarchy qos-interface-set interfaceSetName [ full | brief | summary
] [ filter ]
```

To display information about the scheduler hierarchy for an interface superset:

```
show qos scheduler-hierarchy qos-interface-superset interfaceSupersetName [ full | brief
| summary ] [ filter ]
```

Release Information Command introduced in JunosE Release 7.1.0.

qos-interface-set keyword and *interfaceSetName* variable added in JunosE Release 9.2.0.

qos-interface-superset keyword and *interfaceSupersetName* variable added in JunosE Release 9.2.0.

Description Displays information about the scheduler hierarchy on a specified interface, L2TP session, or tunnel-service interface.

- Options**
- *interfaceType*—Interface type; see Interface Types and Specifiers
 - *interfaceSpecifier*—Particular interface; format varies according to interface type; see Interface Types and Specifiers
 - *atmVpi*—Virtual path identifier of this PVC; number in the range 0–255
 - *s-vlanIdValue*—S-VLAN ID number in the range 0–4095
 - *explicit*—Displays scheduler profiles for the specified interface rather than those stacked above the interface

- *levelNumber*—Number of scheduler levels above specified interface to report; 0 indicates the specified interface
- *trafficClassGroupName*—Name of the traffic-class group for which to display the scheduler hierarchy
- *sessionName* —Name of the L2TP session
- *interfaceSetName*—Name of the interface set
- *interfaceSupersetName*—Name of the interface superset
- *default*—Displays the scheduler hierarchy of the default traffic-class group
- *brief*—Displays information in condensed format
- *full*—Displays information in expanded format
- *summary*—Displays summary of scheduler profiles stacked above the specified interface
- *filter*—See Filtering show Commands

Mode Privileged Exec, User Exec

Related Documentation

- [Monitoring the QoS Scheduler Hierarchy on page 174](#)

show qos shared-shaper

Syntax To display information about shared shapers for a specified interface:

```
show qos shared-shaper interface interfaceType interfaceSpecifier
[ atmVpi | s-vlanIdValue ] [ summary ] [ explicit ] [ brief | full ] [ filter ]
```

To display information about shared shapers on an L2TP session:

```
show qos shared-shaper lt2p session sessionName
[ summary ] [ explicit ] [ brief | full ] [ filter ]
```

To display information about shared shapers on a tunnel-service interface:

```
show qos shared-shaper tunnel-server interfaceType
[ summary ] [ explicit ] [ brief | full ] [ filter ]
```

To display information about shared shapers associated with an interface set:

```
show qos shared-shaper qos-interface-set interfaceSetName [ summary ] [ explicit ] [
brief | full ] [ filter ]
```

To display information about shared shapers associated with an interface superset:

```
show qos shared-shaper qos-interface-superset
interfaceSupersetName [ summary ] [ explicit ] [ brief | full ] [ filter ]
```

Release Information Command introduced before JunosE Release 7.1.0.
atmVpi and *s-vlanIdValue* variables added in JunosE Release 7.1.0.
qos-interface-set keyword and *interfaceSetName* variable added in JunosE Release 9.2.0.
qos-interface-superset keyword and *interfaceSupersetName* variable added in JunosE Release 9.2.0.

Description Displays information about shared shapers for an interface, L2TP session, or tunnel-service interface.

- Options**
- *interfaceType*—Interface type; see Interface Types and Specifiers
 - *interfaceSpecifier*—Particular interface; format varies according to interface type; see Interface Types and Specifiers
 - *atmVpi*—Virtual path identifier of this PVC; number in the range 0–255
 - *s-vlanIdValue*—S-VLAN ID number in the range 0–4095
 - *sessionName* —Name of the L2TP session
 - *interfaceSetName*—Name of the QoS interface set associated with the shared shaper
 - *interfaceSuperSetName*—Name of the QoS interface superset associated with the shared shaper
 - **summary**—Displays summary of shared shapers stacked above the specified interface

- **explicit**—Displays shared shapers for the specified interface rather than those stacked above the interface
- **brief**—Displays information in condensed format
- **full**—Displays information in expanded format
- ***filter***—See Filtering show Commands

Mode Privileged Exec, User Exec

Related Documentation

- [Monitoring Shared Shapers on page 172](#)

show scheduler-profile

Syntax show scheduler-profile [*schedulerProfileName*] [brief | references] [*filter*]

Release Information Command introduced before JunosE Release 7.1.0.

Description Displays information about scheduler profiles configured on the E Series router.

