



**JUNOS[™]e Software
for E Series[™] Broadband Services Routers**

Quality of Service Configuration Guide

Release 11.1.x

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

Published: 2010-03-21

Juniper Networks, the Juniper Networks logo, JUNOS, NetScreen, ScreenOS, and Steel-Belted Radius are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOSe is a trademark 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.

JUNOSe™ Software for E Series™ Broadband Services Routers Quality of Service Configuration Guide

Release 11.1.x

Copyright © 2010, Juniper Networks, Inc.

All rights reserved. Printed in USA.

Writing: Krupa Chandrashekar, Bruce Gillham, Sarah Lesway-Ball, Brian Wesley Simmons, Poornima Goswami, Chander Aima

Editing: Benjamin Mann

Illustration: Nathaniel Woodward

Cover Design: Edmonds Design

Revision History

April 2010—FRS JUNOSe 11.1.x

The information in this document is current as of the date listed in the revision history.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. The JUNOS Software 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

READ THIS END USER LICENSE AGREEMENT ("AGREEMENT") BEFORE DOWNLOADING, INSTALLING, OR USING THE SOFTWARE. BY DOWNLOADING, INSTALLING, OR USING THE SOFTWARE OR OTHERWISE EXPRESSING YOUR AGREEMENT TO THE TERMS CONTAINED HEREIN, YOU (AS CUSTOMER OR IF YOU ARE NOT THE CUSTOMER, AS A REPRESENTATIVE/AGENT AUTHORIZED TO BIND THE CUSTOMER) CONSENT TO BE BOUND BY THIS AGREEMENT. IF YOU DO NOT OR CANNOT AGREE TO THE TERMS CONTAINED HEREIN, THEN (A) DO NOT DOWNLOAD, INSTALL, OR USE THE SOFTWARE, AND (B) YOU MAY CONTACT JUNIPER NETWORKS REGARDING LICENSE TERMS.

1. **The Parties.** The parties to this Agreement are (i) Juniper Networks, Inc. (if the Customer's principal office is located in the Americas) or Juniper Networks (Cayman) Limited (if the Customer's principal office is located outside the Americas) (such applicable entity being referred to herein as "Juniper"), and (ii) the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

2. **The Software.** In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller, or which was embedded by Juniper in equipment which Customer purchased from Juniper or an authorized Juniper reseller. "Software" also includes updates, upgrades and new releases of such software. "Embedded Software" means Software which Juniper has embedded in or loaded onto the Juniper equipment and any updates, upgrades, additions or replacements which are subsequently embedded in or loaded onto the equipment.

3. **License Grant.** Subject to payment of the applicable fees and the limitations and restrictions set forth herein, Juniper grants to Customer a non-exclusive and non-transferable license, without right to sublicense, to use the Software, in executable form only, subject to the following use restrictions:

- a. Customer shall use Embedded Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller.
- b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees; provided, however, with respect to the Steel-Belted Radius or Odyssey Access Client software only, Customer shall use such Software on a single computer containing a single physical random access memory space and containing any number of processors. Use of the Steel-Belted Radius or IMS AAA software on multiple computers or virtual machines (e.g., Solaris zones) requires multiple licenses, regardless of whether such computers or virtualizations are physically contained on a single chassis.
- c. Product purchase documents, paper or electronic user documentation, and/or the particular licenses purchased by Customer may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, registered endpoints, concurrent users, sessions, calls, connections, subscribers, clusters, nodes, realms, devices, links, ports or transactions, or require the purchase of separate licenses to use particular features, functionalities, services, applications, operations, or capabilities, or provide throughput, performance, configuration, bandwidth, interface, processing, temporal, or geographical limits. In addition, such limits may restrict the use of the Software to managing certain kinds of networks or require the Software to be used only in conjunction with other specific Software. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.
- d. For any trial copy of the Software, Customer's right to use the Software expires 30 days after download, installation or use of the Software. Customer may operate the Software after the 30-day trial period only if Customer pays for a license to do so. Customer may not extend or create an additional trial period by re-installing the Software after the 30-day trial period.
- e. The Global Enterprise Edition of the Steel-Belted Radius software may be used by Customer only to manage access to Customer's enterprise network. Specifically, service provider customers are expressly prohibited from using the Global Enterprise Edition of the Steel-Belted Radius software to support any commercial network access services.

The foregoing license is not transferable or assignable by Customer. No license is granted herein to any user who did not originally purchase the applicable license(s) for the Software from Juniper or an authorized Juniper reseller.

4. **Use Prohibitions.** Notwithstanding the foregoing, the license provided herein does not permit the Customer to, and Customer agrees not to and shall not: (a) modify, unbundle, reverse engineer, or create derivative works based on the Software; (b) make unauthorized copies of the Software (except as necessary for backup purposes); (c) rent, sell, transfer, or grant any rights in and to any copy of the Software, in any form, to any third party; (d) remove any proprietary notices, labels, or marks on or in any copy of the Software or any product in which the Software is embedded; (e) distribute any copy of the Software to any third party, including as may be embedded in Juniper equipment sold in the secondhand market; (f) use any 'locked' or key-restricted feature, function, service, application, operation, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, service, application, operation, or capability is enabled without a key; (g) distribute any key for the Software provided by Juniper to any third party; (h) use the Software in any manner that extends or is broader than the uses purchased by Customer from Juniper or an authorized Juniper reseller; (i) use Embedded Software on non-Juniper equipment; (j) use Embedded Software (or make it available for use) on Juniper equipment that the Customer did not originally purchase from Juniper or an authorized Juniper reseller; (k) disclose the results of testing or benchmarking of the Software to any third party without the prior written consent of Juniper; or (l) use the Software in any manner other than as expressly provided herein.

5. **Audit.** Customer shall maintain accurate records as necessary to verify compliance with this Agreement. Upon request by Juniper, Customer shall furnish such records to Juniper and certify its compliance with this Agreement.

6. **Confidentiality.** The Parties agree that aspects of the Software and associated documentation are the confidential property of Juniper. As such, Customer shall exercise all reasonable commercial efforts to maintain the Software and associated documentation in confidence, which at a minimum includes restricting access to the Software to Customer employees and contractors having a need to use the Software for Customer's internal business purposes.

7. **Ownership.** Juniper and Juniper's licensors, respectively, retain ownership of all right, title, and interest (including copyright) in and to the Software, associated documentation, and all copies of the Software. Nothing in this Agreement constitutes a transfer or conveyance of any right, title, or interest in the Software or associated documentation, or a sale of the Software, associated documentation, or copies of the Software.

8. **Warranty, Limitation of Liability, Disclaimer of Warranty.** The warranty applicable to the Software shall be as set forth in the warranty statement that accompanies the Software (the "Warranty Statement"). Nothing in this Agreement shall give rise to any obligation to support the Software. Support services may be purchased separately. Any such support shall be governed by a separate, written support services agreement. TO THE MAXIMUM EXTENT PERMITTED BY LAW, JUNIPER SHALL NOT BE LIABLE FOR ANY LOST PROFITS, LOSS OF DATA, OR COSTS OR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, THE SOFTWARE, OR ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE. IN NO EVENT SHALL JUNIPER BE LIABLE FOR DAMAGES ARISING FROM UNAUTHORIZED OR IMPROPER USE OF ANY JUNIPER OR JUNIPER-SUPPLIED SOFTWARE, EXCEPT AS EXPRESSLY PROVIDED IN THE WARRANTY STATEMENT TO THE EXTENT PERMITTED BY LAW, JUNIPER DISCLAIMS ANY AND ALL WARRANTIES IN AND TO THE SOFTWARE (WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE), INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT DOES JUNIPER WARRANT THAT THE SOFTWARE, OR ANY EQUIPMENT OR NETWORK RUNNING THE SOFTWARE, WILL OPERATE WITHOUT ERROR OR INTERRUPTION, OR WILL BE FREE OF VULNERABILITY TO INTRUSION OR ATTACK. In no event shall Juniper's or its suppliers' or licensors' liability to Customer, whether in contract, tort (including negligence), breach of warranty, or otherwise, exceed the price paid by Customer for the Software that gave rise to the claim, or if the Software is embedded in another Juniper product, the price paid by Customer for such other product. Customer acknowledges and agrees that Juniper has set its prices and entered into this Agreement in reliance upon the disclaimers of warranty and the limitations of liability set forth herein, that the same reflect an allocation of risk between the Parties (including the risk that a contract remedy may fail of its essential purpose and cause consequential loss), and that the same form an essential basis of the bargain between the Parties.

9. **Termination.** Any breach of this Agreement or failure by Customer to pay any applicable fees due shall result in automatic termination of the license granted herein. Upon such termination, Customer shall destroy or return to Juniper all copies of the Software and related documentation in Customer's possession or control.

10. **Taxes.** All license fees payable under this agreement are exclusive of tax. Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software. If applicable, valid exemption documentation for each taxing jurisdiction shall be provided to Juniper prior to invoicing, and Customer shall promptly notify Juniper if their exemption is revoked or modified. All payments made by Customer shall be net of any applicable withholding tax. Customer will provide reasonable assistance to Juniper in connection with such withholding taxes by promptly: providing Juniper with valid tax receipts and other required documentation showing Customer's payment of any withholding taxes; completing appropriate applications that would reduce the amount of withholding tax to be paid; and notifying and assisting Juniper in any audit or tax proceeding related to transactions hereunder. Customer shall comply with all applicable tax laws and regulations, and Customer will promptly pay or reimburse Juniper for all costs and damages related to any liability incurred by Juniper as a result of Customer's non-compliance or delay with its responsibilities herein. Customer's obligations under this Section shall survive termination or expiration of this Agreement.

11. **Export.** Customer agrees to comply with all applicable export laws and restrictions and regulations of any United States and any applicable foreign agency or authority, and not to export or re-export the Software or any direct product thereof in violation of any such restrictions, laws or regulations, or without all necessary approvals. Customer shall be liable for any such violations. The version of the Software supplied to Customer may contain encryption or other capabilities restricting Customer's ability to export the Software without an export license.

12. **Commercial Computer Software.** The Software is "commercial computer software" and is provided with restricted rights. Use, duplication, or disclosure by the United States government is subject to restrictions set forth in this Agreement and as provided in DFARS 227.7201 through 227.7202-4, FAR 12.212, FAR 27.405(b)(2), FAR 52.227-19, or FAR 52.227-14(ALT III) as applicable.

13. **Interface Information.** To the extent required by applicable law, and at Customer's written request, Juniper shall provide Customer with the interface information needed to achieve interoperability between the Software and another independently created program, on payment of applicable fee, if any. Customer shall observe strict obligations of confidentiality with respect to such information and shall use such information in compliance with any applicable terms and conditions upon which Juniper makes such information available.

14. **Third Party Software.** Any licensor of Juniper whose software is embedded in the Software and any supplier of Juniper whose products or technology are embedded in (or services are accessed by) the Software shall be a third party beneficiary with respect to this Agreement, and such licensor or vendor shall have the right to enforce this Agreement in its own name as if it were Juniper. In addition, certain third party software may be provided with the Software and is subject to the accompanying license(s), if any, of its respective owner(s). To the extent portions of the Software are distributed under and subject to open source licenses obligating Juniper to make the source code for such portions publicly available (such as the GNU General Public License ("GPL") or the GNU Library General Public License ("LGPL")), Juniper will make such source code portions (including Juniper modifications, as appropriate) available upon request for a period of up to three years from the date of distribution. Such request can be made in writing to Juniper Networks, Inc., 1194 N. Mathilda Ave., Sunnyvale, CA 94089, ATTN: General Counsel. You may obtain a copy of the GPL at <http://www.gnu.org/licenses/gpl.html>, and a copy of the LGPL at <http://www.gnu.org/licenses/lgpl.html>.

15. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. The provisions of the U.N. Convention for the International Sale of Goods shall not apply to this Agreement. For any disputes arising under this Agreement, the Parties hereby consent to the personal and exclusive jurisdiction of, and venue in, the state and federal courts within Santa Clara County, California. This Agreement constitutes the entire and sole agreement between Juniper and the Customer with respect to the Software, and supersedes all prior and contemporaneous

agreements relating to the Software, whether oral or written (including any inconsistent terms contained in a purchase order), except that the terms of a separate written agreement executed by an authorized Juniper representative and Customer shall govern to the extent such terms are inconsistent or conflict with terms contained herein. No modification to this Agreement nor any waiver of any rights hereunder shall be effective unless expressly assented to in writing by the party to be charged. If any portion of this Agreement is held invalid, the Parties agree that such invalidity shall not affect the validity of the remainder of this Agreement. This Agreement and associated documentation has been written in the English language, and the Parties agree that the English version will govern. (For Canada: Les parties aux présentes confirment leur volonté que cette convention de même que tous les documents y compris tout avis qui s'y rattache, soient rédigés en langue anglaise. (Translation: The parties confirm that this Agreement and all related documentation is and will be in the English language)).

Abbreviated Table of Contents

About the Documentation

xxix

Part 1	QoS on the E Series Router	
Chapter 1	Quality of Service Overview	3
Part 2	Classifying, Queuing, and Dropping Traffic	
Chapter 2	Defining Service Levels with Traffic Classes and Traffic-Class Groups	13
Chapter 3	Configuring Queue Profiles for Buffer Management	17
Chapter 4	Configuring Dropping Behavior with RED and WRED	25
Chapter 5	Gathering Statistics for Rates and Events in the Queue	37
Part 3	Scheduling and Shaping Traffic	
Chapter 6	QoS Scheduler Hierarchy Overview	45
Chapter 7	Configuring Rates and Weights in the Scheduler Hierarchy	51
Chapter 8	Configuring Strict-Priority Scheduling	59
Chapter 9	Shared Shaping Overview	71
Chapter 10	Configuring Simple Shared Shaping of Traffic	79
Chapter 11	Configuring Variables in the Simple Shared Shaping Algorithm	89
Chapter 12	Configuring Compound Shared Shaping of Traffic	99
Chapter 13	Configuring Implicit and Explicit Constituent Selection for Shaping	107
Chapter 14	Monitoring a QoS Scheduler Hierarchy	123
Part 4	Creating a QoS Scheduler Hierarchy on an Interface with QoS Profiles	
Chapter 15	QoS Profile Overview	127
Chapter 16	Configuring and Attaching QoS Profiles to an Interface	131
Chapter 17	Configuring Shadow Nodes for Queue Management	149
Chapter 18	Monitoring a Scheduler Hierarchy on an Interface with QoS Profiles	155
Part 5	Interface Solutions for QoS	
Chapter 19	Configuring an Integrated Scheduler to Provide QoS for ATM	159
Chapter 20	Configuring QoS for Gigabit Ethernet Interfaces and VLAN Subinterfaces	177
Chapter 21	Configuring QoS for 802.3ad Link Aggregation Groups	183

Chapter 22	Configuring QoS for L2TP Sessions	197
Chapter 23	Configuring Interface Sets for QoS	207
Part 6	Managing Queuing and Scheduling with QoS Parameters	
Chapter 24	QoS Parameter Overview	225
Chapter 25	Configuring a QoS Parameter	229
Chapter 26	Configuring Hierarchical QoS Parameters	261
Chapter 27	Configuring IP Multicast Bandwidth Adjustment with QoS Parameters	269
Chapter 28	Configuring the Shaping Mode for Ethernet with QoS Parameters	281
Chapter 29	Configuring Byte Adjustment for Shaping Rates with QoS Parameters	293
Chapter 30	Configuring the Downstream Rate Using QoS Parameters	301
Part 7	Monitoring and Troubleshooting QoS	
Chapter 31	Monitoring QoS on E Series Routers	313
Chapter 32	Troubleshooting QoS	357
Part 8	Index	
	Index	361

Table of Contents

	About the Documentation	xxix
	E Series and JUNOS ^e Documentation and Release Notes	xxix
	Audience	xxix
	E Series and JUNOS ^e Text and Syntax Conventions	xxix
	Obtaining Documentation	xxx
	Documentation Feedback	xxx
	Requesting Technical Support	xxx
	Self-Help Online Tools and Resources	xxxii
	Opening a Case with JTAC	xxxii
Part 1	QoS on the E Series Router	
Chapter 1	Quality of Service Overview	3
	QoS on the E Series Router Overview	3
	QoS Audience	4
	QoS Platform Considerations	4
	Interface Specifiers	5
	QoS Terms	5
	QoS Features	7
	Configuring QoS on the E Series Router	8
	QoS References	9
Part 2	Classifying, Queuing, and Dropping Traffic	
Chapter 2	Defining Service Levels with Traffic Classes and Traffic-Class Groups	13
	Traffic Class and Traffic-Class Groups Overview	13
	Best-Effort Forwarding	13
	Traffic-Class Groups Overview	14
	Configuring Traffic Classes That Define Service Levels	14
	Configuring Traffic-Class Groups That Define Service Levels	15
	Monitoring Traffic Classes and Traffic-Class Groups for Defined Levels of Service	16

Chapter 3	Configuring Queue Profiles for Buffer Management	17
	Queuing and Buffer Management Overview	17
	Static Oversubscription	18
	Dynamic Oversubscription	18
	Color-Based Thresholding	18
	Memory Requirements for Queue and Buffers	19
	Guidelines for Managing Queue Thresholds	19
	Guidelines for Configuring a Maximum Threshold	19
	Guidelines for Configuring a Minimum Threshold	20
	Guidelines for Managing Buffers	20
	Guidelines for Managing Buffer Starvation	21
	Configuring Queue Profiles to Manage Buffers and Thresholds	22
	Monitoring Queues and Buffers	24
Chapter 4	Configuring Dropping Behavior with RED and WRED	25
	Dropping Behavior Overview	25
	RED and WRED Overview	26
	Configuring RED	27
	Example: Configuring Average Queue Length for RED	28
	Example: Configuring Dropping Thresholds for RED	29
	Example: Configuring Color-Blind RED	29
	Configuring WRED	31
	Example: Configuring Different Treatment of Colored Packets for WRED	33
	Example: Defining Different Drop Behavior for Each Traffic Class for WRED	33
	Example: Configuring WRED and Dynamic Queue Thresholds	34
	Monitoring RED and WRED	36
Chapter 5	Gathering Statistics for Rates and Events in the Queue	37
	QoS Statistics Overview	37
	Rate Statistics	38
	Event Statistics	38
	Bulk Statistics Support for QoS Statistics	38
	Configuring Statistic Profiles for QoS	39
	Configuring Rate Statistics	39
	Configuring Event Statistics	40
	Clearing QoS Statistics on the Egress Queue	41
	Clearing QoS Statistics on the Fabric Queue	42
	Monitoring QoS Statistics for Rates and Events	42

Part 3	Scheduling and Shaping Traffic	
Chapter 6	QoS Scheduler Hierarchy Overview	45
	Scheduler Hierarchy Overview	45
	Shaping Rates, Assured Rates, and Relative Weights in a Scheduler Hierarchy	46
	Configuring a Scheduler Hierarchy	47
	Configuring a Scheduler Profile for a Scheduler Node or Queue	48
	Using Expressions for Bandwidth and Burst Values in a Scheduler Profile	48
Chapter 7	Configuring Rates and Weights in the Scheduler Hierarchy	51
	Rate Shaping and Port Shaping Overview	51
	Configuring Rate Shaping for a Scheduler Node or Queue	52
	Configuring Port Shaping	53
	Static and Hierarchical Assured Rate Overview	54
	Configuring an Assured Rate for a Scheduler Node or Queue	55
	Configuring a Static Assured Rate	55
	Configuring a Hierarchical Assured Rate	56
	Changing the Assured Rate to an HRR Weight	56
	Configuring the HRR Weight for a Scheduler Node or Queue	57
Chapter 8	Configuring Strict-Priority Scheduling	59
	Strict-Priority and Relative Strict-Priority Scheduling Overview	59
	Relative Strict-Priority Scheduling Overview	60
	Comparison of True Strict Priority with Relative Strict Priority Scheduling	61
	Schedulers and True Strict Priority	61
	Schedulers and Relative Strict Priority	62
	Relative Strict Priority on ATM Modules	63
	Oversubscribing ATM Ports	64
	Minimizing Latency on the SAR Scheduler	64
	HRR Scheduler Behavior and Strict-Priority Scheduling	64
	Zero-Weight Queues	64
	Setting the Burst Size in a Shaping Rate	65
	Special Shaping Rate for Nonstrict Queues	65
	Configuring Strict-Priority Scheduling	66
	Configuring Relative Strict-Priority Scheduling for Aggregate Shaping Rates	68
Chapter 9	Shared Shaping Overview	71
	Shared Shaping Overview	71
	Shared Shaper Terms	72
	How Shared Shaping Works	72
	Active Constituents for Shared Shaping	73

	Guidelines for Configuring Simple and Compound Shared Shaping	74
	Shared Shaping and Individual Shaping	74
	Shared Shaping and Best-Effort Queues and Nodes	74
	ATM and Shared Shaping	75
	Sharing Bandwidth with the SAR	75
	Shared Shaping and Low-CDV Mode	75
	Logical Interface Traffic Carried in Other Queues	76
	Traffic Starvation and Shared Shaping	76
	Oversubscription and Shared Shaping	77
	Burst Size and Shared Shaping	77
Chapter 10	Configuring Simple Shared Shaping of Traffic	79
	Simple Shared Shaping Overview	79
	Bandwidth Allocation for Simple Shared Shaping	79
	Simple Shared Shaping on the Best-Effort Scheduler Node	79
	Simple Shared Shaping for Triple-Play Networks	80
	Configuring Simple Shared Shaping	81
	Example: Simple Shared Shaping for ATM VCs	83
	Example: Simple Shared Shaping for ATM VPs	85
	Example: Simple Shared Shaping for Ethernet	86
Chapter 11	Configuring Variables in the Simple Shared Shaping Algorithm	89
	Simple Shared Shaping Algorithm Overview	89
	Simple Shared Shaper Algorithm Calculations	90
	Variables of the Simple Shared Shaper Algorithm	91
	Guidelines for Controlling the Simple Shared Shaper Algorithm	93
	Configuring Simple Shared Shaper Algorithm Variables	93
	Sample Process for Controlling the Simple Shared Shaper Algorithm	95
Chapter 12	Configuring Compound Shared Shaping of Traffic	99
	Compound Shared Shaping Overview	99
	Supported Hardware for Compound Shared Shaping	99
	Bandwidth Allocation for Compound Shared Shaping	100
	Configuring Compound Shared Shaping	100
	Example: Compound Shared Shaping for ATM VCs	102
	Example: Compound Shared Shaping for ATM VPs	104

Chapter 13	Configuring Implicit and Explicit Constituent Selection for Shaping	107
	Constituent Selection for Shared Shaping Overview	107
	Types of Shared Shaper Constituents	108
	Implicit Constituent Selection Overview	109
	Implicit Bandwidth Allocation for Compound Shared Shaping	111
	Weighted Compound Shared Shaping Example	113
	Configuring Implicit Constituents for Simple or Compound Shared Shaping	114
	Explicit Constituent Selection Overview	115
	Explicit Shared Shaping Example	116
	Explicit Weighted Compound Shared Shaping Example	117
	Configuring Explicit Constituents for Simple or Compound Shared Shaping	120
Chapter 14	Monitoring a QoS Scheduler Hierarchy	123
	Monitoring QoS Scheduling and Shaping	123
Part 4	Creating a QoS Scheduler Hierarchy on an Interface with QoS Profiles	
Chapter 15	QoS Profile Overview	127
	QoS Profile Overview	127
	Managing System Resources for Nodes and Queues	127
	Scaling Subscribers on the TFA ASIC with QoS	128
Chapter 16	Configuring and Attaching QoS Profiles to an Interface	131
	Supported Interface Types for QoS Profiles	131
	Configuring a QoS Profile	132
	Attaching a QoS Profile to an Interface	134
	Attaching a QoS Profile to a Base Interface	134
	Attaching a QoS Profile to an ATM VP	134
	Attaching a QoS Profile to an S-VLAN	135
	Attaching a QoS Profile to a Port Type	135
	Munged QoS Profile Overview	136
	Sample Munged QoS Profile Process	137
	Example: Port-Type QoS Profile Attachment	139
	Example: QoS Profile Attachment to Port	141
	Example: Diffserv Configuration with Multiple Traffic-Class Groups	143

Chapter 17	Configuring Shadow Nodes for Queue Management	149
	Shadow Node Overview	149
	Shadow Nodes and Scheduler Behavior	150
	Managing System Resources for Shadow Nodes	151
	Configuring Shadow Nodes	152
	Example: Shadow Nodes over VLAN and IP Queues	153
	Example: Shadow Nodes on the Same Traffic-Class Group	154
	Example: Shadow Nodes on Different Traffic-Class Groups	154
Chapter 18	Monitoring a Scheduler Hierarchy on an Interface with QoS Profiles	155
	Monitoring a Scheduler Hierarchy on an Interface with QoS Profiles	155
Part 5	Interface Solutions for QoS	
Chapter 19	Configuring an Integrated Scheduler to Provide QoS for ATM	159
	ATM Integrated Scheduler Overview	159
	Backpressure and the Integrated Scheduler	160
	VP Shaping	161
	Integrating the HRR Scheduler and SAR Scheduler	161
	Per-Packet Queuing on the SAR Scheduler Overview	163
	Operational QoS Shaping Mode for ATM Interfaces Overview	164
	ERX7xx Models, ERX14xx Models, and the ERX310 Router	164
	E120 Router and E320 Router	165
	Guidelines for Configuring QoS over ATM	166
	Configuring Default Integrated Mode for ATM Interface	168
	Configuring Low-Latency Mode for Per-Port Queuing on ATM Interfaces	169
	Configuring Low-CDV Mode for Per-Port Queuing on ATM Interfaces	171
	Configuring the QoS Shaping Mode for ATM Interfaces	174
	Disabling Per-Port Queuing on ATM Interfaces	175
	Monitoring QoS Configurations for ATM	175
Chapter 20	Configuring QoS for Gigabit Ethernet Interfaces and VLAN Subinterfaces	177
	Providing QoS for Ethernet Overview	177
	QoS Shaping Mode for Ethernet Interfaces Overview	178
	Configuring the QoS Shaping Mode for Ethernet Interfaces	179
	Creating a QoS Interface Hierarchy for Bulk-Configured VLAN Subinterfaces with RADIUS	180
	Monitoring QoS Configurations for Ethernet	182

Chapter 21	Configuring QoS for 802.3ad Link Aggregation Groups	183
	QoS for 802.3ad Link Aggregation Interfaces Overview	183
	Types of Load Balancing	183
	Munged QoS Profiles and Load Balancing	184
	802.3ad Link Aggregation and QoS Parameters	184
	QoS and Ethernet Link Redundancy	185
	Active Link Failure and QoS	185
	Administratively Disabling a Link and QoS	185
	Adding a New Link to the LAG and QoS	185
	Hashed Load Balancing for 802.3ad Link Aggregation Groups Overview	186
	Sample Scheduler Hierarchy for Hashed Load Balancing	186
	Subscriber Load Balancing for 802.3ad Link Aggregation Groups	
	Overview	186
	S-VLANs and Subscriber Load Balancing	187
	PPPoE over VLANs and Subscriber Load Balancing	187
	PPPoE over Ethernet (No VLANs) and Subscriber Load Balancing	187
	MPLS over LAG and Subscriber Load Balancing	187
	Sample Scheduler Hierarchy for Subscriber Load Balancing	188
	Subscriber Allocation in 802.3ad Link Aggregation Groups	189
	Guidelines for Configuring QoS over 802.3ad Link Aggregation Groups	190
	Configuring the Scheduler Hierarchy for Hashed Load Balancing in 802.3ad	
	Link Aggregation Groups	191
	Enabling Default Subscriber Load Balancing for 802.3ad Link Aggregation	
	Groups	191
	Configuring the Scheduler Hierarchy for Subscriber Load Balancing in 802.3ad	
	Link Aggregation Groups	192
	Configuring Load Rebalancing for 802.3ad Link Aggregation Groups	193
	Configuring Load-Rebalancing Parameters	193
	Configuring the System to Dynamically Rebalance the LAG	195
	Monitoring QoS Configurations for 802.3ad Link Aggregation Groups	195
Chapter 22	Configuring QoS for L2TP Sessions	197
	Providing QoS for L2TP Overview	197
	Sample Scheduler Hierarchies for L2TP	197
	Configuring QoS for an L2TP Session	199
	Configuring QoS for an L2TP LNS Session	200
	Configuring QoS for an L2TP LAC Session	201
	Configuring QoS for Tunnel-Server Ports for L2TP LNS Sessions	202
	QoS and L2TP TX Speed AVP 24 Overview	203
	Logical Interfaces and Shared-Shaping Rates	203
	Shaping Mode	204
	Monitoring QoS Configurations for L2TP	204

Chapter 23 Configuring Interface Sets for QoS 207

Interface Sets for QoS Overview	207
Interface Set Terms	207
Architecture of Interface Sets for QoS	208
Interface Set Parents and Types	209
Sample Interface Columns and Scheduler Hierarchies	209
Scheduling and Shaping Interface Sets	210
Configuring Interface Sets for Scheduling and Queuing	211
Configuring Interface Supersets for QoS	212
Configuring an Interface Superset	212
Restricting an Interface Superset to an S-VLAN ID or an ATM VP	212
Configuring Interface Sets for QoS	213
Configuring an Interface Set	213
Deleting an Interface Set from an Interface Superset	213
Adding Member Interfaces to an Interface Set	214
Adding Interfaces to an Interface Set with the CLI	214
Adding Interfaces to an Interface Set with RADIUS	215
Changing and Deleting Interface Members in an Interface Set	215
Changing Interface Members with Upper-Layer Protocols in an Interface Set	215
Creating a QoS Parameter on an Interface Superset or Interface Set	216
Configuring a QoS Parameter Definition for an Interface Superset or an Interface Set	216
Creating a QoS Parameter Instance for an Interface Superset	217
Creating a QoS Parameter Instance for an Interface Set	217
Attaching a QoS Profile to an Interface Superset or an Interface Set	218
Configuring a QoS Profile for an Interface Superset or an Interface Set	218
Attaching a QoS Profile to an Interface Superset	218
Attaching a QoS Profile to an Interface Set	219
Deleting an Interface Superset or an Interface Set	219
Deleting an Interface Superset	219
Deleting an Interface Set	220
Example: Configuring Interface Sets for 802.3ad Link Aggregation Groups	220

Part 6 Managing Queuing and Scheduling with QoS Parameters

Chapter 24 QoS Parameter Overview 225

QoS Parameter Overview	225
QoS Parameter Audience	225
QoS Parameter Terms	226

Relationship Among QoS Parameters, Scheduler Profiles, and QoS Profiles	227
QoS Administrator Tasks	227
QoS Client Tasks	228

Chapter 25**Configuring a QoS Parameter 229**

Parameter Definition Attributes for QoS Administrators Overview	229
Naming Guidelines for QoS Parameters	230
Interface Types and QoS Parameters	231
Controlled-Interface Types	231
Instance-Interface Types	232
Subscriber-Interface Types	233
Range of QoS Parameters	234
Applications and QoS Parameters	235
Scheduler Profiles and Parameter Expressions for QoS Administrators	235
Referencing a Parameter Definition in a Scheduler Profile	235
Removing or Modifying a Scheduler Profile	236
Using Expressions for QoS Parameters	236
Operators and Precedence	236
Specifying a Range in Expressions	237
Configuring a Basic Parameter Definition for QoS Administrators	238
Parameter Instances for QoS Clients Overview	240
Global QoS Parameter Instance Overview	240
QoS Parameters for Interfaces Overview	241
Creating Parameter Instances	242
Creating a Global Parameter Instance	242
Creating a Parameter Instance for an Interface	242
Creating a Parameter Instance for an ATM VP	242
Creating a Parameter Instance for an S-VLAN	243
Example: QoS Parameter Configuration for Controlling Subscriber Bandwidth	243
Procedure for QoS Administrators	245
Procedure for QoS Clients	250
Monitoring the Subscriber Configuration	252
Complete Configuration Example	256
QoS Administrator Configuration	256
QoS Client Configuration	258