- Options**
- *schedulerProfileName*—Name of the scheduler profile
 - *brief*—Displays information in a condensed format
 - *references*—Displays QoS profiles that reference this profile
 - *filter*—See Filtering show Commands

Mode Privileged Exec, User Exec

Related Documentation

- [Monitoring the Configuration of Scheduler Profiles on page 171](#)

show statistics-profile

Syntax show statistics-profile [*statisticsProfileName*] [brief | references] [*filter*]

Release Information Command introduced before JunosE Release 7.1.0.

Description Displays information about statistics profiles configured on the E Series router.

- Options**
- *statisticsProfileName*—Name of the statistics profile
 - *brief*—Displays information in a condensed format
 - *references*—Displays QoS profiles that reference this profile
 - *filter*—See Filtering show Commands

Mode Privileged Exec, User Exec

Related Documentation

- Monitoring the Configuration of Statistics Profiles

show traffic-class

Syntax show traffic-class [*trafficClassName*] [brief | references] [*filter*]

Release Information Command introduced before JunosE Release 7.1.0.

Description Displays information about traffic class(es) configured on the E Series router.

- Options**
- *trafficClassName*—Name of the traffic class
 - brief—Displays information in a condensed format
 - references—Displays QoS profiles and traffic class groups that reference this profile
 - *filter*—See Filtering show Commands

Mode Privileged Exec, User Exec

Related Documentation

- [Monitoring Service Levels with Traffic Classes on page 165](#)

show traffic-class-group

Syntax	show traffic-class-group [<i>trafficClassGroupName</i> [slot [<i>trafficClassGroupSlotNumber</i>]]] [brief references] [<i>filter</i>]
Release Information	Command introduced before JunosE Release 7.1.0.
Description	Displays information about a traffic class group configured on the E Series router.
Options	<ul style="list-style-type: none">• <i>trafficClassGroupName</i>—Name of the traffic class group• <i>trafficClassGroupSlotNumber</i>—Number of the slot associated with the group, in the range 0–17• brief—Displays information in a condensed format• references—Displays QoS profiles and traffic class groups that reference this profile• <i>filter</i>—See Filtering show Commands
Mode	Privileged Exec, User Exec
Related Documentation	<ul style="list-style-type: none">• Monitoring Service Levels with Traffic-Class Groups on page 166

PART 4

Troubleshooting

- [Troubleshooting Egress Queues on page 219](#)

Troubleshooting Egress Queues

- Troubleshooting Memory and Processor Use for Egress Queue Rate Statistics and Events on page 219

Troubleshooting Memory and Processor Use for Egress Queue Rate Statistics and Events

Problem The E Series Broadband Services Routers uses shared processing and memory when it gathers egress queue rate statistics and events. If sufficient memory is not available, the statistics gathering is temporarily disabled and the queues are considered to be in *failover mode* until memory becomes available.

The router displays a CLI message whenever queues are put into failover mode and when they recover from failover mode.



NOTE: When an extremely large number of statistics is being gathered over a short period of time, the router might release the processor to perform more important tasks. This can result in longer rate periods than you have configured. For example, if you configured 10,000 queues to gather statistics every second on a line module, the router might actually lengthen the rate to 2 seconds or more.

Solution To display the number of queues that are disabled because of no resources, issue the **show egress-queue rates** command.

Related Documentation

- Monitoring Forwarding and Drop Rates on the Egress Queue
- show egress-queue rates

PART 5

Index

- [Index on page 223](#)

Index

Symbols

10-Gigabit Ethernet	
monitoring.....	161
10-Gigabit Ethernet forwarding ASIC (TFA).....	11, 24
802.3ad link aggregation	
configuring for QoS.....	39
link redundancy and QoS.....	39

A

AAA commands	
show aaa qos downstream-rate.....	200
aaa qos downstream-rate command.....	81
ADSL traffic	
managing cell tax.....	33
ANCP (Access Node Control Protocol)	
shaping downstream rates from.....	80
ASIC scheduler.....	3
assured rate.....	6
assured-rate command.....	75
atm commands	
atm-vp qos-parameter.....	83
atm-vp qos-profile.....	63
ATM commands	
show interfaces.....	202
ATM VP	
interface attachments.....	63
audience for QoS.....	5

B

B-RAS commands	
show aaa qos downstream-rate.....	200
best effort.....	6, 13
best-effort queue.....	6
best-effort scheduler node.....	6
buffer-weight command.....	53
byte adjustment applications	
overview.....	33

C

CDV (cell delay variation).....	6
CDVT (cell delay variation tolerance).....	6

cell delay variation tolerance. See CDVT	
cell delay variation. See CDV	
Cisco HDLC commands	
show interfaces serial.....	202
clear egress-queue command.....	160
clear fabric-queue command.....	160
clearing statistics.....	160
cOCx/STMx commands	
show interfaces.....	202
color-based thresholds.....	15
committed drop threshold.....	17
committed-drop-threshold command.....	87
committed-length command.....	53
committed-threshold command.....	55
configuration examples	
QoS parameters.....	89, 104, 115
configuring. See specific feature, product, or protocol	
conformed drop threshold.....	17
conformed-drop-threshold command.....	87
conformed-fraction command.....	54
conformed-length command.....	53
conformed-threshold command.....	55
constituents, shared-shaping.....	25
controlled-interface-type command.....	75
controlling subscriber bandwidth	
configuration example.....	89, 115
CT3 commands	
show interfaces serial.....	202
customer support.....	xv
contacting JTAC.....	xv

D

DiffServ	
networks.....	3
documentation	
comments on.....	xv
drop profile.....	16
configuring WRED.....	54
dynamic shaping of traffic.....	23

E

effective weight.....	6
egress forwarding ASIC (EFA).....	10
Ethernet	
802.3ad link aggregation and QoS.....	39
configuring for QoS.....	37, 39
overview for QoS.....	37
QoS shaping mode.....	37

Ethernet commands	
member-interface.....	138
show interfaces.....	202

Ethernet link aggregation commands	
member-interface.....	71
event statistics.....	86
exceeded drop threshold.....	17
exceeded-drop-threshold command.....	87
exceeded-fraction command.....	54
exceeded-length command.....	53
exceeded-threshold command.....	55

F

fabric-strict-priority command.....	51
fabric-weight command.....	51
Fast Ethernet	
monitoring.....	161
forwarding classes. <i>See</i> traffic classes	
forwarding rate threshold.....	17
forwarding-rate-threshold command.....	86
frame forwarding ASIC (FFA).....	10

G

Gigabit Ethernet	
monitoring.....	161
group command.....	61, 75
group node.....	6

H

HAR (hierarchical assured rate).....	6
HDLC commands, Cisco	
show interfaces serial.....	202
hierarchical assured rate. <i>See</i> HAR	
hierarchical round-robin. <i>See</i> HRR	
hierarchy, QoS scheduler.....	7, 174
HRR (hierarchical round-robin).....	6

I

instance-interface-type command.....	76
interface commands	
show interfaces.....	202
interface profile	
attachments.....	62
IP multicast bandwidth adjustment	
configuration example.....	104
configuring.....	77
overview.....	31

L

latency.....	6
link aggregation commands	
member-interface.....	138
load balancing	
configuring parameters.....	71
hashed	
configuring.....	70
overview.....	39, 41
munged QoS profiles.....	39
subscriber	
configuring.....	71
enabling default configuration.....	70
overview.....	39, 42
load-rebalance command.....	72, 73

M

manuals	
comments on.....	xv
MLFR commands	
member-interface.....	138
MLPPP commands	
member-interface.....	138
monitoring. <i>See</i> specific feature, product, or protocol	
multiple traffic-class groups.....	14
munged QoS profile	
Ethernet link aggregation.....	39

N

node command.....	58, 61, 70, 71, 75
nodes	
best-effort scheduler.....	6
group.....	6
scheduler.....	7
system resources.....	10

P

port-type profile, QoS.....	7
attachments.....	64
POS commands	
show interfaces.....	202
profile	
drop.....	16
QoS	
attachment.....	7
port-type.....	7
scheduler.....	21
statistics.....	17

profiles.....14

Q

QoS (quality of service)

- administrators of.....5
- clients of.....5
- description of.....3
- differentiated services
 - assured forwarding.....3
 - expedited forwarding.....3
- extends Diffserv.....3
- features.....4
- overview.....3
- parameters.....29
- terms.....6

QoS cell mode application

- configuring.....79

QoS commands

- buffer-weight.....123
- clear egress-queue.....190
- clear fabric-queue.....191