Chapter 26**Configuring Hierarchical QoS Parameters 261**

Hierarchical QoS Parameters Overview	261
Guidelines for Configuring Hierarchical Parameters	261
Configuring a Parameter Definition to Calculate Hierarchical Instances	262
Example: QoS Parameter Configuration for Hierarchical Parameters	263
Procedure for QoS Administrators	264
Procedure for QoS Clients	265
Monitoring Hierarchical QoS Parameters	266

	Complete Configuration Example	267
	QoS Administrator Configuration	267
	QoS Client Configuration	267
Chapter 27	Configuring IP Multicast Bandwidth Adjustment with QoS Parameters	269
	IP Multicast Bandwidth Adjustment for QoS Overview	269
	Guidelines for Configuring IP Multicast Adjustment for QoS	271
	Configuring a Parameter Definition for IP Multicast Bandwidth Adjustment	271
	Example: QoS Parameter Configuration for IP Multicast Bandwidth Adjustment	273
	Monitoring the Configuration	276
	Complete Configuration Example	279
Chapter 28	Configuring the Shaping Mode for Ethernet with QoS Parameters	281
	Cell Shaping Mode Using QoS Parameters Overview	281
	Overriding the QoS Shaping Mode	281
	Module Types and Capabilities for QoS Cell Mode Application	282
	Cell Tax Adjustment Using QoS Cell Mode	282
	Relationship with QoS Downstream Rate Application	283
	Guidelines for Configuring the Cell Shaping Mode with QoS Parameters	283
	Configuring a Parameter Definition to Shape Ethernet Traffic Using Cell Mode	284
	Example: QoS Parameter Configuration for QoS Cell Mode and Byte Adjustment for Cell Shaping	286
	Complete Configuration Example	290
Chapter 29	Configuring Byte Adjustment for Shaping Rates with QoS Parameters	293
	Byte Adjustment for ADSL and VDSL Traffic Overview	293
	Byte Adjustment for Cell Shaping of ADSL Traffic Overview	293
	Calculation and Example of Byte Adjustment for Cell Shaping	294
	Byte Adjustment for Frame Shaping of VDSL Traffic Overview	295
	System Calculation for Byte Adjustment of ADSL and VDSL Traffic	295
	Guidelines for Configuring Byte Adjustment of Cell and Frame Shaping Rates Using QoS Parameters	296
	Configuring a Parameter Definition to Adjust Cell Shaping Rates for ADSL Traffic	297
	Configuring a Parameter Definition to Adjust Frame Shaping Rates for VDSL Traffic	299

Chapter 30 Configuring the Downstream Rate Using QoS Parameters 301

QoS Downstream Rate Application Overview	301
Downstream Rate and the Shaping Mode	301
QoS Adaptive Mode and Downstream Rate	302
Obtaining Downstream Rates from a DSL Forum VSA	302
Guidelines for Configuring QoS Downstream Rate	303
Configuring a Parameter Definition for QoS Downstream Rate	303
Example: QoS Parameter Configuration for QoS Downstream Rate	305
Complete Configuration Example	309

Part 7 Monitoring and Troubleshooting QoS

Chapter 31 Monitoring QoS on E Series Routers 313

Monitoring Service Levels with Traffic Classes	314
Monitoring Service Levels with Traffic-Class Groups	315
Monitoring Queue Thresholds	316
Monitoring Queue Profiles	319
Monitoring Drop Profiles for RED and WRED	320
Monitoring the QoS Scheduler Hierarchy	322
Monitoring the Configuration of Scheduler Profiles	327
Monitoring Shared Shapers	329
Monitoring Shared Shaper Algorithm Variables	330
Monitoring Forwarding and Drop Events on the Egress Queue	331
Monitoring Forwarding and Drop Rates on the Egress Queue	333
Monitoring Queue Statistics for the Fabric	337
Monitoring the Configuration of Statistics Profiles	338
Monitoring the QoS Profiles Attached to an Interface	339
Monitoring the Configuration of QoS Port-Type Profiles	340
Monitoring the Configuration of QoS Profiles	341
Monitoring the QoS Configuration of ATM Interfaces	343
Monitoring the QoS Configuration of IP Interfaces	345
Monitoring the QoS Configuration of Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet Interfaces	347
Monitoring the QoS Configuration of IEEE 802.3ad Link Aggregation Group Bundles	349
Monitoring the Configuration of QoS Interface Sets	350
Monitoring the Configuration of QoS Interface Supersets	351
Monitoring the AAA Downstream Rate for QoS	352
Monitoring QoS Parameter Instances	352
Monitoring QoS Parameter Definitions	355

Chapter 32	Troubleshooting QoS	357
	Troubleshooting Memory and Processor Use for Egress Queue Rate Statistics and Events	357
Part 8	Index	
	Index	361

List of Figures

Part 1

QoS on the E Series Router

Chapter 1	Quality of Service Overview	3
	Figure 1: Traffic Flow Through an E Series Router	4

Part 2

Classifying, Queuing, and Dropping Traffic

Chapter 4	Configuring Dropping Behavior with RED and WRED	25
	Figure 2: Packets Dropped as Queue Length Increases	27
	Figure 3: Color-Blind RED Drop Profile with Colorless Queue Profile	30
	Figure 4: Color-Blind RED Drop Profile with Color-Sensitive Queue Profile	30
	Figure 5: Different Treatment of Colored Packets	33
	Figure 6: Defining Different Drop Behavior for Each Queue	34
	Figure 7: WRED and Dynamic Queue Thresholding	36

Part 3

Scheduling and Shaping Traffic

Chapter 6	QoS Scheduler Hierarchy Overview	45
	Figure 8: QoS Scheduler Hierarchy	46
Chapter 7	Configuring Rates and Weights in the Scheduler Hierarchy	51
	Figure 9: Port Shaping on an Ethernet Module	52
	Figure 10: Hierarchical Assured Rate	55
Chapter 8	Configuring Strict-Priority Scheduling	59
	Figure 11: Sample Strict-Priority Scheduling Hierarchy	60
	Figure 12: True Strict-Priority Configuration	62
	Figure 13: Relative Strict-Priority Configuration	63
	Figure 14: Tuning Latency on Strict-Priority Queues	66
	Figure 15: Sample Strict-Priority Scheduling Hierarchy	67
	Figure 16: Sample Relative Strict-Priority Scheduler Hierarchy	69
Chapter 9	Shared Shaping Overview	71
	Figure 17: Implicit Constituent Selection for Compound Shared Shaper: Mixed Interface Types	76
Chapter 10	Configuring Simple Shared Shaping of Traffic	79
	Figure 18: Simple Shared Shaping over ATM	80
	Figure 19: Simple Shared Shaping over Ethernet	81
	Figure 20: VP Shared Shaping	85
	Figure 21: Hierarchical Simple Shared Shaping over Ethernet	86
Chapter 11	Configuring Variables in the Simple Shared Shaping Algorithm	89
	Figure 22: Simple Shared Shaper Behavior Without Algorithm Controls	89

	Figure 23: Less Conservative Simple Shared Shaper Behavior	90
	Figure 24: More Liberal Simple Shared Shaper Behavior	90
	Figure 25: Dynamic Rate When Video Flow Starts	96
	Figure 26: Dynamic Rate When Video Flow Stops	97
Chapter 12	Configuring Compound Shared Shaping of Traffic	99
	Figure 27: VC Compound Shared Shaping Example	102
	Figure 28: VP Compound Shared Shaping Example	105
Chapter 13	Configuring Implicit and Explicit Constituent Selection for Shaping	107
	Figure 29: Implicit Constituent Selection for Compound Shared Shaper at Best-Effort Node	110
	Figure 30: Implicit Constituent Selection for Compound Shared Shaper at Best-Effort Queue	111
	Figure 31: Weighted Shared Shaping	113
	Figure 32: Implicit Constituent Selection for Compound Shared Shaper: Mixed Interface Types	114
	Figure 33: Explicit Constituent Selection	117
	Figure 34: Case 1: Explicit Constituent Selection with Weighted Constituents	118
	Figure 35: Case 2: Explicit Constituent Selection with Weighted Constituents	119
 Part 4	 Creating a QoS Scheduler Hierarchy on an Interface with QoS Profiles	
Chapter 16	Configuring and Attaching QoS Profiles to an Interface	131
	Figure 36: Munged Profile Example	138
	Figure 37: Attaching QoS Profiles to ATM Subinterfaces	139
	Figure 38: Attaching QoS Profile to ATM Interface and Subinterface	142
	Figure 39: Diffserv Configuration with Multiple Traffic-Class Groups	146
	Figure 40: Diffserv Configuration Without Traffic-Class Groups	147
Chapter 17	Configuring Shadow Nodes for Queue Management	149
	Figure 41: Phantom Nodes	150
	Figure 42: Shadow Nodes	151
 Part 5	 Interface Solutions for QoS	
Chapter 19	Configuring an Integrated Scheduler to Provide QoS for ATM	159
	Figure 43: Integrated ATM Scheduler	161
	Figure 44: Default Integrated Mode	168
	Figure 45: Low-Latency Mode	170
	Figure 46: Low-CDV Mode (per-VP CDVT)	172
	Figure 47: Low-CDV Mode (per-VC CDVT)	172
Chapter 21	Configuring QoS for 802.3ad Link Aggregation Groups	183
	Figure 48: 802.3ad Link Aggregation Scheduler Hierarchy	186
	Figure 49: Subscriber LoadBalanced Scheduler Hierarchy for Port 0	188
	Figure 50: Subscriber LoadBalanced Scheduler Hierarchy for Port 1	189
	Figure 51: Subscriber Allocation and Load Balancing	189
Chapter 22	Configuring QoS for L2TP Sessions	197

	Figure 52: LNS (Non-MLPPP) Scheduler Hierarchy	198
	Figure 53: LNS (MLPPP) QoS Scheduler Hierarchy	198
	Figure 54: LAC over Ethernet (Without VLANs) Scheduler Hierarchy	198
	Figure 55: LAC over Ethernet (With LANs) Scheduler Hierarchy	199
	Figure 56: LAC over ATM	199
Chapter 23	Configuring Interface Sets for QoS	207
	Figure 57: VLAN Interface Column with Interface Sets	210
	Figure 58: Scheduler Hierarchy with Nodes at Interface Set and Superset	210
 Part 6	 Managing Queuing and Scheduling with QoS Parameters	
Chapter 24	QoS Parameter Overview	225
	Figure 59: Relationship of Parameter Definitions, Scheduler Profiles, and QoS Profiles	227
Chapter 25	Configuring a QoS Parameter	229
	Figure 60: Physical Network Topology	244
	Figure 61: QoS Scheduler Hierarchy	244
Chapter 26	Configuring Hierarchical QoS Parameters	261
	Figure 62: Hierarchical Parameters Scheduler Hierarchy	264
Chapter 27	Configuring IP Multicast Bandwidth Adjustment with QoS Parameters	269
	Figure 63: Scheduler Hierarchy with QoS Adjustment for IP Multicast	274
Chapter 28	Configuring the Shaping Mode for Ethernet with QoS Parameters	281
	Figure 64: Byte Adjustment for VC1 and VC2	286
Chapter 29	Configuring Byte Adjustment for Shaping Rates with QoS Parameters	293
	Figure 65: Byte Adjustment Calculation for Ethernet and ATM Encapsulations	294

List of Tables

	About the Documentation	xxix
	Table 1: Notice Icons	xxx
	Table 2: Text and Syntax Conventions	xxx
Part 1	QoS on the E Series Router	
Chapter 1	Quality of Service Overview	3
	Table 3: QoS Terminology	5
	Table 4: QoS Features	7
Part 2	Classifying, Queuing, and Dropping Traffic	
Chapter 3	Configuring Queue Profiles for Buffer Management	17
	Table 5: Egress Memory and Region Size on ASIC Line Modules	19
Part 3	Scheduling and Shaping Traffic	
Chapter 9	Shared Shaping Overview	71
	Table 6: Shared Shaper Terminology Used in This Chapter	72
	Table 7: Comparison of Simple and Compound Shared Shaping	73
Chapter 11	Configuring Variables in the Simple Shared Shaping Algorithm	89
	Table 8: Guidelines for Configuring Simple Shared Shaper Algorithm Variables	93
	Table 9: Rising Edge Sample When Video Flow Starts	95
	Table 10: Data When Video Flow Stops	96
Chapter 13	Configuring Implicit and Explicit Constituent Selection for Shaping	107
	Table 11: Comparison of Implicit and Explicit Shared Shaping	108
	Table 12: Bandwidth Allocation for Case 1 Explicit Constituents	118
	Table 13: Bandwidth Allocation for Case 2 Explicit Constituents	119
Part 4	Creating a QoS Scheduler Hierarchy on an Interface with QoS Profiles	
Chapter 16	Configuring and Attaching QoS Profiles to an Interface	131
	Table 14: Interface Types and Supported Commands	131
Chapter 17	Configuring Shadow Nodes for Queue Management	149
	Table 15: Shadow Node Consumption of Node and Queue Resources	152

Part 5**Interface Solutions for QoS**

Chapter 19	Configuring an Integrated Scheduler to Provide QoS for ATM	159
	Table 16: qos-mode-port Commands	162
	Table 17: Operational Shaping Modes for ERX7xx Models, ERX14xx Models, and the ERX310 Router	165
	Table 18: Operational Shaping Modes for the E120 Router and E320 Router	166
Chapter 20	Configuring QoS for Gigabit Ethernet Interfaces and VLAN Subinterfaces	177
	Table 19: Operational Shaping Modes	178
Chapter 21	Configuring QoS for 802.3ad Link Aggregation Groups	183
	Table 20: Load Balancing Algorithm Parameters	193
Chapter 23	Configuring Interface Sets for QoS	207
	Table 21: Interface Set Terms	208

Part 6**Managing Queuing and Scheduling with QoS Parameters**

Chapter 24	QoS Parameter Overview	225
	Table 22: QoS Parameter Terminology Used in This Chapter	226
Chapter 25	Configuring a QoS Parameter	229
	Table 23: Attributes in Parameter Definitions	229
	Table 24: Sample Parameter Names	230
	Table 25: Valid and Invalid Parameter Names	231
	Table 26: Operators for Parameter Expressions	237
Chapter 28	Configuring the Shaping Mode for Ethernet with QoS Parameters	281
	Table 27: Supported Interfaces for qos-shaping-mode and qos-cell-mode Commands	282
	Table 28: Byte Adjustment for Subscribers VC1 and VC2	286
Chapter 29	Configuring Byte Adjustment for Shaping Rates with QoS Parameters	293
	Table 29: Header Lengths for Ethernet Encapsulation	294
	Table 30: Header Lengths for ATM Encapsulation	295
	Table 31: Byte Adjustment Values for Frame and Cell Shaping Modes	296
Chapter 30	Configuring the Downstream Rate Using QoS Parameters	301
	Table 32: Access Loop Types and Resultant Shaping Mode	302
	Table 33: Shaping Rate and Shaping Mode	306

Part 7**Monitoring and Troubleshooting QoS**

Chapter 31	Monitoring QoS on E Series Routers	313
	Table 34: show traffic-class Output Fields	314
	Table 35: show traffic-class-group Output Fields	315
	Table 36: show qos queue-thresholds Output Fields	319
	Table 37: show queue-profile Output Fields	320
	Table 38: show drop-profile Output Fields	321
	Table 39: show qos scheduler-hierarchy Output Fields	327
	Table 40: show scheduler-profile Output Fields	328
	Table 41: show qos shared-shaper Output Fields	330
	Table 42: show qos shared-shaper-control Output Fields	331

Table 43: show egress-queue events Output Fields	332
Table 44: show egress-queue rates Output Fields	335
Table 45: show fabric-queue Output Fields	337
Table 46: show statistics-profile Output Fields	338
Table 47: show qos interface-hierarchy Output Fields	339
Table 48: show qos-profile Output Fields	342
Table 49: show interfaces atm Output Fields	344
Table 50: show ip interface Output Fields	346
Table 51: show interfaces Output Fields	348
Table 52: show interfaces lag members Output Fields	349
Table 53: show qos-interface-set Output Fields	350
Table 54: show qos-interface-superset Output Fields	351
Table 55: show aaa qos downstream-rate Output Fields	352
Table 56: show qos-parameter Output Fields	354
Table 57: show qos-parameter-define Output Fields	355

About the Documentation

- E Series and JUNOS^e Documentation and Release Notes on page xxix
- Audience on page xxix
- E Series and JUNOS^e Text and Syntax Conventions on page xxix
- Obtaining Documentation on page xxxi
- Documentation Feedback on page xxxi
- Requesting Technical Support on page xxxi

E Series and JUNOS^e Documentation and Release Notes

For a list of related JUNOS^e documentation, see <http://www.juniper.net/techpubs/software/index.html> .

If the information in the latest release notes differs from the information in the documentation, follow the *JUNOS^e 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/>.

Audience

This guide is intended for experienced system and network specialists working with Juniper Networks E Series Broadband Services Routers in an Internet access environment.

E Series and JUNOS^e Text and Syntax Conventions

Table 1 on page xxx defines notice icons used in this documentation.

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 xxx defines text and syntax conventions that we use throughout the E Series and JUNOS documentation.

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents 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 the user must type.	host1(config)# traffic class low-loss1
Fixed-width text like this	Represents information as displayed on your terminal's screen.	<pre>host1#show ip ospf 2 Routing Process OSPF 2 with Router ID 5.5.0.250 Router is an Area Border Router (ABR)</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.	terminal length
<i>Italic text like this</i>	Represents variables.	<i>mask</i> , <i>accessListName</i>
(pipe symbol)	Represents a choice to select one keyword or variable to the left or to the right of this symbol. (The keyword or variable can be either optional or required.)	diagnostic line

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

Convention	Description	Examples
[] (brackets)	Represent optional keywords or variables.	[internal external]
[]* (brackets and asterisk)	Represent optional keywords or variables that can be entered more than once.	[level1 level2 l1]*
{ } (braces)	Represent required keywords or variables.	{ permit deny } { in out } { <i>clusterId</i> <i>ipAddress</i> }

Obtaining Documentation

To obtain the most current version of all Juniper Networks technical documentation, see the Technical Documentation page on the Juniper Networks Web site at <http://www.juniper.net/>.

To download complete sets of technical documentation to create your own documentation CD-ROMs or DVD-ROMs, see the Offline Documentation page at

<http://www.juniper.net/techpubs/resources/cdrom.html>

Copies of the Management Information Bases (MIBs) for a particular software release are available for download in the software image bundle from the Juniper Networks Web site at <http://www.juniper.net/>.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation to better meet your needs. 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

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/customers/support/downloads/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

QoS on the E Series Router

- Quality of Service Overview on page 3

Chapter 1

Quality of Service Overview

The quality of service (QoS) feature enables your E Series router to distinguish traffic with strict timing requirements from traffic that can tolerate delay, jitter, and loss.

QoS topics are discussed in the following sections:

- QoS on the E Series Router Overview on page 3
- QoS Audience on page 4
- QoS Platform Considerations on page 4
- QoS Terms on page 5
- QoS Features on page 7
- Configuring QoS on the E Series Router on page 8
- QoS References on page 9

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 Topics ■ For a list of related RFCs, see *Configuring QoS on the E Series Router* on page 8

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 Topics ■ For information about QoS users and QoS parameters, see *QoS Parameter Audience* on page 225

QoS Platform Considerations

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

Figure 1 on page 4 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 Topics** ■ For more information about supported interface types and specifiers on E Series routers, see [Interface Types and Specifiers](#).

QoS Terms

Table 3 on page 5 defines terms used in this discussion of QoS.

Table 3: 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).

Table 3: QoS Terminology *(continued)*

Term	Description
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.
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.

Table 3: QoS Terminology *(continued)*

Term	Description
WRED	Weighted random early detection congestion avoidance technique.

QoS Features

Table 4 on page 7 describes the major QoS features supported on the E Series router.

Table 4: 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 JUNOS Release 4.0 is deprecated and converted to QoS profiles and scheduler profiles.</p>

Table 4: QoS Features *(continued)*

Feature	Description
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.
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.

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 17.
4. (Optional) To configure RED or WRED, create a drop profile.
See “Dropping Behavior Overview” on page 25.
5. (Optional) To gather rate statistics, create a statistics profile.
See “QoS Statistics Overview” on page 37.
6. Configure a scheduler hierarchy with a scheduler profile.
See “Scheduler Hierarchy Overview” on page 45.
7. (Optional) Configure shaping:
 - Configure shaping and shared shaping using the scheduler profile.
See “Rate Shaping and Port Shaping Overview” on page 51, “Simple Shared Shaping Overview” on page 79, and “Compound Shared Shaping Overview” on page 99.
 - Configure shaping rates independent of the QoS profile and scheduler profile using QoS parameters.
See “Parameter Definition Attributes for QoS Administrators Overview” on page 229.
8. Create a QoS profile. QoS profiles reference queue, drop, statistics, and scheduler profiles.
See “Queuing and Buffer Management Overview” on page 17.
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 17.

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

Part 2

Classifying, Queuing, and Dropping Traffic

- Defining Service Levels with Traffic Classes and Traffic-Class Groups on page 13
- Configuring Queue Profiles for Buffer Management on page 17
- Configuring Dropping Behavior with RED and WRED on page 25
- Gathering Statistics for Rates and Events in the Queue on page 37

Chapter 2

Defining Service Levels with Traffic Classes and Traffic-Class Groups

This chapter provides information for configuring traffic classes and traffic-class groups on the E Series router.

QoS topics are discussed in the following sections:

- Traffic Class and Traffic-Class Groups Overview on page 13
- Configuring Traffic Classes That Define Service Levels on page 14
- Configuring Traffic-Class Groups That Define Service Levels on page 15
- Monitoring Traffic Classes and Traffic-Class Groups for Defined Levels of Service 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 Topics**
- Configuring Traffic Classes That Define Service Levels on page 14
 - Configuring Traffic-Class Groups That Define Service Levels on page 15

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 Topics**
- Monitoring Traffic Classes and Traffic-Class Groups for Defined Levels of Service on page 16
 - fabric-strict-priority
 - fabric-weight
 - traffic-class

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 Topics**
- Configuring Traffic Classes That Define Service Levels on page 14
 - Monitoring Traffic Classes and Traffic-Class Groups for Defined Levels of Service on page 16
 - traffic-class
 - traffic-class-group

Monitoring Traffic Classes and Traffic-Class Groups for Defined Levels of Service

To monitor traffic classes and traffic-class groups:

- Monitoring Service Levels with Traffic Classes on page 314
- Monitoring Service Levels with Traffic-Class Groups on page 315

Chapter 3

Configuring Queue Profiles for Buffer Management

This chapter provides information for configuring queue profiles for buffer management on the E Series router.

QoS topics are discussed in the following sections:

- Queuing and Buffer Management Overview on page 17
- Memory Requirements for Queue and Buffers on page 19
- Guidelines for Managing Queue Thresholds on page 19
- Guidelines for Managing Buffers on page 20
- Configuring Queue Profiles to Manage Buffers and Thresholds on page 22
- Monitoring Queues and Buffers on page 24

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 *JUNOS 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 19.

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 Topics**
- Configuring Queue Profiles to Manage Buffers and Thresholds on page 22
 - Guidelines for Managing Queue Thresholds on page 19

- Guidelines for Managing Buffers on page 20
- RED and WRED Overview on page 26

Memory Requirements for Queue and Buffers

JUNOS software uses 128-byte buffers.

The egress memory available for queues available depends on the ASIC and the line module. Table 5 on page 19 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 Topics**
- To identify the type of ASIC used by a line module, see the *ERX Module Guide* and the *E120 and E320 Module Guide*
 - Guidelines for Managing Queue Thresholds on page 19
 - Guidelines for Managing Buffers on page 20

Guidelines for Managing Queue Thresholds

To prevent the router from setting queue thresholds too low or too high, you can specify minimum and maximum queue thresholds. You can also specify the conformed length and exceeded length as percentages of the committed length.

Guidelines for Configuring a Maximum Threshold

We recommend that you constrain queue thresholds using committed or conformed threshold settings; any unused memory is redistributed to queues whose thresholds are not constrained. This use of thresholds is analogous to the way that shaping rates constrain bandwidth and cause bandwidth redistribution to unconstrained queues.

For example, voice queues are scheduled at strict priority; therefore, they require very little buffering. Configuring a maximum queue threshold enables the system to allocate more buffers to other queues in the system. Video queues are similar but

because they are higher bandwidth, they might require higher maximum committed thresholds.

You might want to limit latency of your multicast traffic by bounding the queue length using a maximum committed threshold. The following example configures the multicast queues so that the committed threshold never exceeds 20 KB, even when the egress memory is lightly loaded. The forfeited buffers are allocated to other queues.

```
host1(config)#queue-profile multicast
host1(config-queue)#committed-length 0 20000
host1(config-queue)#exit
```

Be sure to include 0 in the syntax, or you will configure a minimum threshold.

Guidelines for Configuring a Minimum Threshold

Configuring a minimum threshold does not guarantee that a queue always obtains the minimum buffer allocation. You can configure 1000 queues with a minimum of 1 MB each, but the buffer memory is 32 MB or 128 MB, not 1 GB. In this case, the system moves into higher operating regions (global utilization) if all these queues buffer traffic, until it reaches 90 percent utilization. At that point, the thresholds must reduce to the reserved percentages, and the queue thresholds drop from a high threshold to a very low one. Queues are not guaranteed to obtain any buffering, and are buffered in the order in which they are received.

You can configure a minimum committed threshold by specifying a value such as 1000 with the **committed-length** command:

```
host1(config)#queue-profile multicast
host1(config-queue)#committed-length 1000 20000
host1(config-queue)#exit
```

- Related Topics**
- Memory Requirements for Queue and Buffers on page 19
 - Configuring Queue Profiles to Manage Buffers and Thresholds on page 22

Guidelines for Managing Buffers

Queue profiles enable you to manage queue thresholds and buffers to manage the following common problems:

- Queues that back up and consume too many buffers
- Queues that cannot obtain buffers when they need them (called *buffer starvation*)

You can set the buffer weight to ensure that some sets of queues get higher thresholds than others. Buffer weight is analogous to weight in a scheduler profile. It directs the router to set the queue thresholds proportionately.

This feature provides graceful buffer allocation as the global utilization goes higher; queues with more buffer weight always obtain more buffers, but they do not undergo a dramatic drop in threshold when the system moves from region to region.

JUNOS software uses 128-byte buffers. When setting very small queue thresholds, keep the following guidelines in mind:

- Specifying a maximum queue length of 0 bytes disables queuing of packets on the queue.
- Specifying a maximum queue length of 1–128 bytes creates a single 128-byte buffer for the queue.
- Specifying a maximum queue length of 129–256 bytes creates two 128-byte buffers for the queue.
- Packets and cells consume at least one buffer.

For example, a 64-byte packet consumes a single 128-byte buffer. If you specify a maximum queue length of 256 bytes, then either two packets of 64–128 bytes in length or a single packet of 129–256 bytes can be queued.

For example, suppose a line module with 4000 IP interfaces is configured with four queues per IP interface, corresponding to four traffic classes. Suppose that queues in two of the traffic classes are configured with a buffer weight of 24 to increase burst tolerance. The following example configures the video queue:

```
host1(config)#queue-profile video
host1(config-queue)#buffer-weight 24
host1(config-queue)#exit
host1(config)#
```

When the egress memory is fully loaded, dynamic oversubscription is 0 percent, and the 8000 queues with the default buffer weight strictly partition 25 percent of the 32-MB memory, leaving 75 percent of the memory for the queues weighted 24 (corresponding to the ratio 75 percent:25 percent, or 24:8). Therefore, these queues have committed thresholds of 1 KB each, and queues with the buffer weight of 24 have committed thresholds of 3 KB each. As the egress memory becomes progressively less loaded, all the queue thresholds increase proportionally, based on dynamic oversubscription, but the queues with buffer weight 24 are always set with thresholds three times larger than the default thresholds.

Guidelines for Managing Buffer Starvation

Buffer starvation most commonly occurs when queues or nodes exist in a large round robin, usually in the default traffic-class group. When the round robin congests, the queues back up and require more buffers. The traffic in the round robin starts to burst based on a single node or queue. After a packet is dequeued, the node or queue can wait for thousands of other queues to dequeue a packet before it can dequeue again. During this time, the queue backs up.

If you configure different scheduler profile weights or assured rates for nodes in a large and congested round robin, the buffer starvation becomes apparent. The problem occurs when the heavy weighted nodes wait their turn in the round robin and thousands of other nodes dequeue. While the heavily weighted nodes wait, the

system needs to buffer them. However, all queues receive the same buffer allocation by default. If the system goes to higher buffer regions, it starts dropping packets for all queues. When the heavy weight node finally transmits, it dequeues all buffers, but it cannot dequeue the packets that were dropped. You do not achieve the expected bandwidth based on scheduler profile weights.

To manage buffer starvation, configure buffer weights on queues so they are in the same ratio as the expected bandwidth for the queues. For example, if two queues have scheduler weight (or assured-rate) in the ratio of 2:1, then set the buffer weights to the same ratio.

To manage buffer starvation, set the **maximum-committed-threshold** on queues that do not need buffering, and increase the **buffer-weight** for the heavily weighted queues in the round robin.

The system calculates the correct ratio for you. Issue the **show egress queue rates** command to see the ratio:

```
host1# show egress-queue rates brief interface fastEthernet 9/0.2
```

interface	traffic class	forwarded rate	aggregate drop rate	minimum rate	maximum rate
ip FastEthernet9/0.2	best-effort	0	0	25000	1000000
	videoTrafficClass	0	0	375000	1000000
	multicastTrafficClass	0	0	925000	1000000
	internetTrafficClass	0	0	50000	1000000
Total:		0	0		

```

Queues reported: 4
Queues filtered (under threshold): 0
Queues disabled (no rate period): 0
Queues disabled (no resources): 0
Total queues: 4

```

The minimum rate for each queue is the approximate rate the queue achieves if all configured queues in the line module run infinite traffic. Configure the buffer weights in proportion to the minimum rate displayed by the system.

- Related Topics**
- Memory Requirements for Queue and Buffers on page 19
 - Configuring Queue Profiles to Manage Buffers and Thresholds on page 22
 - Monitoring Forwarding and Drop Rates on the Egress Queue on page 333

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.

1. 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.

2. (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.

3. (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.

4. (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.

5. (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 Topics**
- [Queuing and Buffer Management Overview on page 17](#)
 - [Guidelines for Managing Queue Thresholds on page 19](#)
 - [Guidelines for Managing Buffers on page 20](#)
 - [Memory Requirements for Queue and Buffers on page 19](#)
 - [buffer-weight](#)
 - [committed-length](#)
 - [conformed-fraction](#)
 - [conformed-length](#)
 - [exceeded-fraction](#)
 - [exceeded-length](#)
 - [queue-profile](#)

Monitoring Queues and Buffers

To monitor queues and buffers, see:

- [Monitoring Queue Thresholds on page 316](#)
- [Monitoring Queue Profiles on page 319](#)

Chapter 4

Configuring Dropping Behavior with RED and WRED

This chapter provides information for configuring dropping behavior using RED and WRED on the E Series router.

QoS topics are discussed in the following sections:

- Dropping Behavior Overview on page 25
- RED and WRED Overview on page 26
- Configuring RED on page 27
- Example: Configuring Average Queue Length for RED on page 28
- Example: Configuring Dropping Thresholds for RED on page 29
- Example: Configuring Color-Blind RED on page 29
- Configuring WRED on page 31
- Example: Configuring Different Treatment of Colored Packets for WRED on page 33
- Example: Defining Different Drop Behavior for Each Traffic Class for WRED on page 33
- Example: Configuring WRED and Dynamic Queue Thresholds on page 34
- Monitoring RED and WRED on page 36

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 Topics ■ [Queuing and Buffer Management Overview on page 17](#)

RED and WRED Overview

The scheduler maintains an average queue length for each queue configured for RED. When a packet is enqueued, the current queue length is weighted into the average queue length based on the average-length exponent in the drop profile.

- Small exponent values weight the current queue length heavily, so the average queue length is more responsive to transient bursts.
- Large exponent values weight the current queue length lightly, so the average queue length is less responsive to bursts.