- committed-drop-threshold.....192
- committed-length.....124
- committed-threshold.....125
- conformed-drop-threshold.....193
- conformed-fraction.....126
- conformed-length.....127
- conformed-threshold.....128
- controlled-interface-type.....129
- drop-profile.....130
- exceeded-drop-threshold.....194
- exceeded-fraction.....131
- exceeded-length.....132
- exceeded-threshold.....133
- fabric-strict-priority.....134
- fabric-weight.....135
- forwarding-rate-threshold.....195
- instance-interface-type.....136
- load-rebalance.....137
- node.....139
- qos-parameter82, 140
- qos-parameter-define75, 142
- qos-port-type-profile64, 70, 144
- qos-profile58, 61, 64, 70, 71, 85, 86, 145
- qos-shaping-mode.....146
- queue.....147
- queue-profile.....148
- rate-period.....196
- scheduler-profile.....149

- shaping-rate.....150
- shared-shaping-constituent.....152
- shared-shaping-rate.....153
- show aaa qos downstream-rate.....200
- show drop-profile.....201
- show qos interface-hierarchy.....203
- show qos queue-thresholds.....208
- show qos shared-shaper.....211
- show qos-parameter.....204
- show qos-parameter-define.....206
- show qos-profile.....207
- show scheduler-profile.....213
- show statistics-profile.....214
- show traffic-class.....215
- show traffic-class-group.....216
- statistics-profile.....197
- traffic-class.....155
- traffic-class-group.....156

QoS downstream rate application

- configuring.....80

QoS parameters

- 802.3ad link aggregation.....39
- configuration examples.....89, 104, 115
- overview.....29
- parameter definitions
 - cell byte adjustment.....33
 - configuring.....75
 - frame byte adjustment.....34
 - IP multicast bandwidth
 - adjustment.....31, 77
 - QoS cell mode.....79
 - QoS downstream rate.....80
- parameter instances
 - configuring for an interface.....82
 - configuring globally.....82
 - using with Service Manager.....83

QoS profile

- ATM VP attachments.....63
- attaching.....62
- configuring.....61
- interface attachments.....62
- monitoring.....184
- munged.....39
- port-type attachments.....64
- S-VLAN attachments.....64
- using with Service Manager.....61

QoS shaping mode

- overview.....37
- qos-profile command.....75

queue buffers.....	167
queue command.....	58, 61, 70, 71, 75
queue length.....	167
queue profiles.....	14
color-based thresholds.....	15
configuring.....	53
monitoring.....	167
overview.....	14
queue-profile command.....	53
queues.....	7
system resources.....	10

R

random early detection. See RED	
range command.....	76
rate shaping	
QoS.....	7
rate statistics.....	17
rate-period command.....	85, 86
RED (random early detection).....	7, 16

S

S-VLAN	
interface attachments.....	64
scheduler	
hierarchy.....	7, 22
node, best-effort	
best-effort scheduler.....	6
profile.....	21
configuring.....	57
scheduler hierarchy	
monitoring.....	174
scheduler map. See QoS profile	
scheduler-profile command.....	75
shaping mode	
overview.....	37
shaping-rate command.....	75
shared shaping	
active constituents.....	25
constituents.....	25
active.....	25
limiting bandwidth.....	24
overview.....	23, 24
simple.....	25
controlling the algorithm.....	27
example, basic.....	25
example, on best-effort scheduler	
node.....	25
types, simple versus compound.....	24

shared-shaping-rate command.....	75
show commands	
show fabric-queue	159
show interfaces gigabitEthernet.....	161
show interfaces lag members.....	163
show interfaces tenGigabitEthernet.....	161
show qos-parameter.....	186
show qos-parameter-define	185
show qos-profile.....	180
show scheduler-profile	171
show traffic-class.....	165
show traffic-class-group	166
show qos commands	
show qos interface-hierarchy.....	182
show qos queue-thresholds	167
show qos scheduler-hierarchy.....	174
show qos shared-shaper	172
simple shared shaping. See shared shaping	
SNMP commands	
buffer-weight.	
no buffer-weight See Command	
Reference Guide N to Z, SNMP	
commands	
statistics profile.....	17
clearing.....	160
committed drop threshold.....	17
conformed drop threshold.....	17
event statistics.....	17
exceeded drop threshold.....	17
failover mode.....	219
forwarding rate threshold.....	17
maximum.....	17
overview.....	17
rate period.....	17, 85, 86
rate statistics.....	17
resource use.....	219
thresholds.....	86
statistics-profile command.....	85
strict-priority command.....	58
subscriber-interface-type command.....	76
support, technical See technical support	
svlan commands	
svlan qos-parameter.....	83
svlan qos-profile.....	64

T

technical support	
contacting JTAC.....	xv

traffic classes	
configuring.....	51
overview.....	13
traffic flow.....	9
traffic-class command.....	51, 52, 75
traffic-class groups	
multiple.....	14
overview.....	14
traffic-class-group command.....	52
triple play configurations.....	90

V

variables	
configuring for shared shaping.....	27

W

weight command.....	75
weight, QoS.....	7
weighted random early detection. See WRED	
WRED (weighted random early detection).....	7, 16
configuring.....	54