When the average queue length exceeds the minimum threshold, RED begins randomly dropping packets. While the average queue length increases toward the maximum threshold, RED drops packets with increasing frequency, up to the maximum drop probability. When the average queue length exceeds the maximum drop threshold, all packets are dropped. Figure 2 on page 27 shows this behavior.

Figure 2: Packets Dropped as Queue Length Increases

WRED is an extension of RED that allows you to assign different RED drop thresholds to each color of traffic. The router assigns a color to each packet. Committed means green, conformed means yellow, and exceeded means 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.

- Related Topics**
- Configuring RED on page 27
 - Configuring WRED on page 31

Configuring RED

Each line module supports a default drop profile and 15 configurable drop profiles. You can configure the default drop profile on all E Series line modules except for the ES2 10G LM.

To configure RED:

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 Topics**
- Configuring WRED on page 31
 - Monitoring RED and WRED on page 36
 - average-length-exponent
 - committed-threshold
 - conformed-threshold
 - drop-profile
 - exceeded-threshold

Example: Configuring Average Queue Length for RED

To enable calculation of average queue length, create a drop profile with a nonzero average-length exponent, reference the drop profile within a QoS profile, and attach the QoS profile to an interface.

The following drop profile enables the average queue length calculation, but does not initiate RED dropping behavior:

```
host1(config)#drop-profile averageOnly
host1(config-drop-profile)#average-length-exponent 10
```

- Related Topics**
- Configuring RED on page 27

- Dropping Behavior Overview on page 25
- RED and WRED Overview on page 26

Example: Configuring Dropping Thresholds for RED

You can specify different dropping behavior for committed (green), conformed (yellow), and exceeded (red) packets by specifying a minimum queue threshold, maximum queue threshold, and maximum drop probability for each color of traffic.

By default, conformed threshold and exceeded threshold take the same values as the committed threshold. Therefore, if you specify only a committed threshold, conformed and exceeded traffic is treated like committed traffic. Similarly, if you specify a conformed threshold without an exceeded threshold, exceeded traffic is treated like committed traffic.


The following drop profiles result in identical behavior:

```
host1(config)#drop-profile colorblind1
host1(config-drop-profile)#committed-threshold percent 30 90 5
host1(config-drop-profile)#exit
host1(config)#drop-profile colorblind2
host1(config-drop-profile)#committed-threshold percent 30 90 5
host1(config-drop-profile)#conformed-threshold percent 30 90 5
host1(config-drop-profile)#exit
host1(config)#drop-profile colorblind3
host1(config-drop-profile)#committed-threshold percent 30 90 5
host1(config-drop-profile)#conformed-threshold percent 30 90 5
host1(config-drop-profile)#exceeded-threshold percent 30 90 5
```

- Related Topics**
- Configuring RED on page 27
 - Dropping Behavior Overview on page 25
 - RED and WRED Overview on page 26

Example: Configuring Color-Blind RED

You can configure RED so that packets are dropped without regard to color. To do so, you combine a drop profile that has a committed threshold configured with a queue profile that specifies the same queue length for committed, conformed, and exceeded packets, as shown in Figure 3 on page 30.


Figure 3: Color-Blind RED Drop Profile with Colorless Queue Profile

In the following example, the drop profile and queue profile combine to specify the following:

- When the average queue length is between 30 percent full (30 KB) and 90 percent full (90 KB), up to 5 percent of the packets are randomly dropped regardless of their color.
- When the average queue length is greater than 90 percent, all packets are dropped regardless of color.

```
host1(config)#drop-profile nocolor
host1(config-drop-profile)#c omitted-threshold percent 30 90 5
host1(config-drop-profile)#exit
host1(config)#queue-profile colorless
host1(config-queue)#committed-length 100000 100000
host1(config-queue)#conformed-fraction 100
host1(config-queue)#exceeded-fraction 100
```

To achieve the same drop treatment for each color, you can specify color-blind RED in combination with a color-sensitive queue profile, as shown in Figure 4 on page 30.

Figure 4: Color-Blind RED Drop Profile with Color-Sensitive Queue Profile

In the following example, the drop profile and queue profile combine to specify the following:

- When the average queue length is between 30 percent full (30 KB) and 90 percent full (90 KB), up to 5 percent of the packets are dropped randomly. In this case, the maximum queue length is 100 KB for green packets, 50 KB for yellow packets, and 25 KB for red packets. Therefore, the router randomly drops:
 - Red packets when the average queue length is between 7.5 KB and 22.5 KB
 - Yellow packets when the average queue length is between 15 KB and 45 KB
 - Green packets when the average queue length is between 30 KB and 90 KB
- When the average queue length is greater than 90 percent of the maximum queue length, all packets are dropped. Therefore, the router drops:
 - Red packets when the average queue length is greater than 22.5 KB

- Yellow packets when the average queue length is greater than 45 KB
- Green packets when the average queue length is greater than 90 KB

```
host1(config)#drop-profile colorblindRed
host1(config-drop-profile)#committed-threshold percent 30 90 5
host1(config-drop-profile)#exit
host1(config)#queue-profile colorSensitive
host1(config-queue)#committed-length 100000 100000
```

- Related Topics**
- Configuring RED on page 27
 - Dropping Behavior Overview on page 25
 - RED and WRED Overview on page 26

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 Topics

- Configuring RED on page 27
- Monitoring RED and WRED on page 36
- average-length-exponent
- committed-threshold
- conformed-threshold
- drop-profile
- exceeded-threshold


Example: Configuring Different Treatment of Colored Packets for WRED

Figure 5 on page 33 shows a WRED drop profile that yields progressively more aggressive drop treatment for each color. Exceeded traffic is dropped over a wider range and with greater maximum drop probability than conformed or committed traffic. Conformed traffic is dropped over a wider range and with greater maximum drop probability than committed traffic.

The commands to configure this example are:

```
host1(config)#drop-profile wredColored
host1(config-drop-profile)#committed-threshold percent 30 90 3
host1(config-drop-profile)#conformed-threshold percent 25 90 5
host1(config-drop-profile)#exceeded-threshold percent 20 90 10
```

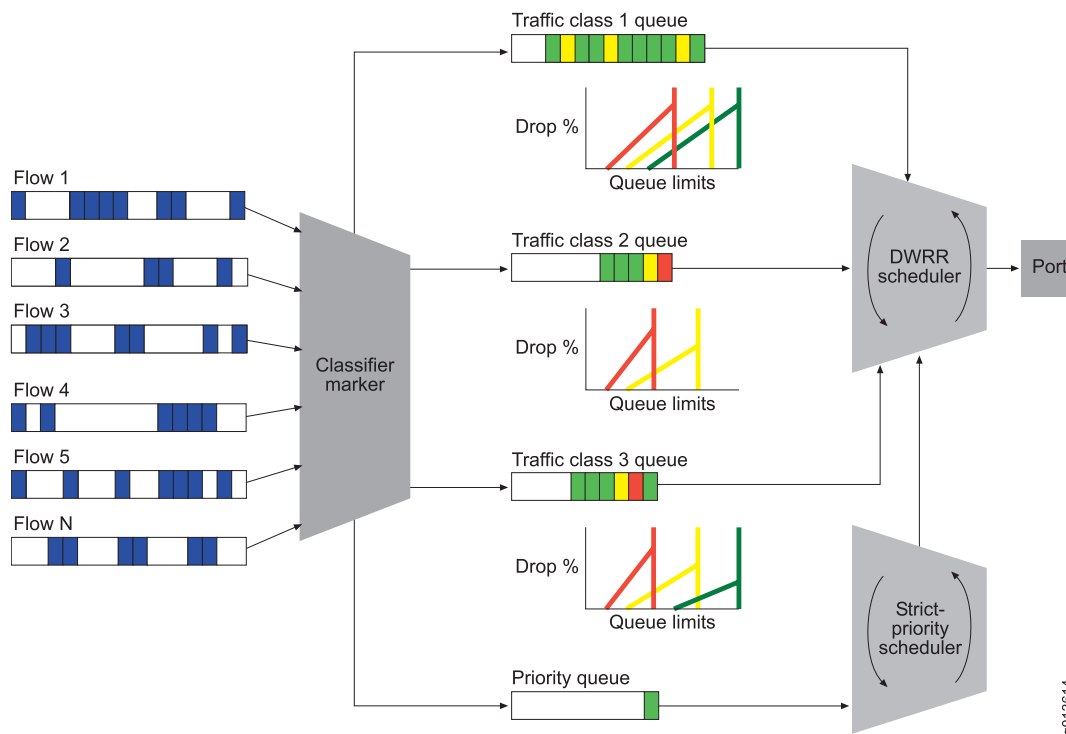
Figure 5: Different Treatment of Colored Packets



- Related Topics**
- Configuring WRED on page 31
 - Dropping Behavior Overview on page 25
 - RED and WRED Overview on page 26

Example: Defining Different Drop Behavior for Each Traffic Class for WRED

You can define different dropping behaviors for each traffic class in the router. By doing so, you can assign less aggressive drop profiles to higher-priority queues and more aggressive drop profiles to lower-priority queues. Figure 6 on page 34 shows an example that classifies packets into one of four traffic classes. Each traffic class has a different queueing behavior, drop treatment, and scheduler treatment.

Figure 6: Defining Different Drop Behavior for Each Queue

- Related Topics**
- Configuring WRED on page 31
 - Dropping Behavior Overview on page 25
 - RED and WRED Overview on page 26

Example: Configuring WRED and Dynamic Queue Thresholds

RED typically operates on fixed-size queues, and you can configure the router to use fixed-size queues. However, by default, the router employs dynamic queue thresholds to provide a good balance between sharing the egress buffer memory between queues and protecting an individual queue's claim on its fair share of the egress memory. Fixed-size queues become problematic as the number of configured queues scales into the thousands, because allocating disjointed partitions of buffer memory to each queue means the allocations become quite small, and most likely not all queues are simultaneously active.


In general, you use queues as follows:

- Fixed-size queues on core routers and core-facing interfaces where the number of queues is relatively small (tens or hundreds, but not thousands).
- Dynamic queues on edge-facing interfaces where the number of queues is relatively large (thousands).

As shown in Figure 7 on page 36, queue lengths extend to oversubscribe memory when aggregate memory utilization is low, and contract to strictly partition memory when memory utilization is high. Dynamic thresholding enforces fairness when free buffers are scarce and promotes sharing when buffers are plentiful. Dynamic queue thresholds are discussed in “Queuing and Buffer Management Overview” on page 17. Figure 7 on page 36 illustrates WRED behavior with dynamic queue thresholding.

To configure WRED to run on queues whose limits dynamically expand and contract, use the **percent** keyword when you configure thresholds in a drop profile. For example:

```
host1(config)#drop-profile internetDropProfile
host1(config-drop-profile)#average-length-exponent 9
host1(config-drop-profile)#committed-threshold percent 30 90 4
host1(config-drop-profile)#conformed-threshold percent 25 90 5
host1(config-drop-profile)#exceeded-threshold percent 20 90 6
```

Figure 7: WRED and Dynamic Queue Thresholding

- Related Topics**
- Configuring WRED on page 31
 - Dropping Behavior Overview on page 25
 - RED and WRED Overview on page 26

Monitoring RED and WRED

To monitor drop profiles, see:

- Monitoring Drop Profiles for RED and WRED on page 320

Chapter 5

Gathering Statistics for Rates and Events in the Queue

This chapter provides information for configuring statistics profiles on the E Series router.

QoS topics are discussed in the following sections:

- QoS Statistics Overview on page 37
- Configuring Statistic Profiles for QoS on page 39
- Configuring Rate Statistics on page 39
- Configuring Event Statistics on page 40
- Clearing QoS Statistics on the Egress Queue on page 41
- Clearing QoS Statistics on the Fabric Queue on page 42
- Monitoring QoS Statistics for Rates and Events on page 42

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 *JUNOS System Basics Configuration Guide, Chapter 4, Configuring SNMP*.

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 Topics**
- Configuring Rate Statistics on page 39
 - Configuring Event Statistics on page 40
 - Monitoring QoS Statistics for Rates and Events on page 42
 - statistics-profile

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 Topics**
- Configuring Statistic Profiles for QoS on page 39
 - Configuring a QoS Profile on page 132
 - Monitoring QoS Statistics for Rates and Events on page 42
 - interface
 - qos-profile
 - queue
 - rate-period
 - statistics-profile

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 Topics**
- Configuring Statistic Profiles for QoS on page 39
 - Configuring a QoS Profile on page 132
 - Monitoring QoS Statistics for Rates and Events on page 42
 - committed-drop-threshold
 - conformed-drop-threshold
 - exceeded-drop-threshold
 - forwarding-rate-threshold
 - qos-profile
 - queue
 - rate-period
 - statistics-profile

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 Topics**
- Monitoring QoS Statistics for Rates and Events on page 42
 - clear egress-queue

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 Topics**
- Monitoring QoS Statistics for Rates and Events on page 42
 - clear fabric-queue

Monitoring QoS Statistics for Rates and Events

To monitor statistics for rates and events in the queue:

- Monitoring Forwarding and Drop Events on the Egress Queue on page 331
- Monitoring Forwarding and Drop Rates on the Egress Queue on page 333
- Monitoring Queue Statistics for the Fabric on page 337
- Monitoring the Configuration of Statistics Profiles on page 338

Part 3

Scheduling and Shaping Traffic

- QoS Scheduler Hierarchy Overview on page 45
- Configuring Rates and Weights in the Scheduler Hierarchy on page 51
- Configuring Strict-Priority Scheduling on page 59
- Shared Shaping Overview on page 71
- Configuring Simple Shared Shaping of Traffic on page 79
- Configuring Variables in the Simple Shared Shaping Algorithm on page 89
- Configuring Compound Shared Shaping of Traffic on page 99
- Configuring Implicit and Explicit Constituent Selection for Shaping on page 107
- Monitoring a QoS Scheduler Hierarchy on page 123

Chapter 6

QoS Scheduler Hierarchy Overview

This chapter provides information for configuring the QoS scheduler hierarchy using scheduler profiles on the E Series router.

QoS topics are discussed in the following sections:


- Scheduler Hierarchy Overview on page 45
- Configuring a Scheduler Hierarchy on page 47
- Configuring a Scheduler Profile for a Scheduler Node or Queue on page 48
- Using Expressions for Bandwidth and Burst Values in a Scheduler Profile on page 48

Scheduler Hierarchy Overview

The egress line module scheduler is an HRR scheduler. Figure 8 on page 46 is an example of a QoS scheduler's hierarchy.

As shown in Figure 8 on page 46, 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 8: 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 Topics**
- Static and Hierarchical Assured Rate Overview on page 54
 - Rate Shaping and Port Shaping Overview on page 51
 - Shared Shaping Overview on page 71
 - Configuring a Scheduler Hierarchy on page 47

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” on page 48.
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” on page 52 or “Configuring Port Shaping” on page 53.
 - Configure an assured rate.
See “Configuring an Assured Rate for a Scheduler Node or Queue” on page 55.
 - Configure the HRR weight.
See “Configuring the HRR Weight for a Scheduler Node or Queue” on page 57.
 - Configure shared shaping.
See “Configuring Simple Shared Shaping” on page 81 and “Configuring Compound Shared Shaping” on page 100.
 - Configure implicit and explicit constituent selection.
See “Configuring Implicit Constituents for Simple or Compound Shared Shaping” on page 114 and “Configuring Explicit Constituents for Simple or Compound Shared Shaping” on page 120.

3. Reference the scheduler profile in a QoS profile and apply to an interface.
See “Configuring a QoS Profile” on page 132 and “Attaching a QoS Profile to an Interface” on page 134.

- Related Topics**
- Scheduler Hierarchy Overview on page 45
 - For information about configuring a scheduling hierarchy with QoS parameters, see Parameter Definition Attributes for QoS Administrators Overview on page 229

Configuring a Scheduler Profile for a Scheduler Node or Queue

To create a scheduler profile for a scheduler hierarchy:

- Create a scheduler profile by assigning a name that represents the type of service and enter Scheduler Profile Configuration mode.

```
host1(config)#scheduler-profile sp-1mbs
host1(config-scheduler-profile)#
```

The router supports up to 1000 scheduler profiles.

- Related Topics**
- Configuring Rate Shaping for a Scheduler Node or Queue on page 52
 - Configuring Port Shaping on page 53
 - Configuring an Assured Rate for a Scheduler Node or Queue on page 55
 - Configuring the HRR Weight for a Scheduler Node or Queue on page 57
 - Configuring Simple Shared Shaping on page 81
 - Configuring Compound Shared Shaping on page 100

Using Expressions for Bandwidth and Burst Values in a Scheduler Profile

Expressions are combinations of constants and operators. You can specify some scheduler profile attributes using an expression, such as the shaping rate. All operations within expressions are performed using 64 bit unsigned math, resulting in a 32 bit, signed integer value.

Expressions consist of both operators and operand values. Operators are mathematical functions, and operand values are the inputs for the mathematical function. Operand values can be an integer. You specify an expression consisting of an operand, followed by zero or more [operator, operand] pairs.

You can specify bandwidth as a percentage and burst in milliseconds or bytes by using expressions with the **shaping-rate**, **shared-shaping-rate**, **assured-rate**, and **weight** commands.

When calculating constant shaping rates, use the following formula to translate burst values from bytes to milliseconds (ms):

$$\text{Time (ms)} = \text{Rate (bps)} \times 1000 \text{ (ms/s)} / (\text{burstValueBytes} \times 8 \text{ bits/byte})$$

Using this formula, a 2 Mbps service with a 500 KB burst yields:

$$(2000000 \times 1000) / (50000 \times 8) = 500 \text{ ms}$$

The shaping rate is calculated when the QoS profile is attached based on the parameter instance. For example:

```
host1(config)# scheduler-profile sp-1mbs
(config-scheduler-profile)# shaping-rate video-bandwidth % 100 burst 500 milliseconds
```

When the shaping rate for video-bandwidth is 2 Mbps, the burst value is calculated using the following formula:

$$\text{Burst Value (bits)} = \text{Rate (bps)} \times 1000 \text{ (ms/s)} / \text{Time (ms)}$$

The burst value in bits is calculated as:

$$\text{Burst Value (bits)} = 2000000 \times 1000 / 500 = 4000000$$

The burst value in bytes is calculated as:

$$\text{Burst Value (bytes)} = 4000000 / 8 = 500000$$

- Related Topics**
- For more information about using expressions within scheduler profiles that are used for QoS parameters, see *Scheduler Profiles and Parameter Expressions for QoS Administrators* on page 235
 - *Configuring Rate Shaping for a Scheduler Node or Queue* on page 52
 - *Configuring Port Shaping* on page 53
 - *Configuring an Assured Rate for a Scheduler Node or Queue* on page 55
 - *Configuring the HRR Weight for a Scheduler Node or Queue* on page 57
 - *Configuring Simple Shared Shaping* on page 81
 - *Configuring Compound Shared Shaping* on page 100

Chapter 7

Configuring Rates and Weights in the Scheduler Hierarchy

This chapter provides information for configuring shaping rates, assured rates, and weights in the QoS scheduler hierarchy using scheduler profiles.


QoS topics are discussed in the following sections:

- Rate Shaping and Port Shaping Overview on page 51
- Configuring Rate Shaping for a Scheduler Node or Queue on page 52
- Configuring Port Shaping on page 53
- Static and Hierarchical Assured Rate Overview on page 54
- Configuring an Assured Rate for a Scheduler Node or Queue on page 55
- Configuring the HRR Weight for a Scheduler Node or Queue on page 57

Rate Shaping and Port Shaping Overview

Rate shaping throttles the rate at which queues transmit packets. Rate shaping is TCP friendly; that is, it buffers packets that are above the rate, rather than dropping them.

Port shaping enables you to shape the aggregate traffic through a port or channel to a rate that is less than the line or port rate. With port shaping, you can configure scheduler nodes at the port level, as shown in Figure 9 on page 52.

Figure 9: Port Shaping on an Ethernet Module

The per-port shaping feature provides the ability to shape the output of a port.

- Related Topics**
- Configuring Rate Shaping for a Scheduler Node or Queue on page 52
 - Configuring Port Shaping on page 53

Configuring Rate Shaping for a Scheduler Node or Queue

The router supports 64,000 rate shapers per line module. Shaping rates are multiples of 1 Kbps.

To configure a shaping rate for a scheduler node or queue:

1. Create a scheduler profile.

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

2. Specify a shaping rate in the scheduler profile.

```
host1(config-scheduler-profile)#shaping-rate 128000 burst 32767 milliseconds
host1(config-scheduler-profile)#shaping-rate 5000 x 90
```

The range for the shared-shaping rate is 1000–1000000000 bps (1 Kbps–1000 Kbps); the default is the minimum shaping rate (1 Kbps). The router rounds the rate to the next higher 8 Kbps.

Use the `operator` and `operandValue` variables to configure a shaping rate with an expression.

You can use the **bps** or **kbps** keywords to specify the unit of the shaping rate. By default, the shaping rate is configured in bps.

Use the **burst** keyword to specify the catch-up number associated with the shaper; the range is 0–522240. Specifying 0 enables the router to select an applicable default value.

Use the **milliseconds** or **bytes** keywords to specify the unit of the burst size.

- Related Topics**
- Rate Shaping and Port Shaping Overview on page 51
 - Configuring a Scheduler Profile for a Scheduler Node or Queue on page 48
 - scheduler-profile
 - shaping-rate

Configuring Port Shaping

To configure port-shaping:

1. Configure the scheduler profile and the shaping rate.

```
host1(config)#scheduler-profile 80mbps
host1(config-scheduler-profile)#shaping-rate 80000000
host1(config-scheduler-profile)#exit
```

2. Configure a QoS profile, specify the **node** command, and reference the scheduler-profile.

```
host1(config)#qos-profile 80mbps
host1(config-qos-profile)#ethernet node scheduler-profile 80mbps
host1(config-qos-profile)#exit
```

3. Attach the QoS profile to the port.

```
host1(config)#interface fastethernet 2/0
host1(config-if)#qos-profile 80mbps
```

The sample configuration shapes Fast Ethernet port 2/0 to a rate no higher than 80 Mbps.

Using the following configuration, you can shape the corresponding HDLC channel down to 20 Mbps:

```
host1(config)#scheduler-profile 20mbps
host1(config-scheduler-profile)#shaping-rate 20000000
host1(config-scheduler-profile)#exit
host1(config)#qos-profile 20mbps
host1(config-qos-profile)#serial node scheduler-profile 20mbps
host1(config-qos-profile)#exit
host1(config)#interface serial 2/0:1/1
host1(config-if)#qos-profile 20mbps
```

- Related Topics**
- Rate Shaping and Port Shaping Overview on page 51
 - Configuring a Scheduler Profile for a Scheduler Node or Queue on page 48
 - For more information about specifying an expression that you can reference within a scheduler profile, see Using Expressions for Bandwidth and Burst Values in a Scheduler Profile on page 48
 - node

- qos-profile
- scheduler-profile
- shaping-rate

Static and Hierarchical Assured Rate Overview

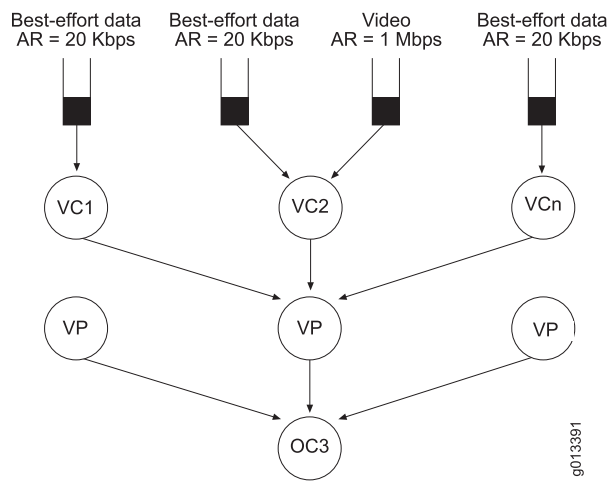
You can configure the effective weight of the scheduler node or queue by configuring a static assured rate or a hierarchical assured rate (HAR). The JUNOS hierarchical assured rate (HAR) feature provides a more powerful and efficient method of configuring assured rates than static assured rates.

When you use static assured rates, a queue is guaranteed to receive its assured rate only when its parent node is configured with an assured rate that equals the sum of all its child assured rates. Therefore, to ensure that a queue receives its specified assured rate, you must frequently recalculate the assured rates on all parent nodes in the queue's hierarchy. This recalculation is necessary because of the number of scheduler nodes and queues that may be dynamically created or deleted through applications such as bandwidth-on-demand. Eventually, this complicated manual recalculation process becomes unreasonable and virtually impossible.

HAR replaces the manual recalculation process by directing the router to dynamically calculate the assured rate for a scheduler node based on the sum of the assured rates of all its child nodes and queues. For example, you might use HAR to increase the effective weight of an ATM-VC scheduler node when a video queue is created, and to later restore the effective rate of the node when the video queue is deleted.

HAR is applicable only to level 1 and level 2 scheduler nodes, and is not applicable to queues or ports. When you configure HAR, the changes take place immediately. When you disable HAR, the scheduler node's previous weight is restored.

Figure 10 on page 55 shows an application of HAR for VC nodes. In the example, VCs, which are configured for HAR, are stacked over virtual path (VP) nodes. The VP nodes are in turn stacked over an OC-3 ATM port. Each VC has a best-effort data queue, which currently has an assured rate of 20 Kbps. The VCs share equal portions of their parent VP's bandwidth. However, when the video queue is added to VC2, HAR enables VC2's share of the VP bandwidth to increase in proportion to the 1-Mbps video queue that was created. The bandwidth of sibling VC nodes, which have only a data queue, is decreased in equal proportions.

Figure 10: Hierarchical Assured Rate

- Related Topics**
- Configuring an Assured Rate for a Scheduler Node or Queue on page 55
 - Configuring the HRR Weight for a Scheduler Node or Queue on page 57

Configuring an Assured Rate for a Scheduler Node or Queue

You can configure the effective weight of the scheduler node or queue by configuring a static assured rate or a hierarchical assured rate (HAR). HAR dynamically adjusts the available bandwidth for a scheduler node based on the creation and deletion of other scheduler nodes.

By default, the HRR weight is configured for the scheduler profile. If the assured rate setting is other than none (the default), then the assured rate is used instead of the HRR weight setting for the scheduler node or queue.

Tasks to configure an assured rate are:

- Configuring a Static Assured Rate on page 55
- Configuring a Hierarchical Assured Rate on page 56
- Changing the Assured Rate to an HRR Weight on page 56

Configuring a Static Assured Rate

To configure a static assured rate:

1. Create a scheduler profile.

```
host1(config)#scheduler-profile static
host1(config-scheduler-profile)#
```

2. Specify a numeric rate with the **assured-rate** command in the scheduler profile.

```
host1(config-scheduler-profile)#assured-rate 56000
```

```
host1(config-scheduler-profile)#assured-rate 50000 - 31000
```

For a static assured rate, specify the bits per second value in the range 25000–1000000000 bps (25 Kbps to 1 Gbps); the default is none (no assured rate).

Use the operator and operandValue variables to configure an assured rate with an expression.

Configuring a Hierarchical Assured Rate

To specify that the HAR is used for scheduler nodes (HAR is not used for queues or ports):

1. Create a scheduler profile.

```
host1(config)#scheduler-profile har
host1(config-scheduler-profile)#
```

2. Specify the **hierarchical** keyword with the **assured-rate** command in the scheduler profile.

```
host1(config-scheduler-profile)#assured-rate hierarchical
```

Changing the Assured Rate to an HRR Weight

To change an assured rate to an HRR weight:

1. Create a scheduler profile.

```
host1(config)#scheduler-profile static
host1(config-scheduler-profile)#
```

2. Delete the configured assured rate.

```
host1(config-scheduler-profile)#no assured-rate
```

The assured rate in the scheduler profile reverts to using the HRR weight specification.

- Related Topics**
- Static and Hierarchical Assured Rate Overview on page 54
 - Configuring a Scheduler Profile for a Scheduler Node or Queue on page 48
 - Configuring the HRR Weight for a Scheduler Node or Queue on page 57
 - For more information about specifying an expression that you can reference within a scheduler profile, see Using Expressions for Bandwidth and Burst Values in a Scheduler Profile on page 48
 - assured-rate
 - scheduler-profile

Configuring the HRR Weight for a Scheduler Node or Queue

By default, the HRR weight is configured for the scheduler profile. You can set a specific HRR weight of the scheduler node or queue. The weight value is used when no assured rate is set.

To configure a static weight:

1. Create a scheduler profile.

```
host1(config)#scheduler-profile relative
host1(config-scheduler-profile)#
```

2. Specify the weight value.

```
host1(config-scheduler-profile)#weight 10
host1(config-scheduler-profile)#weight 800 - 200
```

The weight value is in the range 0–4080. The default weight is 8. Weight 0 (zero) is a special weight that is used for relative strict-priority scheduling.

Use the operator and operandValue variables to configure a weight with an expression.

- Related Topics**
- Static and Hierarchical Assured Rate Overview on page 54
 - For more information about specifying an expression that you can reference within a scheduler profile, see Using Expressions for Bandwidth and Burst Values in a Scheduler Profile on page 48
 - Relative Strict-Priority Scheduling Overview on page 60
 - scheduler-profile
 - weight

Chapter 8

Configuring Strict-Priority Scheduling

This chapter provides information for configuring strict-priority scheduling.

QoS topics are discussed in the following sections:

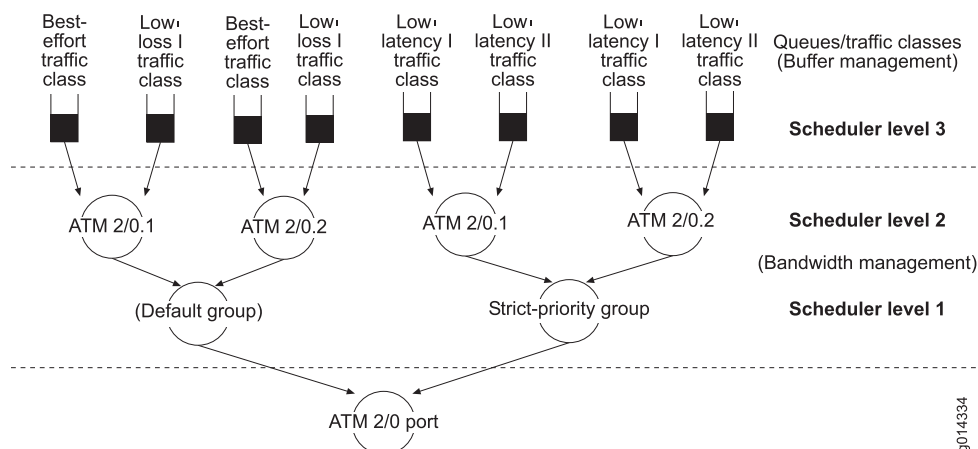
- Strict-Priority and Relative Strict-Priority Scheduling Overview on page 59
- Comparison of True Strict Priority with Relative Strict Priority Scheduling on page 61
- Configuring Strict-Priority Scheduling on page 66
- Configuring Relative Strict-Priority Scheduling for Aggregate Shaping Rates on page 68

Strict-Priority and Relative Strict-Priority Scheduling Overview

You can configure one or more strict-priority queues per interface. Strict-priority scheduling is implemented with a special strict-priority scheduler node that is stacked directly above the port. Queues stacked on top of the strict-priority scheduler node always get bandwidth before other queues.

You can configure only one node at the first scheduler level as strict priority. If any node or queue above the strict-priority node has packets, it is scheduled next. If multiple queues above the strict-priority node have packets, the HRR algorithm selects which strict-priority queue is scheduled next.

Figure 11 on page 60 illustrates an example of a QoS scheduler's hierarchy.

Figure 11: Sample Strict-Priority Scheduling Hierarchy

One strict priority traffic-class group is called the auto-strict-priority group. The scheduler nodes and queues in the auto-strict-priority group receive strict-priority scheduling. If multiple queues above the strict-priority node have packets, the HRR algorithm selects which strict-priority queue is scheduled next.



NOTE: If you configured traffic shaping through traffic shape profiles in JUNOS releases before Release 4.0, traffic shaping is replaced with the rate-shaping feature, which is configured when you configure a scheduler profile.

Relative Strict-Priority Scheduling Overview

Relative strict-priority scheduling provides strict-priority scheduling within a shaped aggregate rate. For example, it allows you to 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.

Relative strict priority differs from true strict priority in that it can implement the aggregate shaping rate for both strict and nonstrict traffic. With true strict priority, you can shape the nonstrict or the strict traffic separately, but you cannot shape the aggregate to a single rate.

The best application of relative strict priority is on Ethernet, where you can shape the aggregate for each VLAN to a specified rate, and provision a strict and nonstrict queue for each VLAN above the shaped VLAN node.

To use relative strict priority, you configure strict-priority queues above the VC or VLAN scheduler node, thereby providing for strict-priority scheduling of the queues within the VC or VLAN. You configure relative strict priority without using QoS traffic-class groups, which causes strict-priority queues to appear in the same scheduler hierarchy as the nonstrict queues.

Relative strict priority provides low latency only if you undersubscribe the port by shaping all VCs on the port so that the sum of the shaping rates is less than the port

rate. The port will not become congested, and the latency caused by the round-robin behavior of both the HRR and cell schedulers is nominal. In these undersubscribed conditions, the latency of a strict-priority queue within each VC is calculated as if the VC were draining onto a wire with bandwidth equal to the shaped rate.

Relative strict priority is carried out in the HRR scheduler on E Series ASIC line modules.


- Related Topics**
- Comparison of True Strict Priority with Relative Strict Priority Scheduling on page 61
 - Configuring Strict-Priority Scheduling on page 66
 - Configuring Relative Strict-Priority Scheduling for Aggregate Shaping Rates on page 68

Comparison of True Strict Priority with Relative Strict Priority Scheduling

This section explains how the HRR and SAR schedulers handle true strict-priority and relative strict-priority configurations.

Schedulers and True Strict Priority

In the strict-priority configuration in Figure 12 on page 62, the queues stacked above the single strict priority scheduler node make up a round-robin separate from the nonstrict queues. All strict queues are drained to completion first, and any residual bandwidth is allocated to the nonstrict round-robin.


Figure 12: True Strict-Priority Configuration

This configuration provides low latency for the strict-priority queues, irrespective of the state of the nonstrict queues. The worst-case latency for a strict packet caused by a nonstrict packet is the propagation delay of a single large packet at the port rate. For a 1500 byte frame at OC3 rate, that latency is less than 100 microseconds.

Because the strict and nonstrict packets for a VC are scheduled in separate round robins, the scheduler cannot enforce an aggregate rate for both of them.

Schedulers and Relative Strict Priority

In the relative strict-priority configuration in Figure 13 on page 63, the scheduler provides relative strict-priority scheduling relative to the VC. If the port is not oversubscribed, the VC round robin does not cause significant latency.

Figure 13: Relative Strict-Priority Configuration

This configuration provides a latency bound for the relative strict-priority queues. The worst-case latency caused by a nonstrict packet is the propagation delay of a single large packet at the VC rate. For a 1500 byte frame at a 2 Mbps rate, that delay is about 6 milliseconds.

This configuration provides for shaping the aggregate of nonstrict and relative strict packets to a single rate, and it is consistent with the traditional ATM model. It does not scale as well as true strict priority, because the nonstrict and relative strict traffic together must not oversubscribe the port rate.

Relative Strict Priority on ATM Modules

You can use relative strict priority on any type of E Series line module; however, on ATM line modules you have an alternative. On ATM line modules you can configure true strict-priority queues in the HRR scheduler and shape the aggregate for the VC in the SAR scheduler. VC backpressure affects only the nonstrict traffic for the VC. For this type of configuration, you should shape the relative strict traffic for each VC in the HRR scheduler to a rate that is less than the aggregate VC rate. This shaping prevents the VC queue in the SAR scheduler from being congested with strict-priority traffic.

The major difference between relative and true strict priority on ATM line modules is that relative strict priority shapes the aggregate for the VC to a pre-cell tax rate, whereas true strict priority shapes the aggregate for the VC to a post-cell tax rate. For example, shaping the VC to 1 Mbps in the HRR scheduler allows 1 Mbps of frame data, but cell tax adds anywhere from 100 Kbps to 1 Mbps additional bandwidth, depending on packet size. Shaping the VC to 1 Mbps in the SAR scheduler allows just 1 Mbps of cell bytes regardless of packet size.

Oversubscribing ATM Ports

You cannot oversubscribe ATM ports and still achieve low latency with relative strict-priority scheduling. There are several ways to ensure that ports are not oversubscribed. The most common is to use a per-VC scheduler by configuring the HRR scheduler with either ATM VP or VC node shaping (using the **atm-vp node** or **atm-vc node** commands), and setting the sum of the shaping rates less than the port rate. In these scenarios, the cell residency in the SAR scheduler is minimal, and cell scheduling does not interfere with relative strict priority.

Minimizing Latency on the SAR Scheduler

There are two methods you can use to control latency on the SAR scheduler. In the first method, you set the ATM QoS port mode to low-latency mode. In low-latency mode, the HRR scheduler controls scheduling, buffering in the SAR scheduler is limited, and latency caused by the SAR scheduler is minimized.

You can also use the default **no qos-mode-port** mode of SAR operation to minimize the latency induced by the SAR. In this method, you set **qos shaping-mode** cell and shape an OC-3 ATM port to 149 Mbps, or an OC-12 ATM port to 600 Mbps. By throttling the rate at which the HRR scheduler delivers packets to the SAR, you bound SAR buffering and latency. This approach retains the flexibility to configure different ATM QoS in the SAR, including shaped VP tunnels, UBR + PCR, nrtVBR, and CBR services.

To set the SAR mode, use the **qos-mode-port** command. For more information about operational modes on ATM interfaces, see “ATM Integrated Scheduler Overview” on page 159.



NOTE: Controlling latency is not normally required. If you undersubscribe the port rate in the HRR scheduler, you can obtain latency bounds without modifying the SAR mode of operation.

HRR Scheduler Behavior and Strict-Priority Scheduling

The HRR scheduler does not offer native strict-priority scheduling above the first scheduler level in the hardware; however, you can configure very large weights in the round robin in the HRR scheduler to obtain approximate strict-priority scheduling. Note that under conditions of low VC bandwidth and large packet sizes, latency and jitter increase because of the inherent propagation delay of large packets over a small shaping rate. The following sections describe additional configuration steps that will ensure that no more than a single nonstrict packet can precede a strict-priority packet on the VC.

Zero-Weight Queues

To reduce latency and jitter, you can configure the relative strict-priority queue with a weight of 0 (zero), which gives the queue a weight of 4080. When a packet arrives at a zero-weighted queue, the queue remains in the active WRR until it is exhausted,

whereas competing queues must leave the active WRR because their weight credits are exhausted. To completely drain the queue, configure the maximum burst size. The zero-weighted queue is eventually alone in the active round robin and is effectively drained at strict priority.

To configure more than one relative strict queue or node, simply configure a maximum weight, and the two relative strict queues or nodes will share bandwidth fairly. You can shape the nonstrict queue, as described in the next section, to keep latency bounded.

Also, configure only a few nonstrict nodes or queues to prevent additional latency and jitter of the relative strict-priority traffic when the nodes or queues are in the round robin and a packet arrives in the zero-weighted queue. The number of nonstrict frames that precede a relative strict frame equals the number of nonzero weighted queues among the sibling scheduler nodes.

Nonstrict queues must still exhaust their weight credits before they leave the active round robin. The result is that occasionally more than one nonstrict frame may precede a relative strict frame, causing more jitter than may be acceptable. You can eliminate this source of latency by shaping the nonstrict queue to the aggregate rate with a burst size of 1.

Setting the Burst Size in a Shaping Rate

The burst value in a shaping rate determines the number of rate credits that can accrue when the queue or scheduler node is held in the inactive round robin. When the queue is back on the active list, the accrued credits allow the queue or node to catch up to the configured rate, up to the burst value.

Normally, the burst size is several packet lengths to allow a queue deprived of bandwidth because of congestion to catch up to its rate. Larger burst sizes allow more bursting to allow the queue to attain its shaped rate under bursty congestion scenarios.

Special Shaping Rate for Nonstrict Queues

To remove additional jitter, you can configure the nonstrict queue with a special shaping rate that causes the hardware to temporarily eject the queue from the active round robin whenever it sends a frame. The result is that at most one nonstrict frame can precede a relative strict-priority frame. The special shaping rate is the same rate as the aggregate rate, but with a configured burst size of 1.

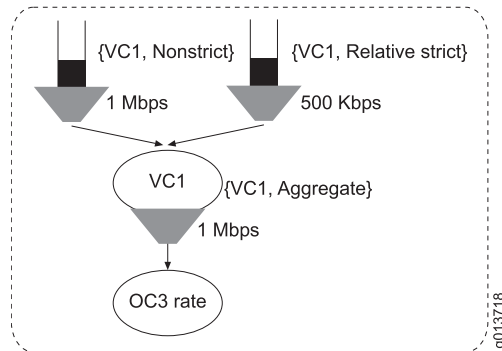
You can still configure a shaping rate for the zero-weighted queue or node. This is useful for limiting starvation of the nonstrict traffic in the aggregate.

In Figure 14 on page 66, the VC node is shaped in the HRR scheduler to 1 Mbps to limit the aggregate traffic for the subscriber. The relative strict traffic is shaped to 500 Kbps. This shaping limits relative strict traffic to 500 Kbps, and prevents the relative strict-priority traffic from starving out the nonstrict traffic.

The third shaper, on the nonstrict queue, is subtle. The rate is 1 Mbps, which allows the nonstrict traffic to consume up to the full aggregate rate of the VC. But the burst size is 1, which causes the nonstrict queue to always yield to the relative strict-priority

queue after sending a packet. This burst size limits the number of nonstrict packets that can precede a relative strict-priority packet to the minimum, one packet.

Figure 14: Tuning Latency on Strict-Priority Queues



- Related Topics**
- Strict-Priority and Relative Strict-Priority Scheduling Overview on page 59
 - Configuring Strict-Priority Scheduling on page 66
 - Relative Strict-Priority Scheduling Overview on page 60

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.

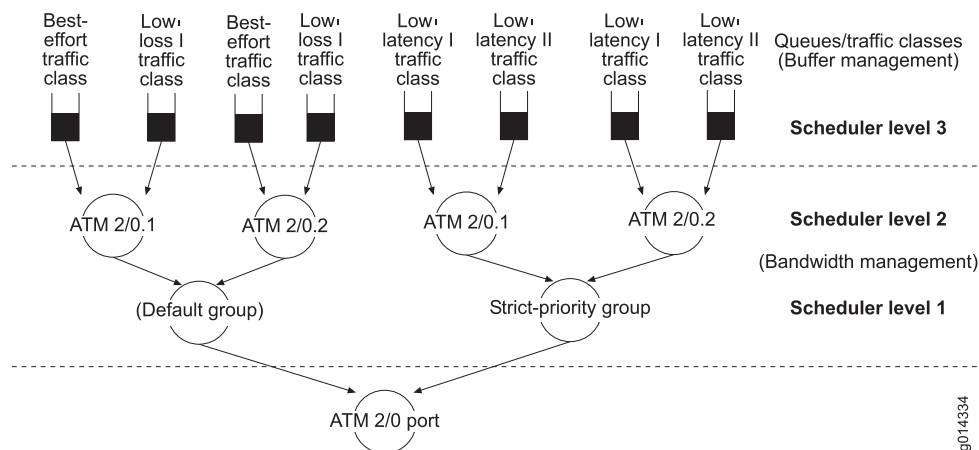
```

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 15 on page 67.

Figure 15: Sample Strict-Priority Scheduling Hierarchy



- Related Topics**
- Strict-Priority and Relative Strict-Priority Scheduling Overview on page 59
 - For more information about specifying an expression that you can reference within a scheduler profile, see Using Expressions for Bandwidth and Burst Values in a Scheduler Profile on page 48
 - group
 - node
 - qos-profile
 - queue
 - scheduler-profile
 - shaping-rate

- strict-priority
- traffic-class
- traffic-class-group

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.

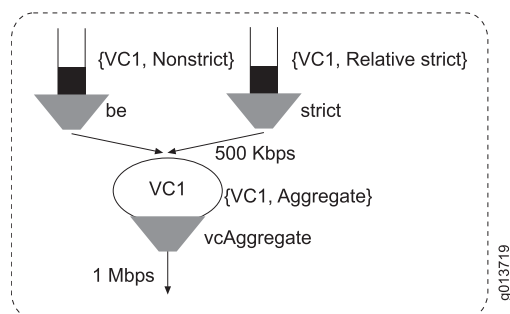
3. Create a scheduler profile for the aggregate bandwidth.

```
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 16 on page 69.

Figure 16: Sample Relative Strict-Priority Scheduler Hierarchy

- Related Topics**
- Strict-Priority and Relative Strict-Priority Scheduling Overview on page 59
 - For more information about specifying an expression that you can reference within a scheduler profile, see Using Expressions for Bandwidth and Burst Values in a Scheduler Profile on page 48
 - node
 - qos-profile
 - scheduler-profile
 - shaping-rate
 - weight

Chapter 9

Shared Shaping Overview

This chapter provides information for configuring shared shaping of traffic on the E Series router.

QoS topics are discussed in the following sections:

- Shared Shaping Overview on page 71
- Shared Shaper Terms on page 72
- How Shared Shaping Works on page 72
- Guidelines for Configuring Simple and Compound Shared Shaping on page 74

Shared Shaping Overview

In the JUNOS 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 Topics**
- Simple Shared Shaping Overview on page 79
 - Compound Shared Shaping Overview on page 99

Shared Shaper Terms

Table 6 on page 72 defines terms used in this discussion of shared shaping.

Table 6: Shared Shaper Terminology Used in This Chapter

Term	Description
Constituent	Scheduler node or queue associated with a logical interface. A shared shaper is configured for a logical interface; all queues and scheduler nodes associated with that logical interface are constituents of the shared shaper.
Active constituent	Constituent that is monitored or controlled by the shared shaper mechanism.
Inactive constituent	Constituent that is ignored by the shared shaper mechanism. Inactive constituents can be indirectly controlled; for example, queues stacked above a node that is an active constituent.
Shared 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.
Implicit shared shaper	Shared shaper where the system automatically selects the active constituents. The system selects scheduler nodes as active; queues above nodes remain inactive.
Explicit shared shaper	Shared shaper where you select the active constituents by issuing the shared-shaping-constituent command in a scheduler profile.
Compound shared shaping	Hardware-assisted mechanism that controls bandwidth for all active constituents.
Simple shared shaping	Software-assisted mechanism that measures the rate of active constituents, and shapes the rate of the best-effort node or queue to the residual shared-shaping rate.

- Related Topics**
- For definitions of other common QoS terms, see QoS Terms on page 5

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 7 on page 73 compares the two types of shared shaping that are available.

Table 7: 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 Topics**
- Simple Shared Shaping Overview on page 79
 - Compound Shared Shaping Overview on page 99
 - Constituent Selection for Shared Shaping Overview on page 107

Guidelines for Configuring Simple and Compound Shared Shaping

When you configure shared shaping, be sure to consider the following behaviors.

Shared Shaping and Individual Shaping

You can use both the **shared-shaping-rate** command and the **shaping-rate** command in a single scheduler profile. For example, you can shape the best-effort node or queue to accept less than the remainder of the shared-shaping rate as in the following commands:

```
(config)#scheduler-profile shared-1mbps
(config-scheduler-profile)#shared-shaping-rate 1000000 simple
(config-scheduler-profile)#shaping-rate 500000
```

If you configure a shaping rate higher than the shared-shaping rate, the rate never exceeds the shared rate, so the router issues the following error message:

```
% shaping-rate cannot be greater than the shared-shaping-rate
```

Although you can configure a shared-shaping rate and a shaping rate in the same scheduler profile, the shaping rate must not exceed the shared-shaping rate. A scheduler profile that includes a shaping rate must not contain a shared-shaping rate that specifies a constituent as weighted.

Shared Shaping and Best-Effort Queues and Nodes

A scheduler profile that includes a shared-shaping rate cannot be associated with a queue other than the best-effort queue or a node other than the best-effort node.

A scheduler profile that is referenced by nodes or queues that are not best effort cannot be modified to include a **shared-shaping rate** command. A scheduler profile

that includes a **shared-shaping rate** command cannot be associated with a group node.

ATM and Shared Shaping

When you configure shared shaping with ATM, be sure to consider the following behaviors.

Sharing Bandwidth with the SAR

On ATM line modules, providers can use the SAR to implement bandwidth sharing for VCs. When the SAR is operating in default mode (that is, when the **no qos-mode-port** command is in effect), the SAR backpressures the VC node in the default traffic-class group, but traffic that is queued through a named traffic-class group is unaffected by VC backpressure. In the absence of voice and video traffic, the VC runs data traffic at the shared rate. When voice and video traffic start streaming, the SAR backpressures just the VC node in the default traffic-class group, thus sharing the bandwidth.

However, providers need to configure shared shaping on more than just ATM VCs. The SAR cannot support shared shaping per virtual path on ATM, and there is no SAR on Ethernet line modules. The shared shaper implemented in the HRR scheduler can support shared shaping for all these different configurations.

Shared Shaping and Low-CDV Mode

JUNOS releases before Release 6.0.0 implemented a *carve-out* scheduling model. If you configured multiple scheduler nodes for a VC or VP, the router added together the shaping rates for each scheduler node and shaped the corresponding VC or VP tunnel in the SAR to the sum of the rates. This implementation forced a strict-priority carve-out model for a logical interface, because the best-effort traffic cannot share unused bandwidth from the strict-priority traffic-class group.

Beginning with JUNOS Release 6.0.0, the router synchronizes the SAR rate for a VC or VP to the shared-shaping rate for the best-effort scheduler node for the VC or VP, so that the default behavior for low-CDV mode becomes shared shaping. Applying shared shaping to the best-effort queue does not synchronize the rate for the corresponding VC or VP in the SAR.

JUNOS releases before Release 6.1.0 had a different behavior than the current shared shaping model when multiple traffic-class groups were configured in low-CDV mode. In those releases, the shaping rates of the VC nodes in each group were added together, and the corresponding VC queue in the SAR was shaped to the sum. The same algorithm was used for shaping VP tunnels in the SAR—the shaping rates of all VP nodes in the hierarchical scheduler were added together to shape the VP tunnel in the SAR. This behavior implements a carve-out model for scheduling into VPs and VCs and generally is not as desirable as the shared shaping model supported in JUNOS Release 6.1.0 and later releases.

Beginning with JUNOS Release 6.1.0, low-CDV mode causes SAR shaping of VCs and VPs only when you specify the **shared-shaping-rate** command for the best-effort VC or VP node in the HRR scheduler.

For more information about configuring low-CDV mode, see “ATM Integrated Scheduler Overview” on page 159.

Logical Interface Traffic Carried in Other Queues

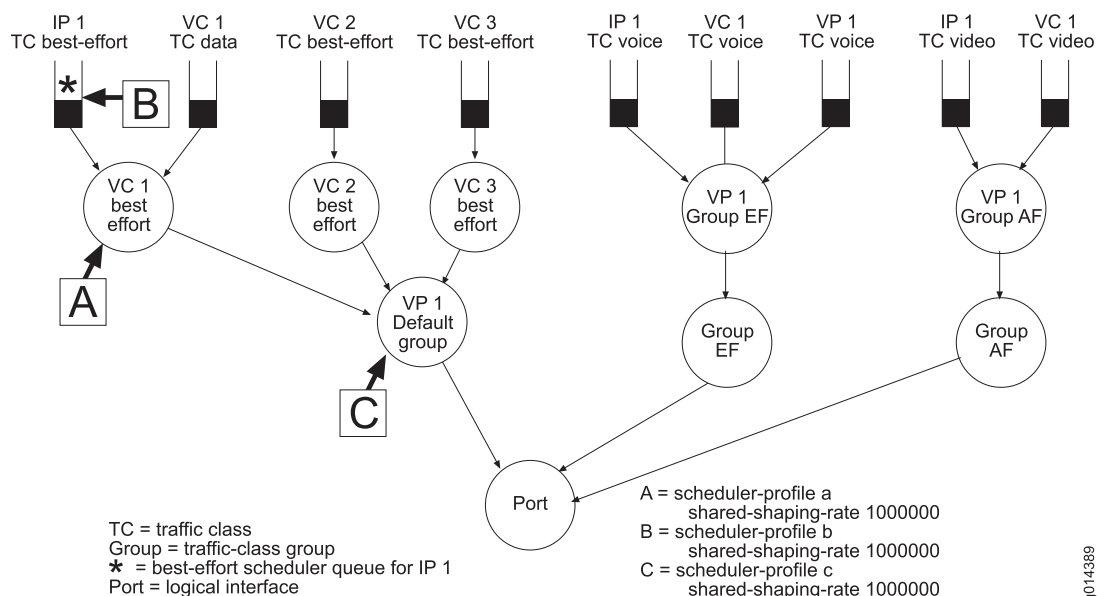
A shared shaper affects only the queues and nodes for a single interface. Queues associated with other interfaces are not constrained by the shared shaper. This behavior should cause no problems if you configure all queues for a single logical interface type. However, if you configure queues for multiple interface types, you may have problems with shared shaping.

For example, a shared shaper for VC 1 does not directly constrain the rate for a queue for IP 1 unless that queue is stacked above a node for VC 1 in the scheduler hierarchy. If the IP queue is stacked above a node for VC 1, then the shared shaper indirectly controls the queue bandwidth through the VC 1 node. But if the IP 1 queue is not stacked above a VC 1 node, it is immune to the shared shaper, and the total bandwidth for VC 1 can exceed the shared rate.

As another example, if a shared queue exists for VP 1 where VC 1 is contained within VP 1, the shared shaper for VC 1 does not constrain the bandwidth of a VP queue. The total bandwidth for VC 1 can again exceed the shared rate.

Figure 17 on page 76 illustrates an example of mixed interface shaping and its implications for implicit constituent selection for compound shared shaping.

Figure 17: Implicit Constituent Selection for Compound Shared Shaper: Mixed Interface Types



Traffic Starvation and Shared Shaping

Traffic in the strict-priority traffic-class group can starve out other traffic competing within the shared shaper. You might want to configure an individual shaping rate

for strict-priority queues, thus reserving the remaining shared bandwidth for nonstrict traffic.

For example, the following scheduler profiles limit the subscriber's strict priority traffic to 1.0 Mbps and limits the subscriber's aggregate traffic to 1.5 Mbps. If scheduler profile strictOne specified a shaping rate greater than or equal to 1.5 Mbps, nonstrict traffic might face starvation.

```
host1(config)#scheduler-profile strictOne
host1(config-scheduler-profile)#shaping-rate 1000000
host1(config-scheduler-profile)#exit
host1(config)#scheduler-profile nonStrictOne
host1(config-scheduler-profile)#shared-shaping-rate 1500000
```

Oversubscription and Shared Shaping

Many providers configure voice and video queues that combine to oversubscribe the shared rate. An external admission control agent, such as RADIUS, controls traffic flows such that the offered load does not ever really oversubscribe the shared rate. The static oversubscribed configuration on the router removes the need for the provider to signal voice or video traffic to the router.

Burst Size and Shared Shaping

The burst size for constituents is typically shaped by the burst value that you specify in the scheduler profile with the **shared-shaping-rate** command. You can override this burst for a particular constituent by applying another scheduler profile to that constituent and specifying the burst value with the **shaping-rate** command.

The following commands configure a VC shared shaper with two constituents, best effort and voice. The best-effort constituent has a burst of 30000 and the voice constituent has a burst of 16384.

```
host1(config)#scheduler-profile bestEffortBurst
host1(config-scheduler-profile)#shared-shaping-rate 1000000 burst 30000
host1(config-scheduler-profile)#exit
host1(config)#scheduler-profile voiceBurst
host1(config-scheduler-profile)#shaping-rate 300000 burst 16384
host1(config-scheduler-profile)#exit
```

Configure the QoS profile that applies the scheduler profiles:

```
host1(config)#qos-profile burstExample
host1(config-qos-profile)#atm-vc node
host1(config-qos-profile)#atm-vc node group EF
host1(config-qos-profile)#atm-vc queue traffic-class best-effort scheduler-profile
bestEffortBurst
host1(config-qos-profile)#atm-vc queue traffic-class voice scheduler-profile voiceBurst
```

- Related Topics**
- For a list of shared shaper terms, see Shared Shaper Terms on page 72
 - Configuring Simple Shared Shaping on page 81
 - Configuring Compound Shared Shaping on page 100

- Configuring Implicit Constituents for Simple or Compound Shared Shaping on page 114
- Configuring Explicit Constituents for Simple or Compound Shared Shaping on page 120

Chapter 10

Configuring Simple Shared Shaping of Traffic

This chapter provides information for configuring simple shared shaping of traffic on the E Series router.

QoS topics are discussed in the following sections:

- Simple Shared Shaping Overview on page 79
- Configuring Simple Shared Shaping on page 81
- Example: Simple Shared Shaping for ATM VCs on page 83
- Example: Simple Shared Shaping for ATM VPs on page 85
- Example: Simple Shared Shaping for Ethernet on page 86

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.

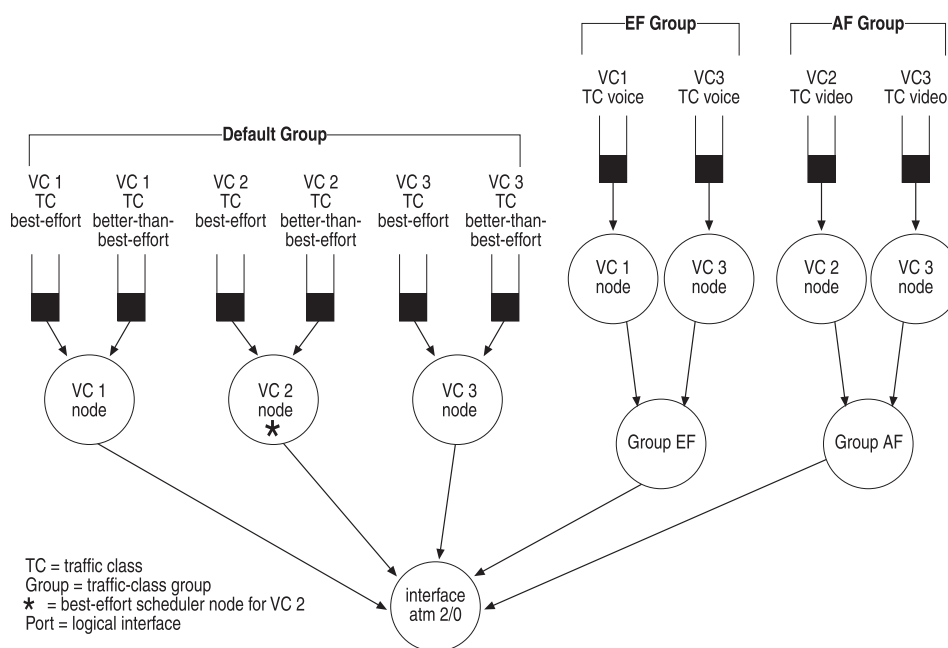
Simple Shared Shaping for Triple-Play Networks

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

In Figure 18 on page 80, 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.

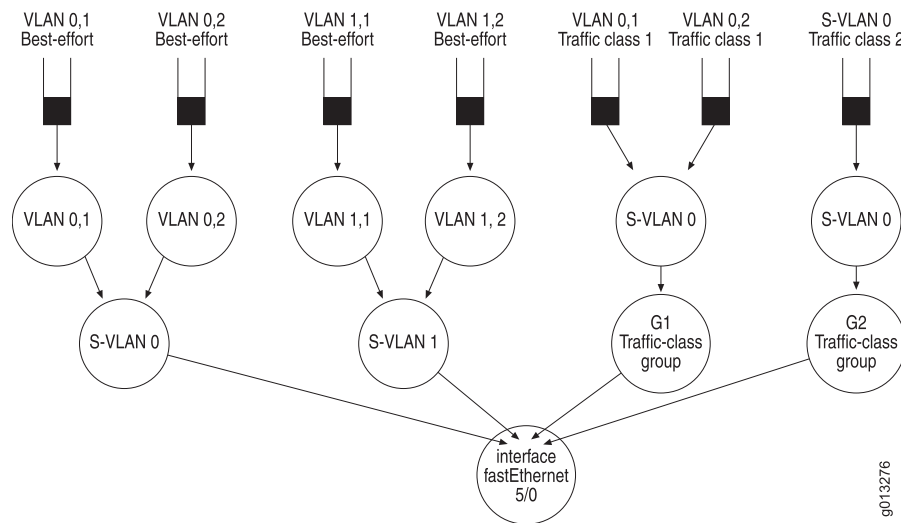
In this example, VC 1 is configured for voice and data. VC 2 is configured for data and video. VC 3 is configured for data, voice, and video. The shared shaper is configured on the best-effort node or queue for VC 1; the corresponding voice queue for VC 1 shares the configured rate.

Figure 18: 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 19 on page 81 illustrates how to shape the subscriber aggregate of voice, video, and data to a single rate in Ethernet.

Figure 19: Simple Shared Shaping over Ethernet



- Related Topics**
- For a list of shared shaper terms, see Shared Shaping Overview on page 71
 - Configuring Simple Shared Shaping on page 81
 - Constituent Selection for Shared Shaping Overview on page 107

Configuring Simple Shared Shaping

This section explains how to configure the shared shaper by specifying a shared-shaping rate for either the best-effort queue or the best-effort scheduler node for the logical interface. The router locates the other queues associated with the logical interface and shapes that set of queues to the shared rate.

You do not explicitly specify shared shaping on the other queues for the logical interface. You can configure individual shaping rates on the other queues that are less than the shared rate. These individual shapers have the effect of reserving some of the shared bandwidth for the other queues.

Before you configure simple shared shaping:

- Configure the traffic classes and traffic-class groups.
See “Configuring Traffic Classes That Define Service Levels” on page 14 and “Configuring Traffic-Class Groups That Define Service Levels” on page 15.

To configure simple shared shaping:

1. Create the scheduler profile.

```
host1(config)#scheduler-profile shared-1mbps
```

2. Configure the shared-shaping rate.

```
host1(config-scheduler-profile)#shared-shaping-rate 128000 burst 32767 simple
host1(config-scheduler-profile)#shared-shaping-rate 80000 + 53000
host1(config-scheduler-profile)#exit
```

The range for the shared-shaping rate is 1000–1000000000 bps (1 Kbps–1000 Kbps); the default is the minimum shaping rate (1 Kbps).

Use the *operator* and *operandValue* variables to specify the shared shaping rate as an expression.

Use the **bps** or **kbps** keywords to specify the unit of the shaping rate. By default, the shaping rate is configured in bps.

Use the **burst** keyword to configure the catch-up number associated with the shaper; the range is 0–522240 (0–510 KB). If you do not specify a burst value, the router selects an applicable default value.

Use the **milliseconds** or **bytes** keywords to specify the unit of the burst size.

You can specify **simple** to shape data queue rates to the value of the shared rate minus the combined voice and video traffic rate. By default, shared shaping is set to **auto**. In this mode, the router selects the type of shared shaping that is applied according to the type of line module. Compound shared shaping is hardware-dependent. If you specify **compound** for line modules that do not support it, an error message is generated and the router applies simple shared shaping.

3. Configure the QoS profile and reference the scheduler profile.

```
(config)#qos-profile subscriber-default-mode
(config-qos-profile)#atm-vc node
(config-qos-profile)#atm-vc node group AF
(config-qos-profile)#atm-vc node group EF
(config-qos-profile)#atm-vc queue traffic-class best-effort scheduler-profile
shared-1mbps
(config-qos-profile)#exit
```



TIP: The scheduler profile that you configured with the shared-shaping rate must be referenced in the best-effort queue or the best-effort scheduler node.

4. Attach the profile to the interface.

```
(config)#interface atm 11/0.10
(config-subif)#qos-profile subscriber-default-mode
(config-scheduler-profile)#exit
```

- Related Topics**
- Simple Shared Shaping Overview on page 79
 - Guidelines for Configuring Simple and Compound Shared Shaping on page 74
 - For more information about specifying an expression that you can reference within a scheduler profile, see Using Expressions for Bandwidth and Burst Values in a Scheduler Profile on page 48
 - Example: Simple Shared Shaping for ATM VCs on page 83
 - Example: Simple Shared Shaping for ATM VPs on page 85
 - Example: Simple Shared Shaping for Ethernet on page 86
 - node
 - qos-profile
 - queue
 - scheduler-profile
 - shared-shaping-rate
 - traffic-class
 - traffic-class-group

Example: Simple Shared Shaping for ATM VCs

The following commands configure a simple shared shaper for a VC, as shown in Figure 18 on page 80. In this example, the best-effort queue for logical interface VC 3 is shaped to a shared rate of 1 Mbps. The voice and video queues for VC 3 share the 1 Mbps with the best-effort traffic. The voice queue has first claim on the shared 1 Mbps, but only up to its individual shaping rate of 200 Kbps. The video queue claims up to the next 300 Kbps. The best-effort queue obtains whatever bandwidth remains of the 1 Mbps after the voice and video traffic have made their claims.

1. Configure the traffic classes and traffic-class groups.

```
(config)#traffic-class voice
(config-traffic-class)#fabric-strict-priority
(config-traffic-class)#exit
(config)#traffic-class video
(config-traffic-class)#exit
(config)#traffic-class-group EF auto-strict-priority
(config-traffic-class-group)#traffic-class voice
(config-traffic-class-group)#exit
((config)#traffic-class-group AF extended
(config-traffic-class-group)#traffic-class video
(config-traffic-class-group)#exit
```

2. Configure the shared shaper.

```
(config)#scheduler-profile 200kbps
(config-scheduler-profile)#shaping-rate 200000
(config-scheduler-profile)#exit
(config)#scheduler-profile 300kbps
```

```

(config-scheduler-profile)#shaping-rate 300000
(config-scheduler-profile)#exit
(config)#scheduler-profile shared-1mbps
(config-scheduler-profile)#shared-shaping-rate 1000000 simple
(config-scheduler-profile)#exit
(config)#qos-profile subscriber-default-mode
(config-qos-profile)#atm-vc node
(config-qos-profile)#atm-vc node group AF
(config-qos-profile)#atm-vc node group EF
(config-qos-profile)#atm-vc queue traffic-class best-effort scheduler-profile
shared-1mbps
(config-qos-profile)#atm-vc queue traffic-class video scheduler-profile 300kbps
(config-qos-profile)#atm-vc queue traffic-class voice scheduler-profile 200kbps
(config-qos-profile)#exit

```

3. Delete the rule in the default port type profile that creates IP best-effort queues by default.

```

(config)#qos-profile atm-default
(config-qos-profile)#no ip queue traffic-class best-effort
(config-qos-profile)#exit

```

4. Attach the profile to the ATM subinterface for VC 3.

```

(config)#interface atm 11/0.10
(config-subif)#qos-profile subscriber-default-mode
(config-scheduler-profile)#exit

```

The **qos-profile subscriber-default-mode** command shown in this example is appropriate if you have configured the SAR to be in default mode (by issuing the **no qos-mode-port** command). If this QoS profile is attached in low-CDV mode, the shaper is effective but the CDV is not correctly bounded, because the VC is not reshaped in the SAR.

The following commands configure a QoS profile different from the one shown in the previous example. In this example, the best-effort scheduler node for VC 3 is shaped to a shared rate of 1 Mbps. The **qos-profile subscriber-low-cdv-mode** command is appropriate if you configure the SAR in low-CDV mode (by issuing the **qos-mode-port low-cdv** command). The VC is reshaped to 1 Mbps in the SAR. If this QoS profile is attached in the SAR default mode, the 1-Mbps shaper is disabled by VC backpressure from the SAR.

```

(config)#qos-profile subscriber-low-cdv-mode
(config-qos-profile)#atm-vc node scheduler-profile shared-1mbps
(config-qos-profile)#atm-vc node group AF
(config-qos-profile)#atm-vc node group EF
(config-qos-profile)#atm-vc queue traffic-class best-effort
(config-qos-profile)#atm-vc queue traffic-class video scheduler-profile 300kbps
(config-qos-profile)#atm-vc queue traffic-class voice scheduler-profile 200kbps
(config-qos-profile)#exit

```


- Related Topics**
- Configuring Simple Shared Shaping on page 81
 - Simple Shared Shaping Overview on page 79

Example: Simple Shared Shaping for ATM VPs

In the example shown in Figure 20 on page 85, VP 1 is shaped to a shared rate of 5 Mbps. The shared shaper requires that voice and video traffic be carried in queues associated with the logical interface, which in this scenario is the VP. VP-level queuing does not guarantee fairness to the voice and video traffic for each VC, but fairness is not a major issue because admission control guarantees that the voice and video queues do not become congested.

This example assumes the same traffic class and traffic-class group configurations that are used in “Example: Simple Shared Shaping for ATM VCs” on page 83.

Figure 20: VP Shared Shaping



The following set of commands configures the shared shaper in Figure 20 on page 85.

```
(config)#scheduler-profile 2mbps
(config-scheduler-profile)#shaping-rate 2000000
(config-scheduler-profile)#exit
(config)#scheduler-profile 400kbps
(config-scheduler-profile)#shaping-rate 400000
(config-scheduler-profile)#exit
(config)#scheduler-profile shared-5mbps
(config-scheduler-profile)#shared-shaping-rate 5000000 simple
(config-scheduler-profile)#exit
(config)#qos-profile vp-subscriber1
(config-qos-profile)#atm-vp node scheduler-profile shared-5mbps
(config-qos-profile)#atm-vp node group AF
(config-qos-profile)#atm-vp node group EF
(config-qos-profile)#atm-vc node
(config-qos-profile)#atm-vc queue traffic-class best-effort scheduler-profile default
```

```
(config-qos-profile)#atm-vp queue traffic-class video scheduler-profile 2mbps
(config-qos-profile)#atm-vp queue traffic-class voice scheduler-profile 400kbps
(config-qos-profile)#exit
```

In this example, the best-effort scheduler node for the VP is shaped to a shared rate of 5 Mbps. The EF and AF queues for the VP share the 5 Mbps with the best-effort traffic. The EF queue has first claim on the shared 5 Mbps, but only up to its individual shaping rate of 400 Kbps. The AF queue claims up to the next 2 Mbps. The VC-level best-effort queues obtain whatever bandwidth remains of the 5 Mbps after the AF traffic and EF traffic have made their claims. This QoS profile is appropriate for low-CDV mode. If the provider configures a shapeless VP tunnel in the SAR, QoS sets the SAR shaper for the VP to match the 5-Mbps shared-shaping rate, and the CDV is bounded for the VP tunnel.

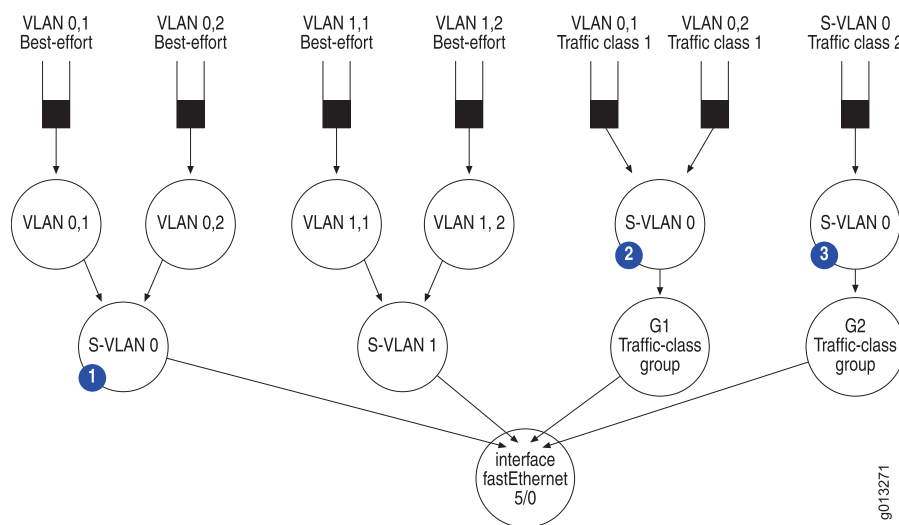
- Related Topics**
- Configuring Simple Shared Shaping on page 81
 - Simple Shared Shaping Overview on page 79

Example: Simple Shared Shaping for Ethernet

In a typical triple-play network configuration over Ethernet, individual subscribers are represented on the B-RAS by VLANs and DSLAMs by SVLANs. In this example, the provider shapes the subscriber aggregate of voice, video, and data to a single rate.

In this example, S-VLAN 0 has traffic in three traffic-class groups: the default group, the TC1 traffic class in the G1 group, and the TC2 traffic class in the G2 traffic-class group.

Figure 21: Hierarchical Simple Shared Shaping over Ethernet



In Figure 21 on page 86, the S-VLANs labeled 1, 2, and 3 indicate the possible constituents for S-VLAN 0. The active constituents for the simple shared shaper are the three nodes for S-VLAN 0 in the three traffic-class groups.



NOTE: This example uses QoS parameters to configure shared shaping.

1. Configure the traffic classes and traffic-class groups.

```
(config)#traffic-class tc1
(config)#exit
(config)#traffic-class tc2
(config)#exit
(config)#traffic-class-group g1
(config-traffic-class-group)#traffic-class tc1
(config-traffic-class-group)#exit
(config)#traffic-class-group g2 extended
(config-traffic-class-group)#traffic-class tc2
(config-traffic-class-group)#exit
```

2. Configure the parameter definitions.

```
(config)#qos-parameter-define vlan-g1-max-rate
(qos-parameter-define)#controlled-interface-type vlan
(qos-parameter-define)#instance-interface-type vlan
(qos-parameter-define)#exit
(config)#qos-parameter-define svlan-g1-max-rate
(qos-parameter-define)#controlled-interface-type svlan
(qos-parameter-define)#instance-interface-type svlan
(qos-parameter-define)#instance-interface-type ethernet
(qos-parameter-define)#exit
(config)#qos-parameter-define vlan-max-rate
(qos-parameter-define)#controlled-interface-type vlan
(qos-parameter-define)#instance-interface-type vlan
(qos-parameter-define)#instance-interface-type svlan
(qos-parameter-define)#exit
(config)#qos-parameter-define svlan-max-rate
(qos-parameter-define)#controlled-interface-type svlan
(qos-parameter-define)#instance-interface-type svlan
(qos-parameter-define)#instance-interface-type ethernet
(qos-parameter-define)#exit
```

3. Configure the shared shaper by referencing parameter definitions in the **shaping-rate** command.

```
(config)#scheduler-profile vlan-be
(config-scheduler-profile)#shared-shaping-rate vlan-max-rate simple
(config-scheduler-profile)#exit
(config)#scheduler-profile svlan-be
(config-scheduler-profile)# shared-shaping-rate svlan-max-rate simple
(config-scheduler-profile)#exit
(config)#scheduler-profile svlan-g1
(config-scheduler-profile)#shaping-rate svlan-g1-max-rate
(config-scheduler-profile)#exit
(config)#scheduler-profile vlan-g1
```

```
(config-scheduler-profile)#shaping-rate vlan-g1-max-rate
(config-scheduler-profile)#exit
(config)#scheduler-profile svlan-g2
(config-scheduler-profile)#shaping-rate svlan-max-rate % 50
(config-scheduler-profile)#exit
```

4. Configure the QoS profile.

```
(config)#qos-profile svlan-4.1
(config-qos-profile)#vlan queue traffic-class best-effort
(config-qos-profile)#vlan node scheduler-profile vlan-be
(config-qos-profile)#svlan node scheduler-profile svlan-be
(config-qos-profile)#vlan queue traffic-class tc1
(config-qos-profile)#svlan node scheduler-profile svlan-g1 group g1
(config-qos-profile)#svlan queue traffic-class tc2
(config-qos-profile)#svlan node scheduler-profile svlan-g2 group g2
(config-qos-profile)#ethernet group g2 scheduler-profile default
```

5. Attach the QoS profile to the S-VLANs on Fast Ethernet interface 11/0.

```
(config)#interface fastEthernet 11/0
(config-if)#svlan 0 qos-parameter svlan-max-rate 4000000
(config-if)#svlan 0 qos-profile svlan-4.1
(config-if)#encapsulation vlan
(config-if)#exit
(config)#interface fastEthernet 11/0.1
(config-if)#svlan id 0 1
(config-if)#ip address 1.2.1.1 255.255.255.0
(config-if)#exit
(config)#interface fastEthernet 11/0.2
(config-if)#svlan id 0 2
(config-if)#ip address 1.3.1.1 255.255.255.0
(config-if)#exit
```

- Related Topics**
- Configuring Simple Shared Shaping on page 81
 - Simple Shared Shaping Overview on page 79

Chapter 11

Configuring Variables in the Simple Shared Shaping Algorithm

This chapter provides information for configuring variables within the simple shared shaper algorithm on the E Series router.

QoS topics are discussed in the following sections:

- Simple Shared Shaping Algorithm Overview on page 89
- Variables of the Simple Shared Shaper Algorithm on page 91
- Guidelines for Controlling the Simple Shared Shaper Algorithm on page 93
- Configuring Simple Shared Shaper Algorithm Variables on page 93
- Sample Process for Controlling the Simple Shared Shaper Algorithm on page 95


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 22 on page 89 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 22: Simple Shared Shaper Behavior Without Algorithm Controls



When the video stream starts in the example displayed by Figure 22 on page 89, 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 23 on page 90. In this example, as the video traffic starts, the best-effort rate is reduced less drastically, and best-effort traffic is not starved.

Figure 23: Less Conservative Simple Shared Shaper Behavior



You can also configure the more liberal simple shared shaper behavior that is displayed in Figure 24 on page 90. 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 24: 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 Topics**
- Variables of the Simple Shared Shaper Algorithm on page 91
 - Configuring Simple Shared Shaper Algorithm Variables on page 93

Variables of the Simple Shared Shaper Algorithm

The formulas the simple shared shaper uses contain values maintained by the simple shared shaper algorithm, and variables that you configure.

The following factors are maintained by the simple shared shaper algorithm:

- **newMeasuredRate**—Sum of bytes enqueued to non-best-effort constituent queues, in bps.
- **oldDynamicRate**—Dynamic shaping rate from the previous rate period, in bits-per-second.
- **sharedShapingRate**—Configured shared shaper rate, in bps. The shared shaping rate is the total rate of all constituents of the shared shaper.

You can configure the following variables, which correspond to the commands described in “Configuring Simple Shared Shaper Algorithm Variables” on page 93.

- **convergenceFactor**—Controls the convergence of the dynamic shaping rate to the calculated shaping rate, expressed as a percentage of the available bandwidth.

The default value of 50 percent causes the dynamic shaping rate to converge by half of the available rate each period. For example, when the dynamic rate of a 5 Mbps simple shared shaper is 1 Mbps, and the measured rate goes from 4 Mbps to 0 Mbps, 4 Mbps of bandwidth becomes available. The simple shared shaper converges from 1 Mbps to 5 Mbps by half of the available bandwidth per period. In this example, the dynamic shaping rates for several periods are 1 Mbps, 3 Mbps, 4 Mbps, 4.5 Mbps, 4.75 Mbps, and so on.

- **maximumVOQL**—Sets the maximum virtual output queue length (VOQL), expressed in milliseconds (ms) of the shared shaping rate.

The default value of 4000 indicates that a 5 Mbps shared shaper does not allow the VOQL to exceed 20 Mbps. Smaller values reduce the effect of the VOQL in the simple shared shaper algorithm.

A maximum VOQL of 0 indicates that the shared shaper ignores the VOQL. This setting is appropriate for configurations where exceeding the shared shaping rate for brief periods of time does not cause downstream queuing.

- **minimumDynamicRate**—Sets the minimum value for the dynamic shaping rate, expressed as a percentage of the shared shaping rate. For example, a value of 25 for a 20 Mbps shared shaper specifies that the dynamic shaping rates never be set to a value less than 5 Mbps. The default value is 0.
- **reactionFactor**—Controls how the simple shared shaper reacts to changing rates, expressed as a percentage. The default value of 200 changes the algorithm to use 200 percent of the changing rate.

This section describes the algorithm tasks in detail.

Step 1: Calculate the New Measured Rate The simple shared shaper uses the following formula to calculate the new measured rate:

$$\text{newMeasuredRate} = \text{bytes enqueued} \times 8 \text{ bits per byte} / \text{rate-period} \times 1000 \text{ ms/sec}$$

Step 2: Calculate the VOQL The simple shared shaper maintains a VOQL, which cannot become less than zero, using the following formulas:

$$\text{VOQL} = \text{VOQL} + (\text{oldDynamicRate} - \text{oldMeasuredRate} - \text{sharedShapingRate})$$

$$\text{If } (\text{VOQL} > \text{maximumVOQL}), \text{ then } (\text{VOQL} = \text{maximumVOQL})$$

Step 3: Calculate the New Dynamic Rate Each rate period, the simple shared shaper calculates the new dynamic rate (the shaping rate of the best-effort node or queue) using the following formula. The system prevents the new dynamic rate from becoming less than zero.

$$\begin{aligned} \text{newDynamicRate} = & (\text{convergenceFactor} \times \text{oldDynamicRate}) + (1 - \text{convergenceFactor}) \\ & \times (\text{sharedShapingRate} - \text{newMeasuredRate}) - \text{reactionFactor} \times (\text{newMeasuredRate} \\ & - \text{oldMeasuredRate}) - \text{VOQL} \end{aligned}$$

Step 4: Determine the Larger Value of the New Dynamic Rate and the Minimum Dynamic Rate The simple shared shaper determines the larger of the new dynamic rate and a minimum dynamic rate, where the minimumDynamicRate is a fraction of the shared-shaping rate, using the following formula:

$$\text{Max}(\text{newDynamicRate}), (\text{minimumDynamicRatePercent} \times \text{sharedShapingRate})$$

- Related Topics**
- Simple Shared Shaping Algorithm Overview on page 89
 - Sample Process for Controlling the Simple Shared Shaper Algorithm on page 95

Guidelines for Controlling the Simple Shared Shaper Algorithm

You can configure the simple shared shaper variables individually, but it is useful to use configuration guidelines to determine how the variables work together to achieve a desired behavior.

Table 8 on page 93 displays guidelines for configuring the most liberal shared shaper to the most conservative shared shaper.

- Most liberal—Appropriate when over-queuing is not a concern
- Liberal—Appropriate when over-queuing is not a concern and a smoother rate adjustments are desirable
- Moderate—Default settings
- Conservative—Appropriate when over-queuing is a major concern
- Most conservative—Rarely appropriate.

Table 8: Guidelines for Configuring Simple Shared Shaper Algorithm Variables

Control	Most Liberal	Liberal	Moderate	Conservative	Most Conservative
convergence-factor	0	25	50	75	99
maximum-voql	0	25	400	600	1000
reaction-factor	0	50	200	300	1000

- Related Topics**
- Simple Shared Shaping Algorithm Overview on page 89
 - Configuring Simple Shared Shaper Algorithm Variables on page 93

Configuring Simple Shared Shaper Algorithm Variables

To configure the variables for all simple shared shapers on the router:

1. Enter QoS Shared Shaper Control Configuration mode.

```
host1(config)#qos-shared-shaper-control
host1(config-qos-shared-shaper-control)#
```

2. (Optional) Configure the convergence factor for all simple shared shapers on the router.

```
host1(config-qos-shared-shaper-control)#convergence-factor 25
```

The convergence factor determines how quickly the dynamic shaping rate converges with the calculated dynamic shaping rate, and is expressed as a percentage of the available bandwidth.

The range for the convergence factor is 0–99 percent, with 0 being the most liberal and 99 the most conservative. The default value is 50.

3. (Optional) Configure the specify the reaction factor for all simple shared shapers on the router.

```
host1(config-qos-shared-shaper-control)#reaction-factor 50
```

The reaction factor determines how the shared shaper reacts to changes in the measured rate.

The range for the reaction factor is 0–1000; 0 is the most liberal and 1000 is the most conservative. The default value is 200.

4. (Optional) Specify the minimum value of the dynamic shaping rate as a percentage of the shared shaping rate for all simple shared shapers on the router.

```
host1(config-qos-shared-shaper-control)#minimum-dynamic-rate-percent 50
```

The range for the minimum dynamic rate value is 0–100 percent. The default value is 0.

5. (Optional) Configure a maximum value for the virtual output queue length (VOQL) for all simple shared shapers on the router.

```
host1(config-qos-shared-shaper-control)#maximum-voql 25
```

The VOQL tracks the amount of data over queued between simple shared-shaper rate periods.

The range for the maximum VOQL value is 0–10000 milliseconds (ms). The default value is 4000.

- Related Topics**
- Simple Shared Shaping Algorithm Overview on page 89
 - Variables of the Simple Shared Shaper Algorithm on page 91
 - Guidelines for Controlling the Simple Shared Shaper Algorithm on page 93
 - Sample Process for Controlling the Simple Shared Shaper Algorithm on page 95
 - Configuring Simple Shared Shaping on page 81
 - convergence-factor
 - maximum-voql
 - minimum-dynamic-rate-percent
 - qos-shared-shaper-control
 - reaction-factor

Sample Process for Controlling the Simple Shared Shaper Algorithm

The simple shared shaper in this example contains two constituents, best-effort and video. The shared-shaping rate is 15 Mbps, and the video rate is 4 Mbps.

The example contains two parts: when the video flow is turned on, and then turned off.



NOTE: The rates in this example are approximate and for illustrative purposes only. Your configuration might yield different results based on network variables.

Starting Video Flow

Table 9 on page 95 lists the dynamic rate when the video flow is turned on for the five classes of simple shared shaper variables. Results vary because the amount of video measured in the first rising period is random, in the range 0–4 Mbps non-inclusive.

Table 9: Rising Edge Sample When Video Flow Starts

Control	Period of Dynamic Rate, in Kbps									
	1	2	3	4	5	6	7	8	9	10
Most liberal	15000	13080	11000	11000	11000	11000	11000	11000	11000	11000
Liberal	15000	9542	8880	10470	10867	10972	10979	10994	10998	10994
Moderate	15000	6510	5606	8303	9651	10329	10628	10814	10967	10953
Conservative	15000	6022	1604	3953	5714	7038	7978	8733	9300	9735
Most conservative	–	–	–	–	–	–	–	–	–	–

In this example, a liberal maximum VOQL value is ineffective because the 15 Mbps shared-shaping rate is much higher than the 4 Mbps video rate. The video rate divided by the shared shaping rate is 26.6 percent, so any value higher than this has no effect.




NOTE: The rates in this example represent approximate egress-queue enqueue rates on an Ethernet line module; therefore, there is no ATM SAR or downstream devices are not used. More liberal configurations can be inappropriate when there might be queuing between the scheduler and the destination. VLAN queuing is used, and saturation rates are offered.

The most liberal case heavily reduces VOQL and changes of rate, leading to a shared shaper that quickly converges. The conservative configuration overreacts to VOQL and the change of rate, and converges very slowly.

Figure 25 on page 96 shows a graph of the dynamic rate when the video flow starts.

Figure 25: Dynamic Rate When Video Flow Starts



Stopping Video Flow


Table 10 on page 96 lists the dynamic rate as the video flow stops for the five classes of simple shared shaper controls. Results might vary because the amount of video measured in the first falling period is random, in the range 0–4 Mbps non-inclusive.

Table 10: Data When Video Flow Stops

Period of Dynamic Rate, in Kbps										
Control	1	2	3	4	5	6	7	8	9	10
Most liberal	11000	12132	15000	15000	15000	15000	15000	15000	15000	15000
Liberal	11000	11584	14146	14786	14946	14986	14996	14999	14999	15000
Moderate	11000	11728	13364	14182	14591	14795	14897	14948	14974	14987
Conservative	10955	11278	12208	12906	13429	13822	14116	14337	14503	14701
Most conservative	–	–	–	–	–	–	–	–	–	–

Figure 26 on page 97 shows a graph of the dynamic rate when the video flow stops.

Figure 26: Dynamic Rate When Video Flow Stops



- Related Topics**
- Simple Shared Shaping Algorithm Overview on page 89
 - Variables of the Simple Shared Shaper Algorithm on page 91
 - Configuring Simple Shared Shaper Algorithm Variables on page 93

Chapter 12

Configuring Compound Shared Shaping of Traffic

This chapter provides information for configuring compound shared shaping of traffic on the E Series router.

QoS topics are discussed in the following sections:

- Compound Shared Shaping Overview on page 99
- Configuring Compound Shared Shaping on page 100
- Example: Compound Shared Shaping for ATM VCs on page 102
- Example: Compound Shared Shaping for ATM VPs on page 104

Compound Shared Shaping Overview

Compound shared shaping is a hardware-assisted mode that can control bandwidth for all scheduler objects associated with the subscriber logical interface. Thus it can manage 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.

Supported Hardware for Compound Shared Shaping

You can configure compound shared shaping on a line module with the EFA2 or TFA hardware.

The EFA2 implementation is different from the EFA ASIC, which does not implement compound shared shaping. Issue the **show qos shared-shaper** command to determine whether compound shared shapers are supported for the line module. Contact your Juniper Networks account representative for more information about line modules with the EFA2 ASIC.

The TFA hardware is only available on the ES2 10G LM on the E120 and E320 Broadband Services Routers.

If you configure a compound shared shaper on hardware that does not support it, the CLI displays the following message:

```
host1(config)#ERROR 02/08/2005 14:06:36 qos: line card in slot 11: EFA2 hardware
not installed. 1 compound shared shaper(s) converted to simple.
```

QoS automatically converts the compound shared shaper to a simple shared shaper.



NOTE: Compound shared shaping is not supported by the frame forwarding ASIC (FFA).

Bandwidth Allocation for Compound Shared Shaping

The compound shared-shaper mechanism actively allocates the bandwidth it receives from the hierarchical scheduler to each active constituent, based on its own rules, independent of the hierarchical scheduler. Constituents are either *priority* constituents or *weighted* constituents. These attributes are specified using the **shared-shaper-constituent** command.

Compound shared-shaper scheduling allocates bandwidth as follows:

1. Priority constituents consume as much of the shared bandwidth as they can, subject to the bandwidth allocated to them by the hierarchical scheduler.
2. Priority constituents are ordered according to their priority.
3. The weighted constituents subdivide the remaining shared bandwidth in proportion to their shared weights, again subject to the bandwidth allocated to them by the hierarchical scheduler.

- Related Topics**
- For a list of shared shaper terms, see Shared Shaping Overview on page 71
 - Configuring Compound Shared Shaping on page 100

Configuring Compound Shared Shaping

Compound shared shaping requires that you set a shared-shaping rate in a scheduler profile associated with a best-effort node or queue.

Before you configure compound shared shaping:

- Configure the traffic classes and traffic-class groups.
See “Configuring Traffic Classes That Define Service Levels” on page 14 and “Configuring Traffic-Class Groups That Define Service Levels” on page 15.

To configure compound shared shaping:

1. Create the scheduler profile.

```
host1(config)#scheduler-profile compound
```

2. Configure the compound shared shaper.

```
host1(config-scheduler-profile)#shared-shaping-rate 128000 burst 32767
compound explicit-constituents
```

The range for the shared-shaping rate is 1000–1000000000 bps (1 Kbps–1000 Kbps); the default is no shaping rate.

Use the *operator* and *operandValue* variables to specify the shared shaping rate as an expression.

Burst is the catch-up number associated with the shaper; the range is 0–522240 (0–510 KB). Specifying 0 enables the router to select an applicable default value.

By default, shared shaping is set to **auto**, where the router selects the type of shared shaping that is configured, depending on the line module. You must specify the **compound** keyword to actively shape voice and video traffic so that the shared rate cannot be exceeded, and shape data queue rates to the value of the shared rate minus the combined voice and video traffic rate. An error message is generated if you specify **compound** for line modules that do not support it, and the router applies simple shared shaping.

3. Configure the QoS profile and reference the scheduler profile.

```
(config)#qos-profile compound
(config-qos-profile)#atm-vc node
(config-qos-profile)#atm-vc node group AF
(config-qos-profile)#atm-vc node group EF
(config-qos-profile)#atm-vc queue traffic-class best-effort scheduler-profile
shared-1mbps
(config-qos-profile)#exit
```



TIP: The scheduler profile that you configured with the shared-shaping rate must be referenced in the best-effort queue or the best-effort scheduler node.

4. Attach the profile to the interface.

```
(config)#interface atm 11/0.10
(config-subif)#qos-profile subscriber-default-mode
(config-scheduler-profile)#exit
```

- Related Topics**
- Compound Shared Shaping Overview on page 99
 - Guidelines for Configuring Simple and Compound Shared Shaping on page 74
 - For more information about specifying an expression that you can reference within a scheduler profile, see Using Expressions for Bandwidth and Burst Values in a Scheduler Profile on page 48
 - Example: Compound Shared Shaping for ATM VCs on page 102
 - Constituent Selection for Shared Shaping Overview on page 107
 - node
 - qos-profile
 - queue

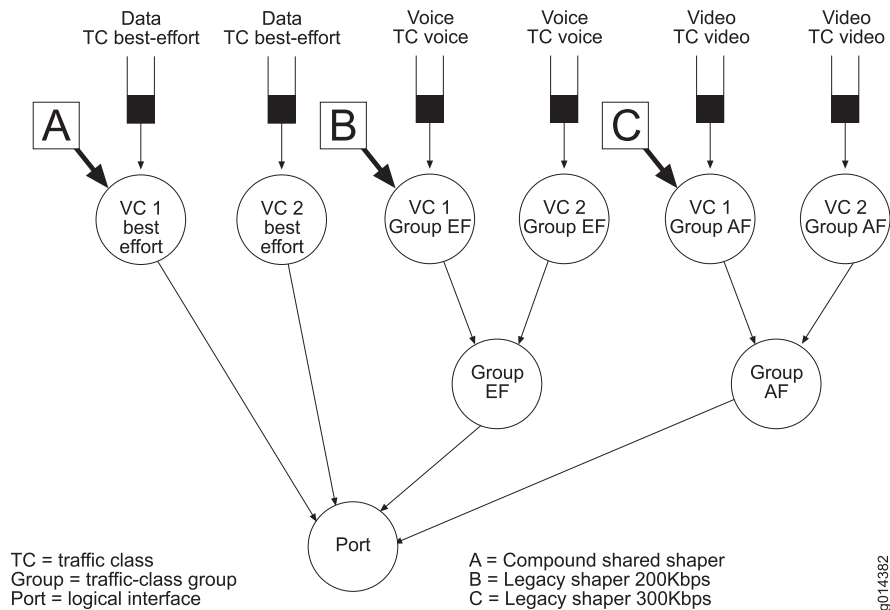
- scheduler-profile
- shared-shaping-rate
- traffic-class
- traffic-class-group

Example: Compound Shared Shaping for ATM VCs

Figure 27 on page 102 illustrates a typical DSL triple-play configuration, involving voice, video, and data traffic. In this example, a total of 1 Mbps of bandwidth is allocated to voice, video, and best-effort data traffic associated with the VC 1 logical interface.

The voice queue in the EF traffic-class group for VC 1 is a strict constituent that has first claim on up to 200 Kbps of the shared bandwidth. The video queue in the AF traffic-class group for VC 1 is a strict constituent that can claim up to 300 Kbps of the remaining 800–1000 Kbps of shared bandwidth. The best-effort queue for logical interface VC 1 is a strict constituent that has the last claim to the remaining 500–1000 Kbps of shared bandwidth.

Figure 27: VC Compound Shared Shaping Example



To configure VC compound shared shaping:

1. Configure the traffic classes, traffic-class groups, and additional scheduler profiles.
2. Configure the scheduler profile that defines the shared shaper and the profiles that apply the legacy shaper.

```
host1(config)#scheduler-profile shared-1Mbps
```

```

host1(config-scheduler-profile)#shared-shaping-rate 1000000 burst 32768 auto
host1(config)#scheduler-profile 300Kbps
host1(config-scheduler-profile)#shaping-rate 300000
host1(config)#scheduler-profile 200Kbps
host1(config-scheduler-profile)#shaping-rate 200000

```

3. Configure the QoS profile.

```

host1(config)#qos-profile vcSharedShaping

```

4. Create group nodes.

```

host1(config-qos-profile)#atm group AF scheduler-profile default
host1(config-qos-profile)#atm group EF scheduler-profile default

```

5. Create VC nodes for each group and for traffic in the default group.

```

host1(config-qos-profile)#atm-vc node
host1(config-qos-profile)#atm-vc node group AF
host1(config-qos-profile)#atm-vc node group EF

```

6. Create queues for the best-effort, video, and voice traffic. Apply the scheduler profile that defines the shared-shaping rate to the best-effort queue. Apply the legacy shaper profiles to the voice and video traffic queues.

```

host1(config-qos-profile)#atm-vc queue traffic-class best-effort scheduler-profile
shared-1mbps
host1(config-qos-profile)#atm-vc queue traffic-class video scheduler-profile
300Kbps
host1(config-qos-profile)#atm-vc queue traffic-class voice scheduler-profile
200Kbps
host1(config-qos-profile)#exit

```

7. Attach the QoS profile to an ATM subinterface.

```

host1(config)#interface atm 11/0.1
host1(config-interface)#qos-profile vcSharedShaping
host1(config-interface)#exit

```

In this example, the constituents of the VC shared shaper are the VC 1 best effort node, the VC 1 Group EF node, and the VC 1 Group AF node. The available bandwidth is strictly allocated in the following order:

1. VC 1 EF group node
2. VC 1 AF group node
3. VC 1 best effort node

To display the sample shared shaper configuration:

```

host1# show shared-shaper atm 11/0.1

```

interface	shared shaping rate	current shaping rate	resource	shaping rate
-----	-----	-----	-----	-----

```

atm-vc ATM11/0.1 1000000 compound best-effort atm-vc queue
                                atm-vc best-effort node
                                EF voice atm-vc queue      200000
                                AF video atm-vc queue      300000
atm-vc ATM11/0.2 1000000 compound best-effort atm-vc queue
                                atm-vc best-effort node
                                EF voice atm-vc queue      200000
                                AF video atm-vc queue      300000

Total shared shapers: 2
Total constituents:  8
Total failovers:     0

```

- Related Topics**
- Configuring Compound Shared Shaping on page 100
 - Compound Shared Shaping Overview on page 99

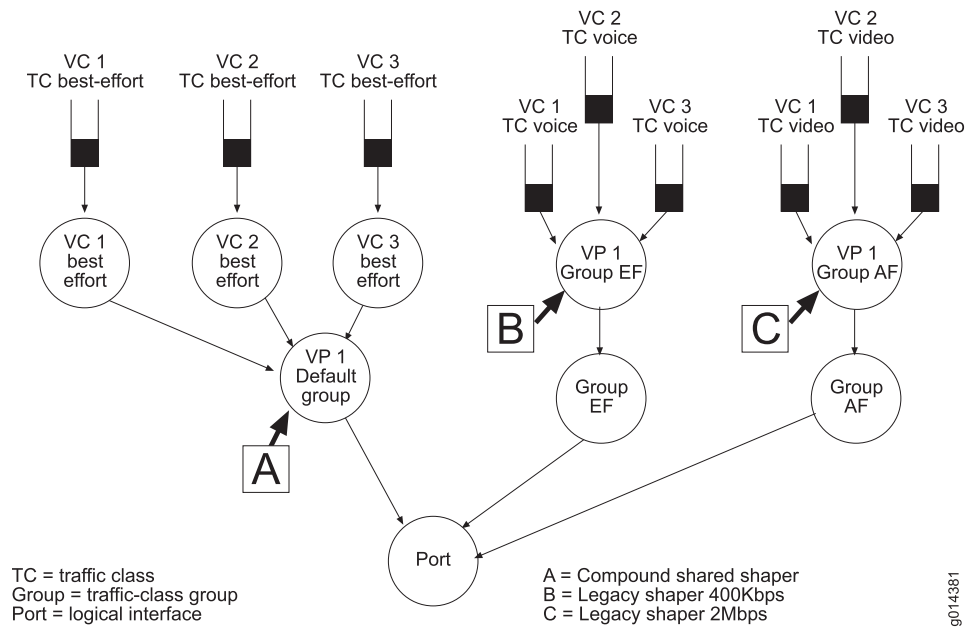
Example: Compound Shared Shaping for ATM VPs

Figure 28 on page 105 shows a compound shared shaper for a VP interface. VP shared shaping enables a shared shaper to apply to all the aggregate rates of all VCs within the VP.

In this example, the VP is shaped to a compound shared rate of 5 Mbps. The voice traffic gets strict priority scheduling for up to 400 Kbps of the shared rate on the VP. The video traffic gets up to 2 Mbps of the remaining 4.6–5 Mbps on the VP. Finally, the data traffic has the last claim to the remaining 2.6–3 Mbps of shared VP bandwidth.

This configuration enables data traffic to flow at 2.6 Mbps when voice and video are both using their limit. When both voice and video are quiescent, data can flow at the full 5 Mbps shared rate.

The QoS profile used in this example is appropriate for low-CDV mode. If the provider configures a shapeless VP tunnel in the SAR, QoS sets the SAR shaper for the VP to match the 5 Mbps shared-shaping rate, and the CDV is bounded for the VP tunnel. VP-level queuing does not guarantee fairness to the voice and video for each VC.

Figure 28: VP Compound Shared Shaping Example

To configure VP compound shared shaping:

1. Configure the traffic classes, traffic-class groups, and additional scheduler profiles.
2. Configure the scheduler profile that defines the shared shaper and the profiles that apply the legacy shaper.

```
host1(config)#scheduler-profile shared-5Mbps
host1(config-scheduler-profile)#shared-shaping-rate 5000000 burst 32768 auto
host1(config-scheduler-profile)#exit
```

3. Configure the scheduler profile for AF (video) traffic.

```
host1(config)#scheduler-profile 2Mbps
host1(config-scheduler-profile)#shaping-rate 2000000
```

4. Configure the scheduler profile for EF (voice) traffic.

```
host1(config)#scheduler-profile 400Kbps
host1(config-scheduler-profile)#shaping-rate 400000
host1(config-scheduler-profile)#exit
```

5. Configure the QoS profile.

```
host1(config)#qos-profile vpSharedShaping
```

6. Create group nodes.

```
host1(config-qos-profile)#atm group AF scheduler-profile default
host1(config-qos-profile)#atm group EF scheduler-profile default
```

7. Create VP nodes for each group and for traffic in the default group. The scheduler profile containing the shared-shaping rate is applied to the VP node that is in the default group and contains the best-effort queue.

```
host1(config-qos-profile)#atm-vp node scheduler-profile shared-5Mbps
host1(config-qos-profile)#atm-vp node group AF scheduler-profile 2Mbps
host1(config-qos-profile)#atm-vp node group EF scheduler-profile 400Kbps
```

8. Create a VC node for the default group.

```
host1(config-qos-profile)#atm-vc node
```

9. Create queues for the best-effort, video, and voice traffic.

```
host1(config-qos-profile)#atm-vc queue traffic-class best-effort
host1(config-qos-profile)#atm-vc queue traffic-class AF
host1(config-qos-profile)#atm-vc queue traffic-class EF
host1(config-qos-profile)#exit
```

10. Attach the QoS profile to an ATM subinterface.

```
host1(config)#interface atm 11/0.1
host1(config-interface)#qos-profile vpSharedShaping
```

In this example, the constituents of the VP shared shaper are the VP 1 default group node, the VP 1 Group EF node, and the VP 1 Group AF node. The available bandwidth is strictly allocated in the following order:

1. VP1 EF group node
2. VP1 AF group node
3. VP1 default group node

- Related Topics**
- Configuring Compound Shared Shaping on page 100
 - Compound Shared Shaping Overview on page 99

Chapter 13

Configuring Implicit and Explicit Constituent Selection for Shaping

This chapter provides information for configuring implicit and explicit constituents on the E Series router.

QoS topics are discussed in the following sections:

- Constituent Selection for Shared Shaping Overview on page 107
- Implicit Constituent Selection Overview on page 109
- Configuring Implicit Constituents for Simple or Compound Shared Shaping on page 114
- Explicit Constituent Selection Overview on page 115
- Configuring Explicit Constituents for Simple or Compound Shared Shaping on page 120

Constituent Selection for Shared Shaping Overview

Shared shaping supports both *implicit* and *explicit* constituent selection. Implicit constituent selection is the easier of the two methods and works well for most cases. With implicit selection, you configure a shared-shaping rate on the best-effort node or queue and QoS locates the other constituents automatically.

Use explicit constituent selection when you want to shape a subset of the interface traffic to the shared rate. An example of this is when you want the sum of best-effort and voice traffic to be shaped to the shared rate, but want video traffic to be exempt from the shared-shaping rate.

Active constituents are selected either implicitly by QoS or explicitly by the user. Active constituents of the simple shared shaper can be any node and queues in named traffic-class groups. Active constituents of the compound shared shaper can be nodes or queues. If you choose a node as an active constituent, queues above it are not active constituents.

Inactive constituents are queues that are stacked above an active node or nodes stacked below active queues. For both of these situations, the shared shaper controls the active constituents, and the legacy scheduler indirectly controls the inactive constituents to achieve the shared rate. The other case for inactive constituents is when you use explicit constituent selection and some of the nodes and queues are explicitly not included in the shared shaper.

To use implicit constituent selection, you specify only the shared-shaping rate and the logical interface. The router identifies the constituents associated with the logical interface type and their allocated bandwidth. This method is appropriate for the typical case where the intent is to shape all subscriber queues to the shared rate.

If you want instead to shape a subset of the queues for a subscriber to the shared rate, the explicit selection process is appropriate. Explicit selection is also useful when you want queues as the active constituents instead of the node below them. By choosing queues you can assign appropriate priority or weights.

Types of Shared Shaper Constituents

The **shared-shaping-constituent** command in a scheduler profile specifies constituents and their attributes. The command has two aspects. For explicit constituent selection, this command specifies the constituents. For the compound shared shaper only, this command specifies scheduling attributes of shared shaping: the shared priority and the shared weight.

A shared shaper can be one of the following four types:

- Simple implicit—Constituents are best-effort node or queues, and all nodes and queues in named traffic-class groups.
- Simple explicit—The software selects constituents based on the **shared-shaping-constituent** command. The weight and priority attributes of the **shared-shaping-constituent** command are ignored, because the simple shared shaper does not allocate bandwidth among constituents; instead it controls just the best-effort queue or node.
- Compound implicit—Constituents are selected automatically by the software. If a node exists in a given traffic-class group, the node is active and the queues stacked above it are inactive constituents. The **shared-shaping-constituent** command does not affect constituent selection. However, if the command is present for a constituent that was implicitly selected, the software configures that constituent with the shared priority and shared weight as indicated.
- Compound explicit—The software selects constituents based on the shared priority and shared weight configured with the **shared-shaping-constituent** command. If no attributes are specified, the software supplies a shared priority consistent with the legacy scheduler configuration.

Table 11 on page 108 compares implicit and explicit shared shaping.

Table 11: Comparison of Implicit and Explicit Shared Shaping

Implicit Shared Shaping	Explicit Shared Shaping
<ul style="list-style-type: none"> ■ To specify the logical interface for shared shaping, associate a scheduler profile that includes the shared-shaping-rate command or the shared-shaping-rate simple command with a best-effort node or queue. 	<ul style="list-style-type: none"> ■ To specify the logical interface for shared shaping, associate a scheduler profile that includes the shared-shaping-rate rate explicit-constituents command or the shared-shaping-rate rate simple explicit-constituents command with a best-effort node or queue.

Table 11: Comparison of Implicit and Explicit Shared Shaping *(continued)*

Implicit Shared Shaping	Explicit Shared Shaping
<ul style="list-style-type: none"> ■ Constituents consist of all nodes and queues for the same logical interface type. 	<ul style="list-style-type: none"> ■ Constituents consist of all nodes and queues for the same logical interface type.
<ul style="list-style-type: none"> ■ Active constituents are automatically selected from all constituents according to the implicit shared shaping rules. 	<ul style="list-style-type: none"> ■ Active constituents are explicitly selected from all constituents by association with a scheduler profile that includes the shared-shaper-constituent command. ■ If the scheduler profile associated with a constituent does not include this command, then the constituent is not active and is not shaped by the shared shaper.

- Related Topics**
- Implicit Constituent Selection Overview on page 109
 - Configuring Implicit Constituents for Simple or Compound Shared Shaping on page 114
 - Explicit Constituent Selection Overview on page 115
 - Configuring Explicit Constituents for Simple or Compound Shared Shaping on page 120

Implicit Constituent Selection Overview

The implicit selection process for simple and compound shared shaping are the same. The process operates according to the following rules:

1. The point at which the scheduler profile that contains a **shared-shaping-rate** command is associated with a best-effort node or best-effort queue determines the logical interface type that the shared shaper applies to. Logical interface types include IP, VP, VC, VLAN, S-VLAN, and so on.
2. All nodes and queues for the same logical interface are potential constituents.
3. The best-effort node is selected if you configure node-based shared shaping. The best-effort queue is selected if you configure queue-based shared shaping. If you configure both, then the best-effort node is selected over the best-effort queue.
4. Non-best-effort queues are selected.

The implicit selection process for compound shared shaping operates according to the following rules:

1. The point at which the scheduler profile that contains a **shared-shaping-rate** command is associated with a best-effort node or best-effort queue determines the logical interface type that the shared shaper applies to. Logical interface types include IP, VP, VC, VLAN, and S-VLAN.
2. All nodes and queues for the same logical interface are potential constituents.

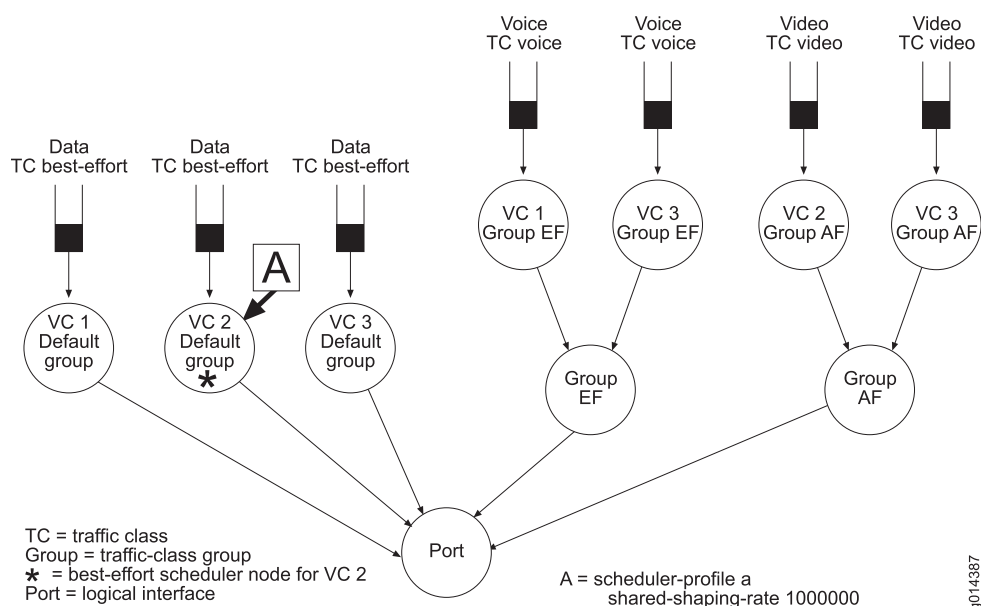
3. Nodes are selected over queues.

For example, suppose a shared shaper is associated with a particular interface type. A node for that interface type is present and has a queue for that interface type stacked above it. The node is selected and becomes an active constituent; the queue is not selected.

Now suppose a shared shaper is associated with a logical interface at the best-effort node, and a second shared shaper is simultaneously associated with the same interface at the best-effort queue. In this case, the node is selected as the constituent, because nodes are selected over queues.

In Figure 29 on page 110, scheduler profile A includes a shared-shaping rule, and is associated with the best-effort node for VC 2. The constituents are all the scheduler objects associated with VC 2: VC 2 nodes and VC 2 queues. Nodes are selected over queues, so the implicitly selected active constituents are the VC 2 default group node, the VC 2 Group EF node, and the VC 2 Group AF node.

Figure 29: Implicit Constituent Selection for Compound Shared Shaper at Best-Effort Node



In Figure 30 on page 111, scheduler profile B is associated with the best-effort queue for VC 3. This association indicates that the logical interface type being shared is VC. The constituents are all the scheduler objects associated with VC 3: VC 3 nodes and VC 3 queues. Nodes are selected over queues, so the implicitly selected active constituents for profile B's shared shaper are the VC 3 default group queue, the VC 3 Group EF node, and the VC 3 Group AF node. The VC 3 default group queue is selected instead of the VC 3 default group node because the shared shaper is associated with that best-effort queue.


Figure 30: Implicit Constituent Selection for Compound Shared Shaper at Best-Effort Queue

Figure 30 on page 111 illustrates other examples of implicit constituent selection. It does not reflect typical configurations, but includes a mixture of interface types: IP, VC, and VP. If only scheduler profile A is applied, the associated interface is VC 1. The selected constituents then consist of the VC 1 best-effort node, the VC 1 TC voice queue, and the VC 1 TC video queue.

If only scheduler profile B is applied, the associated interface is IP 1. The selected constituents then consist of the IP 1 best-effort queue, the IP 1 TC voice queue, and the IP 1 TC video queue.

If only scheduler profile C is applied, the associated interface is VP 1. The selected constituents then consist of the VP 1 default group node, the VP 1 Group EF node, and the VP 1 Group AF node.

Implicit Bandwidth Allocation for Compound Shared Shaping

After selecting the implicit constituents for compound shared shaping, the router places the constituents in an order that determines how the constituents can claim a share of the available shared bandwidth.

When it implements compound implicit shared shapers, the software selects attributes for the active constituents consistent with the hierarchical scheduler.

- Auto-strict nodes and queues have the highest priority.
- Nodes and queues in extended traffic-class groups are next.
- Nodes and queues in the default traffic-class group have the lowest priority.

For example, suppose a compound shared shaper has a rate of 2 Mbps. The shared shaper has three active constituents: the best-effort node, a voice queue in the auto-strict traffic-class group, and a video queue in an extended traffic-class group.

For compound implicit shared shaping, the shared shaper assigns the voice queue all the 2 MB, the video queue the next priority, and the best-effort node the last priority. The voice queue is unlikely to drop because it has highest priority in the hierarchical scheduler as well as highest priority within its shared shaper. The video queue is less likely to drop, but you must still take care that the hierarchical scheduler is provisioned to allocate the proper assured bandwidth to video. The shared shaper can shape, or deny, bandwidth to its constituents, but it cannot allocate assured bandwidth in the hierarchical scheduler.

The compound shared-shaper mechanism also works as follows. In the legacy scheduler, weight and shaping rate are independent attributes that together determine bandwidth allocation. The scheduler allocates bandwidth based on relative weights, and the shaper can deny that bandwidth when the shaping rate is reached. With the shared shaper in effect, two independent shaping rates must be satisfied for the queue or node to dequeue. A deficit in either type of shaping bounds the bandwidth.

As a general way of predicting the scheduler behavior, if the physical port is congested because many queues and nodes are competing in the hierarchical scheduler, the legacy weights and shaping rates dominate the scheduler outcome. If the hierarchical scheduler is not congested, a shared shaper configured for a logical interface dominates the outcome for the traffic scheduled through that logical interface.

The compound shared shaper orders constituents, and allocates shared bandwidth to them, according to the following rules:

1. Strict constituents in the auto-strict-priority traffic-class group

For multiple strict-priority traffic-class groups, bandwidth allocation order is the same order in which the additional strict traffic class groups were configured. You can issue the **show traffic-class-groups** command to view this order.

2. Strict constituents in extended traffic-class groups

For multiple extended traffic class groups, bandwidth allocation order is the same order in which the traffic class groups were configured. You can issue the **show traffic-class-groups** command to view this order.

3. Strict constituents in the default group
4. Weighted constituents in the auto-strict-priority traffic class group
5. Weighted constituents in extended traffic class groups
6. Weighted constituents in the default group

By default, strict constituents transmit traffic at a rate up to the lesser of their shared-shaping rate or the legacy shaping rate. Individual strict constituents can be allocated any bandwidth value less than the shared rate. The sum of all constituent rate credits does not have to be less than the shared rate. Individual constituent rates are not capped, because a particular traffic class often does not exceed a limit because of admission control, or because the class is policed at some point in the path.

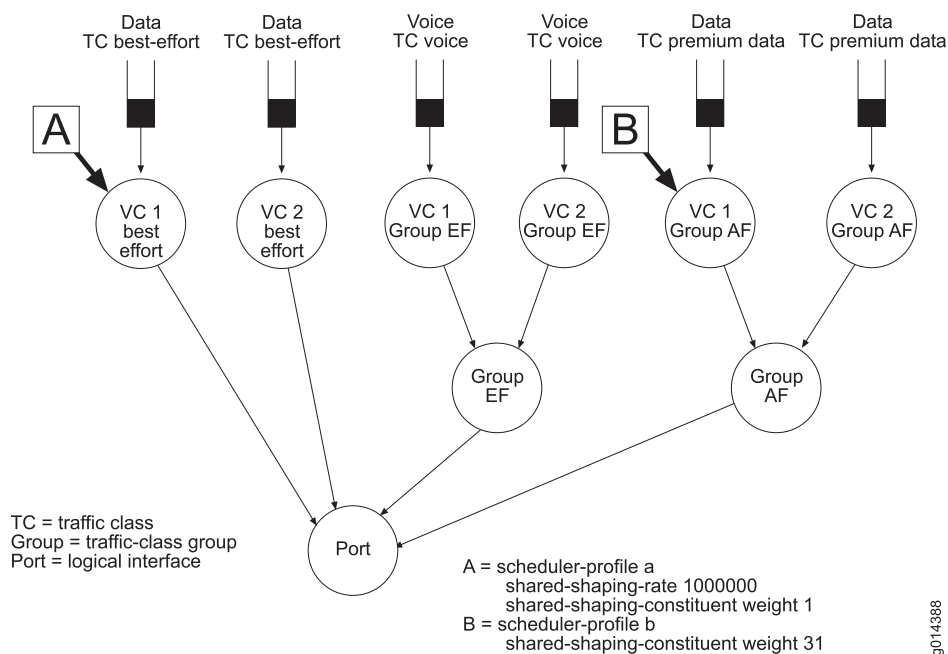
Unlike strict constituents, which can consume bandwidth up to the legacy shaping rate or the shared-shaping rate, weighted constituents share bandwidth with their peers solely in proportion to their shared-shaping-weight. A higher weight value grants the constituent a greater proportion of the available bandwidth.

Although a shared shaper can be applied to up to eight constituents, only four of these can be weighted constituents. If you configure more than four weighted constituents as part of the same shared shaper, the first four are treated as weighted constituents but the remainder are handled as strict constituents, generating a warning message.

Weighted Compound Shared Shaping Example

Weighted shared shaping is most useful for sharing bandwidth between traffic classes carrying TCP data. Figure 31 on page 113 shows an application of weighted shared shaping where weighted constituents span multiple traffic class groups, making them ineligible for legacy weighted scheduling. Best-effort data and premium data constituents are weighted.

Figure 31: Weighted Shared Shaping



Scheduler profile A specifies the shared-shaping rate of 1 Mbps for the best-effort node, which is associated with a VC logical interface. The node is further configured with a weight of 1. Scheduler profile B specifies the VC 1 AF node as a weighted constituent with a weight of 31.

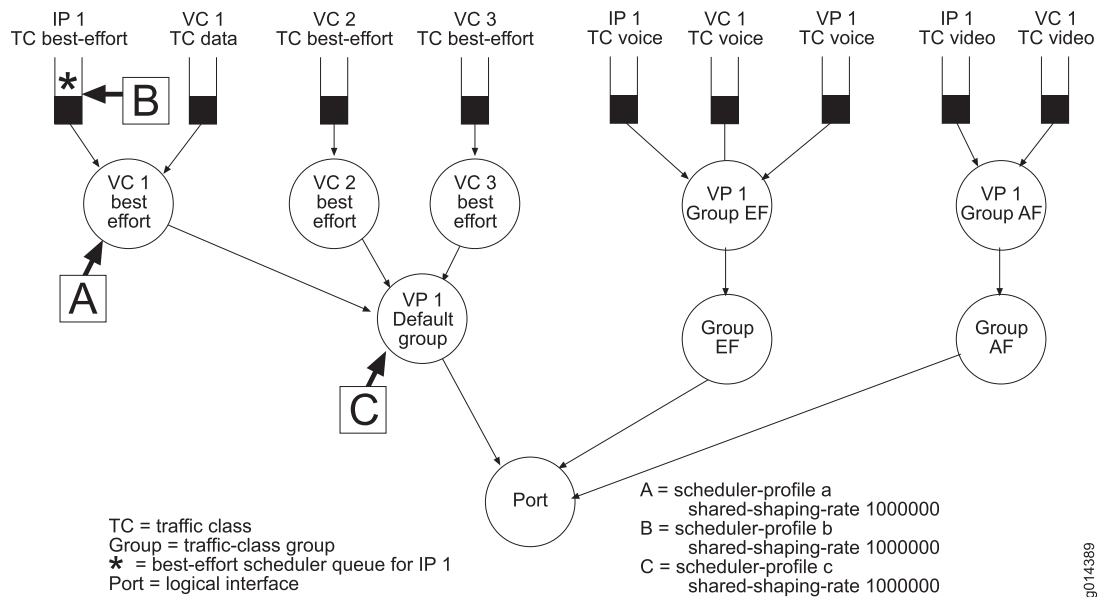
The implicitly selected constituents of the shared shaper are the VC 1 best-effort node, the VC 1 AF group node, and the VC 1 EF group node. Bandwidth is allocated as follows:

- The VC 1 EF group node is strict and can transmit up to the shared-shaping rate of 1 Mbps. Any remaining bandwidth is available to the remaining constituents.
- The VC 1 AF group node is weighted with the VC 1 best-effort node. The sum of the constituent weights is 32. With a weight of 31, the VC 1 AF group node can transmit 31/32nds of the available bandwidth when both constituents are competing for bandwidth.

- The VC 1 best-effort node is weighted with VC 1 AF group node. The sum of the constituent weights is 32. With a weight of 1, the VC 1 best-effort node can transmit 1/32 of the available bandwidth when both constituents are competing for bandwidth.

Figure 32 on page 114 illustrates an example of mixed interface shaping and its implications for implicit constituent selection for compound shared shaping.

Figure 32: Implicit Constituent Selection for Compound Shared Shaper: Mixed Interface Types



- Related Topics** ■ Configuring Implicit Constituents for Simple or Compound Shared Shaping on page 114

Configuring Implicit Constituents for Simple or Compound Shared Shaping

There are two types of implicit constituents:

- Simple implicit—Constituents are best-effort node or queues, and all nodes and queues in named traffic-class groups.
- Compound implicit—Constituents are selected automatically by the software. If a node exists in a given traffic-class group, the node is active and the queues stacked above it are inactive constituents.

Before you configure implicit constituents:

- Configure the traffic classes and traffic-class groups.

See “Configuring Traffic Classes That Define Service Levels” on page 14 and “Configuring Traffic-Class Groups That Define Service Levels” on page 15.

To configure implicit constituents:

1. Create the scheduler profile.

```
host1(config)#scheduler-profile implicit
```

2. Configure the shared shaper.

To configure a simple shared shaper:

```
host1(config-scheduler-profile)#shared-shaping-rate 128000 bps
```

To configure a compound shared shaper:

```
host1(config-scheduler-profile)#shared-shaping-rate 128000 burst 32767
compound
```

3. (Optional) For compound shared shapers, specify the attributes for the constituent.

```
host1(config-scheduler-profile)#shared-shaping-constituent weight 28
```

Including this command does not affect how the system selects the compound implicit constituent. If the command is present for a constituent that was implicitly selected, the software configures that constituent using the strict-priority or weight attributes.

After you configure implicit constituents:

- Configure the scheduler hierarchy with the best-effort nodes and queues.

See “Configuring a QoS Profile” on page 132.

Related Topics

- Constituent Selection for Shared Shaping Overview on page 107
- Implicit Constituent Selection Overview on page 109
- scheduler-profile
- shared-shaping-constituent
- shared-shaping-rate

Explicit Constituent Selection Overview

If you want only a subset of the queues for a subscriber to be shaped to the shared rate, then you must explicitly identify the desired constituents rather than accepting the implicitly selected constituents.

For compound shared shaping, explicit selection is also useful when you want queues as the active constituents instead of the node below them. By choosing queues you can assign appropriate priority or weights.

In the set of nodes and queues for a logical interface, only scheduler objects associated with a scheduler profile that includes a **shared-shaping-constituent** command are considered constituents. Objects that are not explicitly selected are exempt from the shared shaper.

To identify the constituents for simple shared shaping, include the **explicit-constituents** keyword with the **shared-shaping-rate simple** command in a scheduler profile that you associate with a best-effort node or queue to identify the logical interface.




NOTE: If you explicitly specify both a node and the queues stacked above the node as constituents in a scheduler profile, compound shared shaping uses the node as the constituent.

For compound shared shaping, omit the **simple** keyword. For a compound shared shaper, you can further designate the explicit constituents as priority or weighted.

Explicit Shared Shaping Example

In Figure 33 on page 117, two scheduler profiles are applied to scheduler objects VC 1 best effort node, VC 1 AF node, and VC 1 EF node. The **shared-shaping-constituent** command in each profile specifies that the associated object is an explicit constituent of the shared shaper.


Figure 33: Explicit Constituent Selection

In this example, the VC shared shaper has two explicit constituents, the VC 1 best effort node and the VC 1 Group EF node. By default, these constituents are considered to be strict constituents with a priority of 8.

If implicit selection rules are followed in this example, the association of the shared shaper with the VC 1 best-effort node selects the VC 1 best effort node, the VC 1 Group EF node, and the VC 1 Group AF node.

Explicit Weighted Compound Shared Shaping Example

Figure 34 on page 118 illustrates a case where scheduler profiles A, B, C, D, and E are applied to scheduler objects.

Figure 34: Case 1: Explicit Constituent Selection with Weighted Constituents

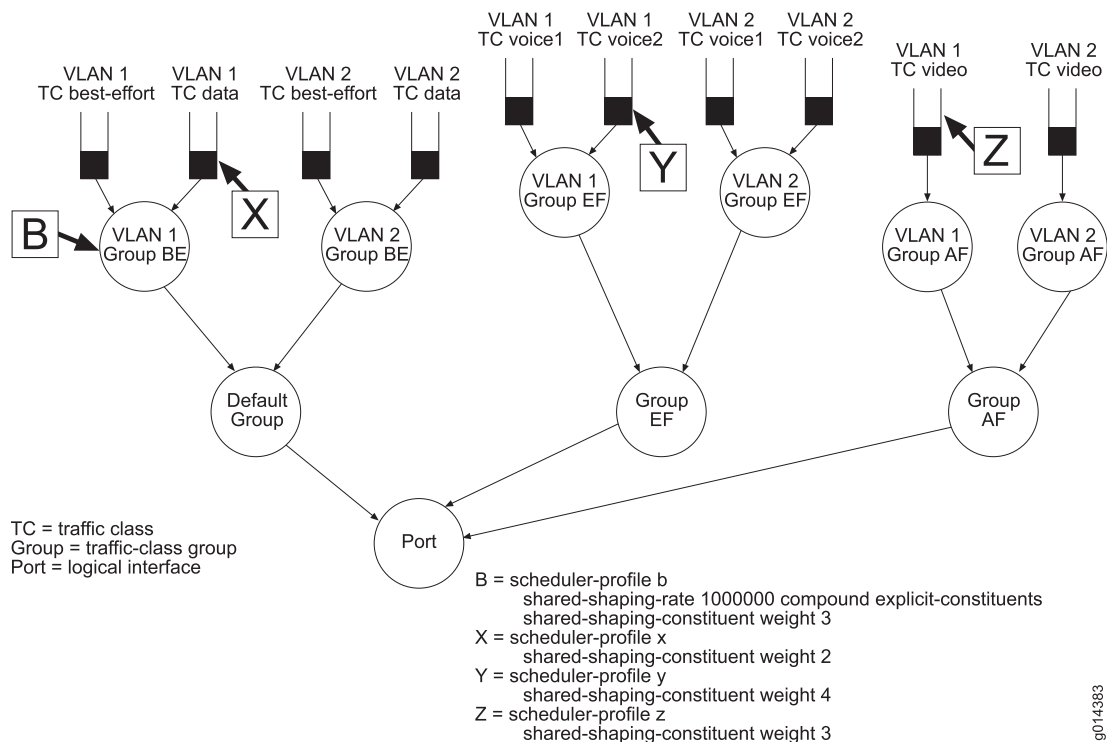
In Case 1, scheduler profile A associates the shared-shaping rate with the VLAN 1 best-effort queue. Table 12 on page 118 lists the explicit constituents of the shared shaper and the bandwidth allocated to each constituent:

Table 12: Bandwidth Allocation for Case 1 Explicit Constituents

Explicit Constituent	Bandwidth Allocation
VLAN 1 TC voice1 queue	Strict constituent that can consume up to its legacy shaping-rate.
VLAN 1 TC voice2 queue	Weighted constituent that shares bandwidth with its weighted shared shaper siblings in a proportion of 4/10.
VLAN 1 TC video queue	Weighted constituent that shares bandwidth with its weighted shared shaper siblings in a proportion of 3/10.
VLAN 1 TC data queue	Weighted constituent that shares bandwidth with its weighted shared shaper siblings in a proportion of 2/10.
VLAN 1 TC best-effort queue	Weighted constituent that shared bandwidth with weighted shared shaper siblings in a proportion of 1/10.

Figure 35 on page 119 illustrates another case where scheduler profiles B, X, Y, and Z are applied to scheduler objects. Each profile assigns a weight to an explicit constituent.

Figure 35: Case 2: Explicit Constituent Selection with Weighted Constituents



In Case 2, scheduler profile B associates the shared-shaping rate with the best-effort node for VLAN 1. Table 13 on page 119 lists the explicit constituents of the shared shaper and the bandwidth allocated to each constituent:

Table 13: Bandwidth Allocation for Case 2 Explicit Constituents

Explicit Constituent	Bandwidth Allocation
VLAN 1 TC voice1 queue	Strict constituent that can consume up to its legacy shaping-rate.
VLAN 1 TC voice2 queue	Weighted constituent that shares bandwidth with its weighted shared shaper siblings in a proportion of 4/10.
VLAN 1 TC video queue	Weighted constituent that shares bandwidth with its weighted shared shaper siblings in a proportion of 3/10.
Best-effort node for VLAN 1	Weighted constituent that shared bandwidth with weighted shared shaper siblings in a proportion of 3/10.

NOTE: The node is selected as the constituent when both the node and the queues stacked over node are specified in a scheduler profile.

- Related Topics** ■ Configuring Explicit Constituents for Simple or Compound Shared Shaping on page 120

Configuring Explicit Constituents for Simple or Compound Shared Shaping

You can specify explicit constituents and set the attributes of both implicit and explicit shared-shaping constituents that determine how bandwidth is allocated to them.

There are two types of explicit constituents:

- Simple explicit constituents—The software selects constituents based on the **shared-shaping-constituent** command. The weight and priority attributes of the **shared-shaping-constituent** command are ignored, because the simple shared shaper does not allocate bandwidth among constituents; instead it controls just the best-effort queue or node.
- Compound explicit—The software selects constituents based on the configured shared priority and shared weight in the **shared-shaping-constituent** command. If no attributes are specified, the software supplies a shared priority consistent with the legacy scheduler configuration. You can specify a constituent as strict (priority) or weighted. Strict-priority constituents are allocated bandwidth ahead of weighted constituents.

Before you configure explicit constituents:

- Configure the traffic classes and traffic-class groups.
See “Configuring Traffic Classes That Define Service Levels” on page 14 and “Configuring Traffic-Class Groups That Define Service Levels” on page 15.

To configure explicit constituents:

1. Create the scheduler profile.

```
host1(config)#scheduler-profile explicit
```

2. Configure the shared-shaper and specify that you do not want the router to identify shared shaper constituents associated with the logical interface.

To configure a simple shared shaper:

```
host1(config-scheduler-profile)#shared-shaping-rate 128000 bps
```

To configure a compound shared shaper:

```
host1(config-scheduler-profile)#shared-shaping-rate 128000 burst 32767
compound explicit-constituents
```

3. Specify the attributes for the explicit constituent.

```
host1(config-scheduler-profile)#shared-shaping-constituent weight 28
```

You can specify a constituent as strict (priority) or weighted. Strict-priority constituents are allocated bandwidth ahead of weighted constituents.

You can optionally set a value that determines the precedence of a constituent among its peers (strict or weighted) for claiming bandwidth.

For strict-priority constituents, the priority range is 1–8 and the default value is 8. A lower value correlates to a higher claim.

For weighted constituents, the range is 1–31 and the default value is 8. The weights of all sibling weighted constituents are added together. Then each weighted constituent is allocated bandwidth according to the proportion of its weight to the total.

- Related Topics**
- Constituent Selection for Shared Shaping Overview on page 107
 - Explicit Constituent Selection Overview on page 115
 - scheduler-profile
 - shared-shaping-constituent
 - shared-shaping-rate

Chapter 14

Monitoring a QoS Scheduler Hierarchy

This chapter provides information for configuring the QoS scheduler hierarchy using scheduler profiles on the E Series router.

QoS topics are discussed in the following section:

- Monitoring QoS Scheduling and Shaping on page 123

Monitoring QoS Scheduling and Shaping

To monitor QoS scheduling, see:

- Monitoring the QoS Scheduler Hierarchy on page 322
- Monitoring the Configuration of Scheduler Profiles on page 327
- Monitoring Shared Shapers on page 329
- Monitoring Shared Shaper Algorithm Variables on page 330

Part 4

Creating a QoS Scheduler Hierarchy on an Interface with QoS Profiles

- QoS Profile Overview on page 127
- Configuring and Attaching QoS Profiles to an Interface on page 131
- Configuring Shadow Nodes for Queue Management on page 149
- Monitoring a Scheduler Hierarchy on an Interface with QoS Profiles on page 155

Chapter 15

QoS Profile Overview

This chapter provides information for configuring an interface with QoS profiles on E Series routers.

QoS topics are discussed in the following sections:

- QoS Profile Overview on page 127
- Managing System Resources for Nodes and Queues on page 127
- Scaling Subscribers on the TFA ASIC with QoS on page 128

QoS Profile Overview

You create an interface hierarchy for QoS by configuring a QoS profile that specifies queue profiles, drop profiles, statistics profiles, and scheduler profiles in combination with interface types. A QoS profile specifies the queue, drop statistics gathering, and scheduler configuration for a subtree of the interface hierarchy. The QoS profile controls the way scheduler nodes, queues, and shadow nodes are bound to the interfaces above its attachment point in the interface hierarchy.

You can attach a QoS profile to the interface at the base of the subtree hierarchy, an ATM VP, or an S-VLAN. For example, a QoS profile attached to an ATM port specifies queuing attributes for interfaces of all types that are stacked over the port.

- Related Topics**
- Supported Interface Types for QoS Profiles on page 131
 - Configuring a QoS Profile on page 132

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 Topics**
- To identify the type of ASIC used by a line module, see the *ERX Module Guide* and the *E120 and E320 Module Guide*
 - Scaling Subscribers on the TFA ASIC with QoS on page 128
 - For more information about system resource requirements for shadow nodes, see Managing System Resources for Shadow Nodes on page 151
 - For information about egress memory available on ASIC line modules, see Memory Requirements for Queue and Buffers on page 19

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 Topics**
- For more information about system resource requirements for shadow nodes, see *Managing System Resources for Shadow Nodes* on page 151
 - For QoS system maximums, see *JUNOS Release Notes, Appendix A, System Maximums*
 - *Monitoring the QoS Profiles Attached to an Interface* on page 339

Chapter 16

Configuring and Attaching QoS Profiles to an Interface

This chapter provides information for configuring and attaching QoS profiles to an interface.

QoS topics are discussed in the following sections:

- Supported Interface Types for QoS Profiles on page 131
- Configuring a QoS Profile on page 132
- Attaching a QoS Profile to an Interface on page 134
- Munged QoS Profile Overview on page 136
- Example: Port-Type QoS Profile Attachment on page 139
- Example: QoS Profile Attachment to Port on page 141
- Example: Diffserv Configuration with Multiple Traffic-Class Groups on page 143

Supported Interface Types for QoS Profiles

Each QoS profile command begins with a keyword that designates an interface type.

Table 14 on page 131 lists the interface types and the commands that you can use with them.

Table 14: Interface Types and Supported Commands

Interface Type	Queue	Node	Group	Shadow Node
atm	✓	✓	✓	✓
atm-vc	✓	✓	–	✓
atm-vp	✓	✓	–	✓
bridge	✓	✓	–	✓
ethernet	✓	✓	✓	✓
fr-vc	✓	✓	–	✓

Table 14: Interface Types and Supported Commands *(continued)*

Interface Type	Queue	Node	Group	Shadow Node
ip	✓	✓	–	✓
ip-tunnel	✓	✓	–	✓
ipv6	✓	✓	–	✓
l2tp-session	✓	✓	–	✓
l2tp-tunnel	✓	✓	–	✓
lsp	✓	✓	–	✓
serial	✓	✓	✓	✓
server-port	✓	✓	✓	✓
svlan	✓	✓	–	✓
vlan	✓	✓	–	✓

- Related Topics**
- Supported Interface Types for QoS Profiles on page 131
 - Configuring a QoS Profile on page 132
 - Configuring Shadow Nodes on page 152

Configuring a QoS Profile

Before you configure a QoS profile:

- Configure the traffic classes.
See “Configuring Traffic Classes That Define Service Levels” on page 14.
- Configure the queuing hierarchy.
See “Configuring Queue Profiles to Manage Buffers and Thresholds” on page 22.
- Configure the scheduler hierarchy and shaping with scheduler profiles.
See “Configuring a Scheduler Hierarchy” on page 47.

To configure a QoS profile:

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

```
host1(config)#qos-profile qos-pc-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 Topics**
- Attaching a QoS Profile to an Interface on page 134
 - Supported Interface Types for QoS Profiles on page 131
 - Configuring Shadow Nodes on page 152
 - Monitoring a Scheduler Hierarchy on an Interface with QoS Profiles on page 155
 - For information about configuring QoS profiles with Service Manager, see *JUNOS Broadband Access Configuration Guide*
 - group
 - node

- qos-profile
- queue

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 134
- Attaching a QoS Profile to an ATM VP on page 134
- Attaching a QoS Profile to an S-VLAN on page 135
- Attaching a QoS Profile to a Port Type on page 135

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
```

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 Topics**
- Supported Interface Types for QoS Profiles on page 131
 - Configuring a QoS Profile on page 132
 - For information about attaching a QoS profile using Service Manager, see *JUNOS Broadband Access Configuration Guide*
 - atm-vp qos-profile
 - atm vp-tunnel
 - encapsulation vlan
 - interface
 - qos-port-type-profile
 - qos-profile
 - svlan id
 - svlan qos-profile

Munged QoS Profile Overview

QoS profile attachments affect the queuing configuration of all the forwarding interfaces stacked above the attachment point. The subtree of the interface hierarchy stacked above the attachment point is the scope of the attachment. When multiple QoS profiles are attached beneath a forwarding interface, the forwarding interface lies in the scope of all the QoS profiles. Rules from all the QoS profiles are combined in a process called mungeing. The set of rules used for a given forwarding interface is called the munged QoS profile.

When a QoS profile is attached to an interface, the router searches the interface stack, from the point of attachment down to the port interface at the base of the interface hierarchy, to find all QoS profiles attached under that interface. The rules are combined to form the munged QoS profile. The router reconfigures queues for all forwarding interfaces in the scope of the attachment to conform to the munged profile.

The munge algorithm works as follows:

1. Start with the rules in the QoS profile being attached.
2. Traverse down the stack of interfaces until another QoS profile attachment is found.
3. Add rules from the lower-attached QoS profile to the munged QoS profile. Conflicting rules from the lower-attached QoS profile are not added: rules in higher-attached QoS profiles override or eclipse rules in lower-attached QoS profiles.
4. Repeat Steps 2 and 3 until a port interface is reached at the bottom of the interface stack.
 - a. If there is a QoS profile attached at the port, add the profile's rules to the munged QoS profile, and the munge algorithm is then complete.

- b. If there is no QoS profile attached at the port, then locate the QoS profile indicated in the **qos-port-type-profile** command that corresponds to the interface type of the port. For example, if the port is an ATM interface, the default QoS port-type profile for type ATM is named atm-default. Add the rules in the QoS port-type profile to the munged QoS profile.

The entries in the QoS profile specified in the corresponding **qos-port-type-profile** command have the lowest precedence.

After the munged QoS profile is complete, the router reprocesses the queues for all forwarding interfaces in the scope of the attachment, adding, deleting, or modifying the scheduler hierarchy as required by the munged QoS profile rules. Conflicting node rules operate differently than this.

With conflicting node rules, the mungeing algorithm for QoS nodes start at the base of the interface hierarchy (usually near the physical interface), instead of at the top of the interface column. If a QoS profile is not attached to the port, nodes are added to the interface column according to the QoS port-type profile. Nodes are subsequently added from profiles that are attached higher in the interface column until all node rules from the interface column have been added, or the maximum hierarchy of three nodes has been reached. Higher level nodes can not eclipse lower-attached nodes. For example, if a QoS hierarchy is Ethernet node > Ethernet group node > VLAN node > queue, an IP node from a higher-attached QoS profile cannot eclipse the VLAN node.

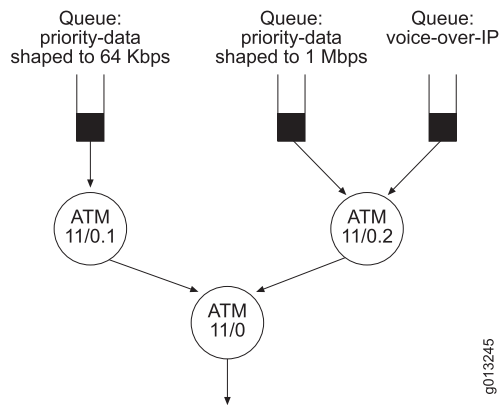
In Step 3, the router must decide which rules from a QoS profile conflict with rules already contained within the munged QoS profile. Queue rules are identified by their {interface type, traffic class} pair; two queue rules with the same interface type and traffic class are deemed conflicting. Node rules are identified by their {interface type, traffic-class group} pair; two node rules with the same interface type and traffic-class *group* are deemed conflicting.



NOTE: The munge algorithm is modified when you configure QoS for 802.3ad link aggregation interfaces.

Sample Munged QoS Profile Process

Figure 36 on page 138 shows the relationship between a port-attached QoS profile and a QoS profile that is attached to the specific interface, ATM 11/0.2.

Figure 36: Munged Profile Example

The port-attached QoS profile on ATM 11.0 contains the following queue rule:

```
host1(config)#qos-profile atmPort
host1(config-qos-profile)#ip queue traffic-class priority-data scheduler-profile 64kbps
host1(config-qos-profile)#exit
```

All forwarding interfaces stacked above the port are within the scope of the attachment, so all IP interfaces stacked above the port will be provisioned with a queue in the priority-data traffic class, shaped to 64 Kbps.

The QoS profile attached at subinterface ATM 11/0.2 contains the following two rules:

```
host1(config)#qos-profile atmVc
host1(config-qos-profile)#ip queue traffic-class priority-data scheduler-profile 1mbps
host1(config-qos-profile)#ip queue traffic-class voice-over-ip
host1(config-qos-profile)#exit
```

The queue rule for {interface type IP, traffic-class priority-data} in the QoS profile that is attached to ATM 11/0.2 effectively overrides the queue rule for the same interface type and traffic class in the port-attached QoS profile on ATM 11.0.

The second queue rule, which is for the voice-over-ip traffic-class, is not conflicting. In this configuration, the provider has configured a 64 Kbps priority-data queue for each IP interface stacked above the port. But the IP interface above the ATM 11/0.2 attachment provides 1 Mbps for priority-data, and also has a second queue provisioned for VoIP.



NOTE: When a QoS profile is attached to an interface, the router first searches to determine if a munged QoS profile already exists. If you modify an existing QoS profile, the router automatically updates all munged QoS profiles that are dependent on the modified profile.

- Related Topics** ■ For more information about the munge algorithm and 802.3ad link aggregation interfaces, see Munged QoS Profiles and Load Balancing on page 184

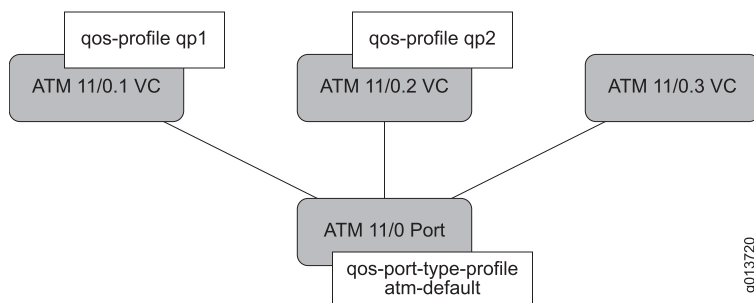
Example: Port-Type QoS Profile Attachment

In this example, three ATM subinterfaces are configured on an ATM port:

- ATM 11/0.1—QoS profile qp1 is attached
- ATM 11/0.2—QoS profile qp2 is attached
- ATM 11/0.3—No QoS profile is attached

The major ATM interface, 11/0, does not have a QoS profile explicitly attached. Therefore, by default the atm-default QoS port-type profile is attached.

Figure 37: Attaching QoS Profiles to ATM Subinterfaces



To configure this example:

1. Create and configure QoS profile qp1.

```
host1(config)#qos-profile qp1
host1(config-qos-profile)#atm-vp node scheduler-profile sp1
host1(config-qos-profile)#atm-vc queue traffic-class tc1 scheduler-profile sp1
queue-profile qp1
host1(config-qos-profile)#atm-vc queue traffic-class tc2 scheduler-profile sp2
queue-profile qp2
host1(config-qos-profile)#atm-vc queue traffic-class tc3 scheduler-profile sp3
queue-profile qp3
host1(config-qos-profile)#atm-vc queue traffic-class tc4 scheduler-profile sp4
queue-profile qp4
host1(config-qos-profile)#atm-vc queue traffic-class tc5 scheduler-profile sp5
queue-profile qp5
host1(config-qos-profile)#exit
```

2. Create and configure QoS profile qp2.

```
host1(config)#qos-profile qp2
host1(config-qos-profile)#atm-vp node scheduler-profile sp1
host1(config-qos-profile)#atm-vc queue traffic-class tc1 scheduler-profile sp1
queue-profile qp1
```

```

host1(config-qos-profile)#atm-vc queue traffic-class tc2 scheduler-profile sp2
queue-profile qp2
host1(config-qos-profile)#atm-vc queue traffic-class tc3 scheduler-profile sp3
queue-profile qp3
host1(config-qos-profile)#exit

```

3. Attach the QoS profiles to the ATM subinterfaces, as shown in Figure 37 on page 139.

```

host1(config)#interface atm 11/0.1
host1(config-subif)#qos-profile qp1
host1(config-subif)#exit
host1(config)#interface atm 11/0.2
host1(config-subif)#qos-profile qp2
host1(config-subif)#exit

```

4. Display the QoS interface hierarchy for ATM interface 11/0. This display shows all QoS attachments above interface 11/0.

If no QoS profiles are attached above the specified interface, the router shows the first attachment below the specified interface.

```
host1# show qos interface-hierarchy interface atm 11/0
```

```
attachment@ atm-vc ATM11/0.2:
```

qos profile	interface type	rule type	traffic class	scheduler profile	queue profile	t-class group
qp2@ATM11/0.2	atm-vp	node		sp1	default	
qp2@ATM11/0.2	atm-vc	queue	tc1	sp1	qp1	
qp2@ATM11/0.2	atm-vc	queue	tc2	sp2	qp2	
qp2@ATM11/0.2	atm-vc	queue	tc3	sp3	qp3	
atm-default @atm ip		node		default	default	
atm-default @atm atm-vc		node		default	default	
atm-default @atm Bridge		node		default	default	
atm-default @atm ipv6		node		default	default	
atm-default @atm ip		queue	best-effort	default	default	
atm-default @atm atm		queue	best-effort	default	default	
atm-default @atm atm-vc		queue	best-effort	default	default	
atm-default @atm Bridge		queue	best-effort	default	default	
atm-default @atm ipv6		queue	best-effort	default	default	

```
attachment@ atm-vc ATM11/0.1:
```

qos profile	interface type	rule type	traffic class	scheduler profile	queue profile	t-class group
qp1@ATM11/0.1	atm-vp	node		sp1	default	
qp1@ATM11/0.1	atm-vc	queue	tc1	sp1	qp1	
qp1@ATM11/0.1	atm-vc	queue	tc2	sp2	qp2	
qp1@ATM11/0.1	atm-vc	queue	tc3	sp3	qp3	
qp1@ATM11/0.1	atm-vc	queue	tc4	sp4	qp4	
qp1@ATM11/0.1	atm-vc	queue	tc5	sp5	qp5	
atm-default @atm ip		node		default	default	
atm-default @atm atm-vc		node		default	default	
atm-default @atm Bridge		node		default	default	
atm-default @atm ipv6		node		default	default	
atm-default @atm ip		queue	best-effort	default	default	
atm-default @atm atm		queue	best-effort	default	default	
atm-default @atm atm-vc		queue	best-effort	default	default	

```

atm-default @atm Bridge      queue  best-effort  default  default
atm-default @atm ipv6        queue  best-effort  default  default

```

ATM subinterface 11/0.3 was not shown because no QoS profile is attached to it. You can display the QoS interface hierarchy for subinterface 11/0.3 by specifying the subinterface, as shown below. In this case, the QoS port-type profile, atm-default, is attached (by default) to the ATM major interface, ATM 11/0, below ATM subinterface 11/0.3. Because no QoS profile is attached to this ATM subinterface, the QoS port-type profile is applied.

The @atm in the qos profile column indicates that the row comes from a default QoS port-type profile that is below the interfaces shown: subinterfaces ATM 11/0.2 and ATM 11/0.1 in this example.

You can explicitly show the ATM subinterface that has no explicit QoS profile attachment, as shown below. In this case, attachment@ indicates the ATM major interface (11/0) below the subinterface.

```

host1# show qos interface-hierarchy interface atm 11/0.3
attachment@ atm ATM11/0:
  qos      interface  rule  traffic      scheduler  queue  t-class
profile    type      type  class        profile    profile group
-----
atm-default@atm  ip      node
atm-default@atm  atm-vc  node
atm-default@atm  Bridge  node
atm-default@atm  ipv6    node
atm-default@atm  ip      queue best-effort
atm-default@atm  atm     queue best-effort
atm-default@atm  atm-vc  queue best-effort
atm-default@atm  Bridge  queue best-effort
atm-default@atm  ipv6    queue best-effort

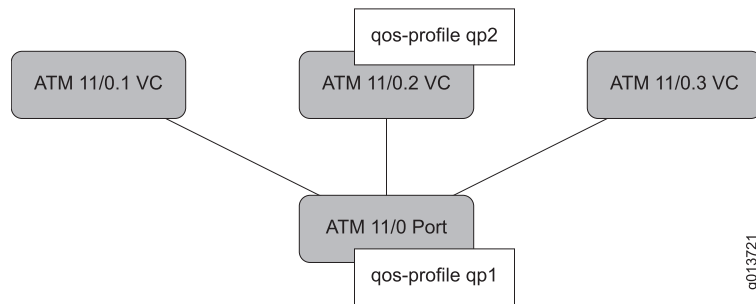
```

Example: QoS Profile Attachment to Port

In Figure 38 on page 142, the major ATM interface, 11/0, has QoS profile qp1 explicitly attached. The major ATM interface has three ATM subinterfaces configured:

- ATM 11/0.1—No QoS profile is explicitly attached
- ATM 11/0.2—QoS profile qp2 is attached
- ATM 11/0.3—No QoS profile is explicitly attached

The qp1 profile overrides the QoS port-type profile, atm-default, on subinterfaces 1 and 3. It does not override profile qp2, which was explicitly attached to subinterface 2.

Figure 38: Attaching QoS Profile to ATM Interface and Subinterface

To configure this example:

1. Create and configure QoS profile qp1.

```

host1(config)#qos-profile qp-1
host1(config-qos-profile)#atm-vp node scheduler-profile sp1
host1(config-qos-profile)#atm-vc queue traffic-class tc1 scheduler-profile sp1
queue-profile qp1
host1(config-qos-profile)#atm-vc queue traffic-class tc2 scheduler-profile sp2
queue-profile qp2
host1(config-qos-profile)#atm-vc queue traffic-class tc3 scheduler-profile sp3
queue-profile qp3
host1(config-qos-profile)#atm-vc queue traffic-class tc4 scheduler-profile sp4
queue-profile qp4
host1(config-qos-profile)#atm-vc queue traffic-class tc5 scheduler-profile sp5
queue-profile qp5
host1(config-qos-profile)#exit

```

2. Create and configure QoS profile qp2.

```

host1(config)#qos-profile qp2
host1(config-qos-profile)#atm-vp node scheduler-profile sp1
host1(config-qos-profile)#atm-vc queue traffic-class tc1 scheduler-profile sp1
queue-profile qp1
host1(config-qos-profile)#atm-vc queue traffic-class tc2 scheduler-profile sp2
queue-profile qp2
host1(config-qos-profile)#atm-vc queue traffic-class tc3 scheduler-profile sp3
queue-profile qp3
host1(config-qos-profile)#exit

```

3. Attach QoS profile qp1 to ATM interface 11/0.

```

host1(config)#interface atm 11/0
host1(config-if)#qos-profile qp1
host1(config-if)#exit

```

4. Attach QoS profile qp2 to ATM subinterface 11/0.2.

```

host1(config)#interface atm 11/0.2
host1(config-subif)#qos-profile qp2
host1(config-subif)#exit
host1(config)#exit

```

5. Display the QoS interface hierarchy for ATM 11/0.

```

host1#show qos interface-hierarchy interface atm 11/0
qos          interface rule  traffic      scheduler   queue  t-class
profile      type      type  class       profile     profile group
-----
@ATM11/0     atm        queue  best-effort  default    default
qp1@ATM11/0  atm-vp     node                   sp1         default
qp1@ATM11/0  atm-vc     queue  tc1          sp1         qp1
qp1@ATM11/0  atm-vc     queue  tc2          sp2         qp2
qp1@ATM11/0  atm-vc     queue  tc3          sp3         qp3
qp1@ATM11/0  atm-vc     queue  tc4          sp4         qp4
qp1@ATM11/0  atm-vc     queue  tc5          sp5         qp5
attachment@  atm-vc ATM11/0.2:
qos          interface rule  traffic      scheduler   queue  t-class
profile      type      type  class       profile     profile group
-----
qp2@ATM11/0.2 atm-vp     node                   sp1         default
qp2@ATM11/0.2 atm-vc     queue  tc1          sp1         qp1
qp2@ATM11/0.2 atm-vc     queue  tc2          sp2         qp2
qp2@ATM11/0.2 atm-vc     queue  tc3          sp3         qp3
@ATM11/0     atm        queue  best-effort  default    default
qp1@ATM11/0  atm-vc     queue  tc4          sp4         qp4
qp1@ATM11/0  atm-vc     queue  tc5          sp5         qp5

```

Note that:

- ATM best-effort queues are created on ATM interface @ATM11/0 and ATM 11/0.2.
- ATM 11/0.2 subinterface has three queues (traffic classes tc1, tc2, and tc3) that come from QoS profile qp2. Traffic class tc3 is defined in both QoS profile qp1 and qp2. The QoS profile attached closest to the leaf node is used, however. Traffic class tc3 comes from QoS profile qp2, which is attached to ATM subinterface ATM 11/0.2.
- Queues for traffic classes tc4 and tc5 come from QoS profile qp1, which is attached at the ATM major interface.

Example: Diffserv Configuration with Multiple Traffic-Class Groups

In this example configuration, a service provider offers three types of service: data, video-on-demand, and voice. Each service has different QoS requirements. The data users log in and can dynamically subscribe to video and voice services. The data service is a best-effort service. The video service is a *better than best effort* service, which corresponds to assured forwarding PHB. The voice service is a low-latency service, which corresponds to expedited forwarding PHB.

You can meet these varying traffic requirements by creating a traffic class group for each of the three services. Creating groups enables you to apply QoS to the group nodes. For example, you could specify the following:

- The voice service gets low-latency, strict priority treatment through the fabric and on egress. You configure an assured rate of 20 Mbps, and shape the traffic to 20 Mbps. Each voice user is shaped to 1 Mbps to support up to 20 voice

subscribers without oversubscription. Call admission control ensures that there are no more than 20 simultaneous voice service subscribers. Unused bandwidth is divided among the video and best-effort users.

- The video service is scheduled by the HRR scheduler and gets the hierarchical assured rate. You shape the video traffic to 50 Mbps. Each video service user is assured 1 Mbps, and is shaped to 1 Mbps to support up to 50 video subscribers without oversubscription. Call admission control ensures that there are no more than 50 simultaneous video service subscribers. Unused bandwidth is divided among the best-effort users.
- The best-effort data service is scheduled by the HRR scheduler and gets the bandwidth left over from the voice and video services.

Configure this implementation as follows.

1. Create the video and voice traffic classes. Assign the voice traffic class a strict-priority treatment within the fabric. Manually creating a best-effort traffic class is superfluous because the router creates this class by default.

```
(config)#traffic-class video
(config-traffic-class)#exit
(config)#traffic-class voice
(config-traffic-class)#fabric-strict-priority
(config-traffic-class)#exit
(config)#traffic-class best-effort
(config-traffic-class)#exit
```

2. Create scheduler profiles for the assured forwarding, expedited forwarding, and best-effort groups. Specify strict priority scheduling for the expedited forwarding traffic and shape it to 20 Mbps.

```
(config)#scheduler-profile expeditedGroup
(config-scheduler-profile)#strict-priority
(config-scheduler-profile)#shaping-rate 20000000
(config-scheduler-profile)#assured-rate 20000000
(config-scheduler-profile)#exit
```

3. Assured traffic is not strict, so it is scheduled by the HRR scheduler. Shape the assured traffic to 50 Mbps, and specify the hierarchical assured rate to give assured traffic preferential treatment over best-effort traffic.

```
(config)#scheduler-profile assuredGroup
(config-scheduler-profile)#shaping-rate 50000000
(config-scheduler-profile)#assured-rate hierarchical
(config-scheduler-profile)#exit
```

4. Best effort traffic is also scheduled by the HRR scheduler. You do not apply any shaping for this traffic because it simply gets the leftover bandwidth.

```
(config)#scheduler-profile bestEffortGroup
(config-scheduler-profile)#exit
```

5. Create scheduler profiles for the voice, video, and best-effort service classes. Shape voice and video to 1 Mbps. Because you do not specify a shaping rate, the best-effort traffic can borrow unused bandwidth.


```

(config)#scheduler-profile voice
(config-scheduler-profile)#shaping-rate 1000000
(config-scheduler-profile)#exit
(config)#scheduler-profile video
(config-scheduler-profile)#shaping-rate 1000000
(config-scheduler-profile)#exit
(config)#scheduler-profile best-effort
(config-scheduler-profile)#exit

```

6. Put the video traffic class into the assured-forwarding traffic-class group and specify the group as strict priority. Put the voice traffic class into the expedited-forwarding traffic-class group. Put the best-effort traffic class into the best-effort traffic-class group.

```

(config)#traffic-class-group assured-forwarding auto-strict-priority
(config-traffic-class-group)#traffic-class video
(config-traffic-class-group)#exit
(config)#traffic-class-group expedited-forwarding extended
(config-traffic-class-group)#traffic-class voice
(config-traffic-class-group)#exit
(config)#traffic-class-group best-effort extended
(config-traffic-class-group)#traffic-class best-effort
(config-traffic-class)#exit

```

7. Create a QoS profile that contains the group rules for the assured-forwarding, expedited-forwarding, and best-effort traffic-class groups.

```

(config)#qos-profile qpDiffServExample
(config-qos-profile)#ethernet group assured-fwd scheduler-profile assuredGroup
(config-qos-profile)#ethernet group expedited-fwd scheduler-profile
expeditedGroup
(config-qos-profile)#ethernet group best-effort scheduler-profile bestEffortGroup
(config-qos-profile)#ip node group assured-fwd scheduler-profile default
(config-qos-profile)#ip node group expedited-fwd scheduler-profile default
(config-qos-profile)#ip node group best-effort scheduler-profile default
(config-qos-profile)#ip queue traffic-class voice scheduler-profile voice
(config-qos-profile)#ip queue traffic-class video scheduler-profile video
(config-qos-profile)#ip queue traffic class best-effort scheduler-profile best-effort
(config-qos-profile)#exit

```

8. Attach the QoS profile to an Ethernet port.

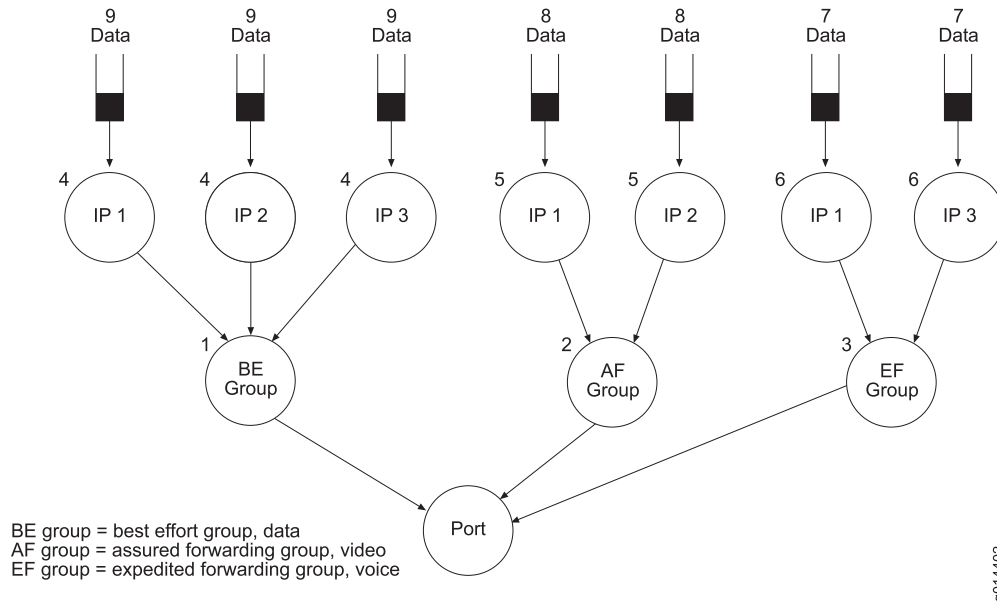
```

(config)#interface fastEthernet 9/0
(config-if)#qos-profile qpDiffServExample
(config-if)#exit

```

Figure 39 on page 146 shows this configuration with 3 users: IP 1, IP 2, and IP 3.

- IP 1 subscribes to data, video, and voice services.
- IP 2 subscribes to data and video services.
- IP 3 subscribes to data and voice services.

Figure 39: Diffserv Configuration with Multiple Traffic-Class Groups

The following set of commands configures the QoS profile detailed in Step 7 previously. Each line in the profile is known as a *profile rule*. The numbers associated with each rule correspond to the numbers in Figure 39 on page 146.

```
(config)#qos-profile qpDiffServExample
(1) (config-qos-profile)#ethernet group best-effort scheduler-profile bestEffortGroup
(2) (config-qos-profile)#ethernet group assured-fwd scheduler-profile assuredGroup
(3) (config-qos-profile)#ethernet group expedited-fwd scheduler-profile expeditedGroup
(4) (config-qos-profile)#ip node group best-effort scheduler-profile default
(5) (config-qos-profile)#ip node group assured-fwd scheduler-profile default
(6) (config-qos-profile)#ip node group expedited-fwd scheduler-profile default
(7) (config-qos-profile)#ip queue traffic-class voice scheduler-profile voice
(8) (config-qos-profile)#ip queue traffic-class video scheduler-profile video
(9) (config-qos-profile)#ip queue traffic class best-effort scheduler-profile best-effort
```

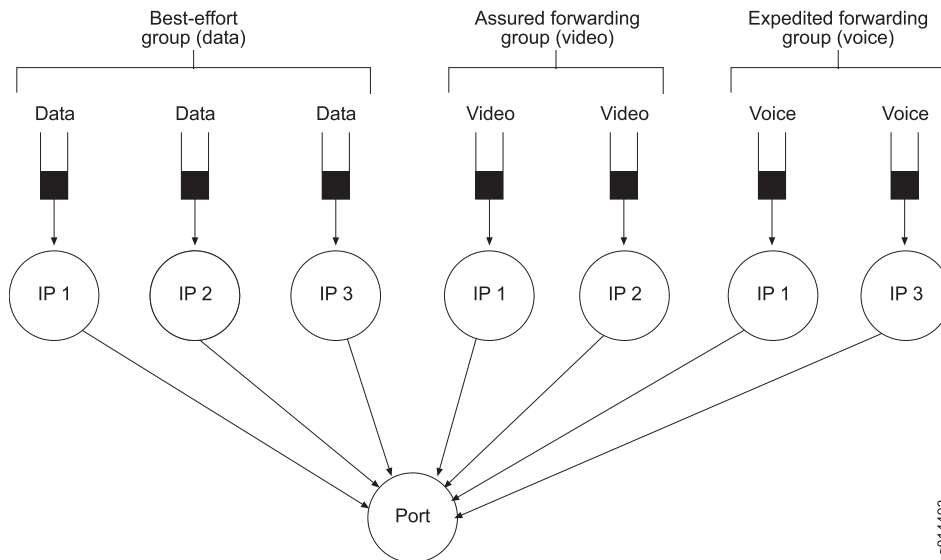
When you specify a group rule within an attached QoS profile, nodes and queue may be attached to group nodes. If the qpDiffServExample QoS profile used in the preceding example did not contain group rules, then the groups would exist with no attachments.

For example, the following set of commands configures the same QoS profile, but with the group removed, as shown in Figure 40 on page 147.

```
(config)#qos-profile qpDiffServExample
(config-qos-profile)#ip node scheduler-profile default
config-qos-profile)#ip queue traffic-class voice scheduler-profile voice
config-qos-profile)#ip queue traffic-class video scheduler-profile video
config-qos-profile)#ip queue traffic class best-effort scheduler-profile best-effort
```

In this case, the configuration creates the groups but does not place any of the traffic classes into the groups. Figure 40 on page 147 shows that IP 1, IP 2, and IP 3 contain the ungrouped traffic classes, data, video, and voice.

Figure 40: Diffserv Configuration Without Traffic-Class Groups



Because the BE, AF, and EF groups have no queues, their scheduler attributes (weight, assured rate, shaping rate) do not affect the HRR scheduler's distribution of bandwidth.

Chapter 17

Configuring Shadow Nodes for Queue Management

This chapter provides information for configuring shadow nodes on E Series routers.

QoS topics are discussed in the following sections:


- Shadow Node Overview on page 149
- Shadow Nodes and Scheduler Behavior on page 150
- Managing System Resources for Shadow Nodes on page 151
- Configuring Shadow Nodes on page 152
- Example: Shadow Nodes over VLAN and IP Queues on page 153
- Example: Shadow Nodes on the Same Traffic-Class Group on page 154
- Example: Shadow Nodes on Different Traffic-Class Groups on page 154

Shadow Node Overview

The frame forwarding ASIC (FFA) and the 10-Gigabit Ethernet forwarding ASIC (TFA) require that all queues be above the port scheduler node with two additional scheduler nodes. The router implicitly creates *phantom nodes* when you do not specify two scheduler nodes above the port interface. Phantom nodes cannot be monitored using **show** commands.

Phantom nodes have the same weight as the associated queues and are not shaped, which preserves the behavior of the queues as if they are at their original level.

Figure 41 on page 150 compares a scheduler hierarchy with and without phantom nodes.

Figure 41: Phantom Nodes

The first scheduler hierarchy displayed in Figure 41 on page 150 shows Queue A, Queue B, and Node C at the same scheduler level and with the same weight of 8. They equally share the bandwidth available to the level 1 node.

The second scheduler hierarchy in Figure 41 on page 150 shows the phantom nodes the router added for Queue A and Queue B. It also shows the weight associated with Queue A and Queue B. As the result, Phantom A, Phantom B, and Node C share the bandwidth of the level 1 node. The phantom nodes do not change the behavior of Queue A and Queue B.

Shadow Nodes and Scheduler Behavior

You can configure *shadow nodes* when you want to explicitly set the queues at the required scheduler level for any line module with the EFA, EFA2, FFA, or TFA hardware. Shadow nodes enable you to specify the weight and the shaping rate of the added node. Shadow nodes can also conserve scheduler node resources.

You define the shadow node by referencing the shadow node in the QoS profile. Like phantom nodes, the router creates shadow nodes only when the additional node is required to meet the proper queue level.

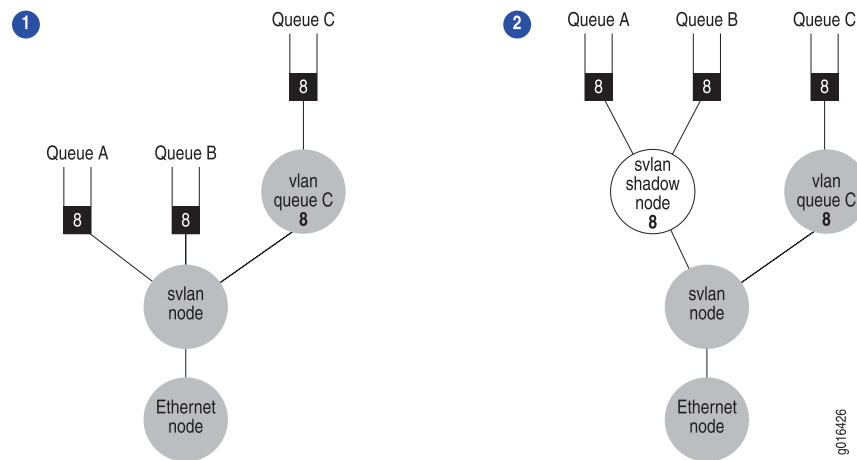
The router creates shadow nodes after all the nodes and group nodes are created, and only when a node of the same interface type has existed in the same group of the scheduler hierarchy.

Shadow nodes can be configured for all interface types available for nodes.



NOTE: Shadow nodes ignore any shared-shaping rates in a scheduler profile.

Figure 42 on page 151 compares a scheduler hierarchy with and without shadow nodes.

Figure 42: Shadow Nodes

Unlike phantom nodes, shadow nodes can alter the behavior of the scheduler.

The first scheduler hierarchy in Figure 42 on page 151 shows VLAN interfaces A, B and C stacked above the same S-VLAN interface. Interfaces A and B have the same scheduler hierarchy (referencing qos-profile AB) and have a VLAN queue stacked directly above the S-VLAN node. In this case, VLAN interfaces A, B and C share the same 33 percent bandwidth available to the S-VLAN node.

Interface C has a VLAN queue stacked above a VLAN node and the S-VLAN node (referencing qos-profile C).

Specifying a shadow node forces the VLAN queue to the proper scheduler level. The second scheduler hierarchy in Figure 42 on page 151 shows the shadow node that is applied after QoS profile AB-shadow is assigned to interfaces A and B. As a result, interfaces A and B have 25 percent of the S-VLAN bandwidth and interface C has 50 percent of the S-VLAN bandwidth.

The S-VLAN shadow node uses the same scheduler profile as the queue.

To provide interfaces A and B with the proper weight, configure the weight of the shadow node to the sum of its queue weight. You can use hierarchical parameter instances and weight expressions to configure an appropriate weight.

- Related Topics**
- For a list of interface types supported for shadow nodes, see Supported Interface Types for QoS Profiles on page 131
 - For more information about hierarchical parameters, see Hierarchical QoS Parameters Overview on page 261

Managing System Resources for Shadow Nodes

Each ASIC hardware type provides different node and queue resources.

Level 1 queues stack directly above the port; level 2 queues stack above a node and the port. The router implicitly creates the level 1 and level 2 queues.

Shadow node queues stack above a port node, a level 1 node, and a shadow node. Therefore, the shadow node queue is at level 3. The router does not implicitly create any nodes for the queues.

You can configure 64,000 level 1 queues using shadow nodes by specifying the group and shadow node rules in the QoS profile. Each level 1 queue is stacked above the port, the group node, and the shadow node; therefore, it requires 64,002 descriptors.

Table 15 on page 152 lists the number of nodes required to create a queue.

Table 15: Shadow Node Consumption of Node and Queue Resources

	Level 1 Queues (at Port)	Level 2 Queues (at Node)	Shadow Node Queue
Required Nodes	3	2	1

- Related Topics**
- Managing System Resources for Nodes and Queues on page 127
 - Scaling Subscribers on the TFA ASIC with QoS on page 128

Configuring Shadow Nodes

Before you configure shadow nodes:

- Configure the traffic classes.
See “Configuring Traffic Classes That Define Service Levels” on page 14.
- Configure the queuing hierarchy.
See “Configuring Queue Profiles to Manage Buffers and Thresholds” on page 22.
- Configure the scheduler hierarchy and shaping with scheduler profiles.
See “Configuring a Scheduler Hierarchy” on page 47.

To add a shadow node to a QoS profile:

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

```
host1(config)#qos-profile shadowNode
host1(config-qos-profile)#
```

2. Configure a scheduler node for each interface of the specified type.

```
host1(config-qos-profile)#atm node scheduler-profile default
```

3. Configure a shadow node for each interface of the specified type.


```
host1(config-qos-profile)#atm shadow-node scheduler-profile default
```

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

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

5. (Optional) Configure a traffic-class group and reference a scheduler profile in the QoS profile.

```
host1(config-qos-profile)#atm group default scheduler-profile default
```

The router creates the shadow node when the following conditions are met:

- After all the nodes and group nodes are created.
- If the queues are not at the required scheduler level.
- When a node of the same interface type has existed in the same group of the scheduler hierarchy.

Related Topics

- Shadow Node Overview on page 149
- Shadow Nodes and Scheduler Behavior on page 150
- Managing System Resources for Shadow Nodes on page 151
- Monitoring a Scheduler Hierarchy on an Interface with QoS Profiles on page 155
- group
- node
- qos-profile
- queue
- shadow-node

Example: Shadow Nodes over VLAN and IP Queues

This example illustrates how and when the system creates shadow node after you configure it.

In the first part, you specify an Ethernet node, an Ethernet group node, a VLAN node, a VLAN shadow node, and an IP queue. Because the IP queue is at a proper scheduler level without the shadow node, the system does not create a shadow node.

```
host1(config-qos-profile)#ethernet node
host1(config-qos-profile)#ethernet group default scheduler-profile default
host1(config-qos-profile)#vlan node
host1(config-qos-profile)#vlan shadow-node
host1(config-qos-profile)#ip queue traffic-class best-effort scheduler-profile default
```

In the second part, you specify an Ethernet node, a VLAN node, a shadow node, and a VLAN queue. The system creates the shadow node so that the VLAN queue is at the proper scheduler level.

```
host1(config-qos-profile)#ethernet node
host1(config-qos-profile)#vlan node
host1(config-qos-profile)#vlan shadow-node
host1(config-qos-profile)#vlan queue traffic-class best-effort scheduler-profile default
```

Example: Shadow Nodes on the Same Traffic-Class Group

This example demonstrates how to configure the shadow nodes on the same traffic-class group.

You specify a VLAN node, an IP node, an IP video queue, and a best-effort Ethernet queue. The system adds the Ethernet node, the VLAN node, the IP node, and the IP video queue to the scheduler hierarchy. Even though the two queues belong to the same traffic-class group, the Ethernet best-effort queue is stacked above the shadow node and the IP video queue is stacked above the IP node.

```
host1(config-qos-profile)#ethernet node
host1(config-qos-profile)#ethernet shadow-node scheduler profile shadow
host1(config-qos-profile)#ethernet queue traffic-class best-effort scheduler-profile default
host1(config-qos-profile)#vlan node
host1(config-qos-profile)#ip node
host1(config-qos-profile)#ip queue traffic-class video scheduler-profile default
```

Example: Shadow Nodes on Different Traffic-Class Groups

This example shows how to configure shadow nodes on different traffic-class groups. After adding the voice queue in the auto-strict priority group named strict, the system stacks the IP voice queue above the Ethernet port, the voice group, and the phantom node.

```
host1(config-qos-profile)#ethernet node
host1(config-qos-profile)#ethernet shadow-node scheduler profile shadow
host1(config-qos-profile)#ethernet queue traffic-class best-effort scheduler-profile default
host1(config-qos-profile)#vlan node
host1(config-qos-profile)#ip node
host1(config-qos-profile)#ip queue traffic-class video scheduler-profile default
host1(config-qos-profile)#ethernet group voice-group scheduler-profile strict
host1(config-qos-profile)#ip queue traffic-class voice scheduler-profile default
```

Chapter 18

Monitoring a Scheduler Hierarchy on an Interface with QoS Profiles

This chapter provides information for monitoring a scheduler hierarchy on an interface with QoS profiles.

QoS topics are discussed in the following sections:

- Monitoring a Scheduler Hierarchy on an Interface with QoS Profiles on page 155

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 339
- Monitoring the Configuration of QoS Port-Type Profiles on page 340
- Monitoring the Configuration of QoS Profiles on page 341
- Monitoring the Configuration of Scheduler Profiles on page 327
- Monitoring QoS Parameter Instances on page 352

Part 5

Interface Solutions for QoS

- Configuring an Integrated Scheduler to Provide QoS for ATM on page 159
- Configuring QoS for Gigabit Ethernet Interfaces and VLAN Subinterfaces on page 177
- Configuring QoS for 802.3ad Link Aggregation Groups on page 183
- Configuring QoS for L2TP Sessions on page 197
- Configuring Interface Sets for QoS on page 207

Chapter 19

Configuring an Integrated Scheduler to Provide QoS for ATM

This chapter provides information for configuring an integrated scheduler to provide QoS for ATM.

QoS topics are discussed in the following sections:

- ATM Integrated Scheduler Overview on page 159
- Integrating the HRR Scheduler and SAR Scheduler on page 161
- Per-Packet Queuing on the SAR Scheduler Overview on page 163
- Guidelines for Configuring QoS over ATM on page 166
- Configuring Default Integrated Mode for ATM Interface on page 168
- Configuring Low-Latency Mode for Per-Port Queuing on ATM Interfaces on page 169
- Configuring Low-CDV Mode for Per-Port Queuing on ATM Interfaces on page 171
- Configuring the QoS Shaping Mode for ATM Interfaces on page 174
- Disabling Per-Port Queuing on ATM Interfaces on page 175
- Monitoring QoS Configurations for ATM on page 175

ATM Integrated Scheduler Overview

The E Series Broadband Services Router provides extended ATM QoS functionality through its integrated scheduler. The integrated scheduler consists of two schedulers in series—the hierarchical round robin (HRR) scheduler and the segmentation and reassembly (SAR) scheduler.

The integrated scheduler enables you to configure QoS on your ATM networks using the HRR scheduler that is used on all E Series ASIC-enabled line modules. In addition, you can use the commercial SAR scheduler to configure traditional ATM cell-based QoS.



NOTE: The term *HRR scheduler* is used in this chapter to describe the scheduling performed by the ASIC on the ATM line module. Although the ASIC might differ depending on the ATM line module, the configuration and performance of the HRR scheduler are the same. For example, the ERX7xx models, ERX14xx models, and ERX310 Broadband Services router use the egress forwarding ASIC (EFA); and the E120 and E320 Broadband Services Routers use the frame forwarding ASIC (FFA) on the ES2 4G LM.

The HRR scheduler and the SAR scheduler work together as an integrated scheduler for ATM traffic. The HRR scheduler is configured by default with per-VC and per-IP interface scheduler nodes, and one best-effort class queue for each IP interface. The SAR scheduler implements weighted round-robin scheduling with one queue per VC. The VC queues are grouped into round robins based on the ATM service classes and the VP tunnels you have configured.

In the default integrated mode, controlled by the ATM application, the SAR scheduler controls the scheduling via the VC backpressure messages it sends to the HRR scheduler. When the HRR scheduler receives a backpressure message from the SAR scheduler, the HRR scheduler disables the node regardless of the node weight or shaping rate. When the HRR scheduler receives a backpressure release, the scheduler node is reenabled.

Backpressure and the Integrated Scheduler

ATM packets are initially scheduled through the HRR scheduler and then sent to the SAR scheduler, from where the cells are scheduled onto the circuit. If a SAR VC queue begins to fill up, the SAR scheduler issues *VC backpressure* messages to the HRR scheduler. The backpressure messages control the amount of traffic the HRR scheduler sends to the SAR scheduler. The SAR scheduler can also exert *port backpressure* on the HRR scheduler.

In default integrated mode, the SAR sends VC backpressure messages as well as port backpressure messages. Port backpressure messages are sent to the port node in the hierarchical scheduler.

Backpressure is a critical mechanism that enables the two schedulers in series to operate as a single integrated scheduler. Backpressure ensures that packets do not drain over internal data paths at an unmanageable rate from the HRR scheduler to the SAR scheduler. Without backpressure from the SAR scheduler, the HRR scheduler does not detect congestion even if the SAR scheduler is completely saturated.

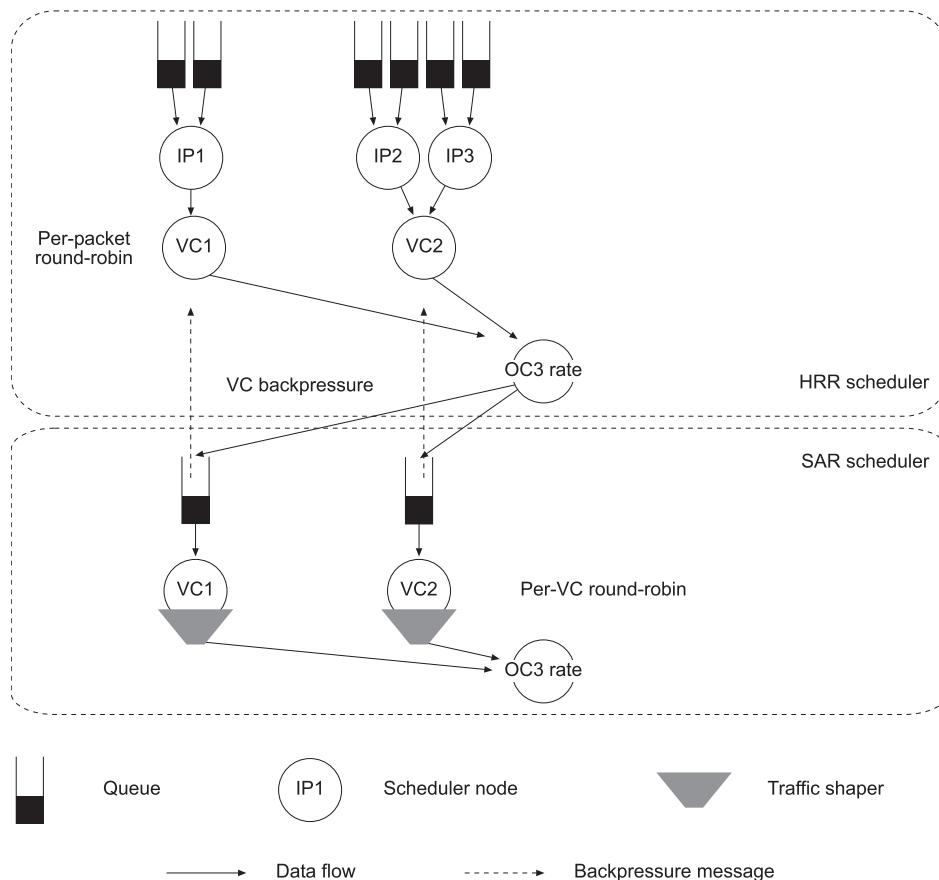


NOTE: The default QoS profile for ATM (atm-default) contains the *atm-vc node* command, which creates the scheduler node that is required by the SAR VC backpressure mechanism. If the SAR scheduler is operating in default integrated mode, this command must be in QoS profiles that are attached to ATM ports.

Figure 43 on page 161 shows the HRR and SAR schedulers working together to form the integrated scheduler. When the SAR VC queues start to back up, the SAR exerts VC backpressure to the corresponding VC node in the HRR scheduler.

VC backpressure affects only VC nodes that are in the default traffic-class group. As a consequence, VC nodes that are in named traffic-class groups within the scheduler hierarchy are not affected by VC backpressure.

Figure 43: Integrated ATM Scheduler



VP Shaping

VP shaping can be performed either in the SAR or by using the QoS shaping application using QoS profiles. Configuring VP shaping in the SAR enables traffic to be sent out of the port at a rate that closely matches strict ATM contract rates. SAR VP shaping is configured for the physical port using the **atm vp-tunnel** command.

- Related Topics**
- Integrating the HRR Scheduler and SAR Scheduler on page 161
 - Per-Packet Queuing on the SAR Scheduler Overview on page 163

Integrating the HRR Scheduler and SAR Scheduler

The proper integration of the two schedulers is an important element of the router's ATM QoS support. Three QoS port modes control integration of the two schedulers:

- Default integrated QoS port mode—ATM application controls the scheduling facilities of the SAR scheduler.
- Low-latency QoS port mode—HRR scheduler controls the traffic rate.
- Low-CDV QoS port mode—HRR scheduler and the SAR scheduler work together to schedule traffic.

Improper configuration of the two schedulers might create an inefficient scenario in which extra latency is introduced, or might cause the scheduler to underuse the link.

To configure integration of the schedulers, use the **qos-mode-port** commands listed in Table 16 on page 162.

Table 16: qos-mode-port Commands

Command	Backpressure	SAR Buffering	Scheduling
no qos-mode-port (default integrated mode)	VC and port	significant	SAR
qos-mode-port low-cdv	port	normal	SAR and HRR
qos-mode-port low-latency	port	minimal	HRR
qos-mode-port	port	minimal	HRR



NOTE: For ERX7xx models, ERX14xx models, and the ERX310 router, the *qos-mode-port* commands are valid only for the major interface on port 0.

To properly integrate the schedulers, make sure that the HRR and the SAR schedulers shape packets at the same rate. If the HRR scheduler sends packets at a higher rate than the SAR scheduler shapes them, the SAR scheduler can become congested and block the entire port.

To manage the integration of the HRR and the SAR schedulers:

1. Specify the cell-based shaping mode.
See “Configuring the QoS Shaping Mode for ATM Interfaces” on page 174.
2. Configure low-CDV QoS port mode to ensure that the HRR and SAR schedulers are configured at the same rate.
See “Configuring Low-CDV Mode for Per-Port Queuing on ATM Interfaces” on page 171.
3. Configure the QoS application to control the SAR scheduler’s operation. In this mode you configure both schedulers using scheduler profiles and QoS profiles. The E Series router then ensures that VPs and VCs are shaped to the same rates in both schedulers.



NOTE: You can also use the QoS cell mode application with QoS parameters to manage the integration of HRR and SAR schedulers.

Specifying the QoS cell mode application with the **qos-parameter-define** command enables you to configure a port with either frame or cell shaping mode and then configure the port for low-CDV port mode.

- Related Topics**
- For more information about scheduler profiles, see Scheduler Hierarchy Overview on page 45
 - For more information about configuring QoS profiles, see QoS Profile Overview on page 127
 - For more information about configuring QoS parameters, see QoS Parameter Overview on page 225

Per-Packet Queuing on the SAR Scheduler Overview

You can configure port queuing on the SAR scheduler, enabling per-packet rather than per-circuit scheduling. Port queuing mode allows you to use more of the facilities of the HRR scheduler, which are effectively disabled in default integrated mode, while at the same time making the SAR scheduler more transparent. In port queuing mode, you use the QoS application to configure the three levels of the HRR scheduler, including weighted round robin, traffic shaping, and strict priority scheduling.

You can configure the following modes:

- **Default integrated mode**—The ATM SAR scheduler does the scheduling. Both VC and port backpressure are enabled, and the HRR scheduler does minimal scheduling. The SAR scheduler performs significant buffering.
- **Low-latency mode**—The HRR scheduler does the scheduling. All QoS configurations are supported. VC backpressure is disabled, port backpressure is set as aggressive, and the SAR scheduler does minimal buffering. This mode enables the lowest latency for packets scheduled in the HRR scheduler with strict priority. Because the SAR scheduler is running with minimal buffering, there is no head-of-line blocking.
- **Low-CDV mode**—The HRR and SAR schedulers both perform scheduling; QoS synchronizes the rates of the two schedulers. All QoS configurations are supported. VC backpressure is disabled, and port backpressure is set to the default thresholds of 6 MB per OC3 port and 24 MB per OC12 port. In this mode, you can configure shaping in both the SAR scheduler and the HRR scheduler; low-cdv mode works with cell shaping mode only and enables relative weighted VCs and hierarchical shaping in the HRR scheduler. The SAR scheduler performs normal buffering and can shape either the VC or VP, but not both.

Operational QoS Shaping Mode for ATM Interfaces Overview

The E Series router enables you to shape ATM traffic based on either frames or cells. The default frame shaping mode provides compatibility with previous versions of the E Series software. 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 ATM network.

ATM policing is sensitive to cell delay variation tolerance (CDVT). If the cells on a particular VC or VP arrive too closely spaced, an ATM switch might drop cells. However, the cell scheduler reduces CDVT by ensuring cell spacing. The router enables you to use techniques such as WRR on the HRR scheduler to achieve the proper packet scheduling. You use the SAR scheduler in series with the HRR scheduler to even out cell bursts into smoother per-VC and per-VP traffic profiles that bound CDVT. You accomplish this by using the **qos-shaping-mode cell** command to configure the QoS shaping mode, and the **qos-mode-port low-cdv** command to configure the port queuing mode.

The QoS shaping mode also determines how QoS statistics are reported. Frame shaping reports QoS statistics such as transmitted bytes and dropped bytes based on bytes within frames. Cell shaping reports the statistics in bytes within cells and also accounts for cell encapsulation and padding overhead.

ERX7xx Models, ERX14xx Models, and the ERX310 Router

The ERX7xx models, ERX14xx models, and the ERX310 router use an operational shaping mode that is based on the following two commands:

- The QoS shaping mode you set with the **qos-shaping-mode** command on port 0 and on the specific port
- The port queuing mode you set with the **qos-mode-port** command on port 0

The router uses the following rules to determine the operational shaping mode used for a port:

1. If the specific port has a QoS shaping mode configured, the operational shaping mode for that port is the same as the QoS shaping mode.
2. If the specific port has no QoS shaping mode configured, the operational shaping mode is the same as the QoS shaping mode for port 0, if one is configured.
3. If both the specific port and port 0 have no QoS shaping mode configured, the operational shaping mode is based on the port 0 queuing mode. If the port 0 queuing mode (set by the **qos-mode-port** command) is low-cdv, the operational shaping mode is cell; otherwise the operational shaping mode is frame.

Table 17 on page 165 lists the possible combinations of the two commands and the resultant operational shaping mode.

Table 17: Operational Shaping Modes for ERX7xx Models, ERX14xx Models, and the ERX310 Router

Rule	qos-shaping-mode for the Specific Port	qos-shaping-mode for Port 0	qos-mode-port for Port 0	Operational Shaping Mode for the Specific Port
Rule 1	Cell	Cell	low-cdv	Cell
	Frame	Frame	low-latency or none	Frame
Rule 2	No shaping mode	Cell	low-cdv	Cell
	No shaping mode	Frame	low-latency or none	Frame
Rule 3	No shaping mode	No shaping mode	low-cdv	Cell
	No shaping mode	No shaping mode	low-latency or none	Frame

E120 Router and E320 Router

The E120 router and the E320 router use an operational shaping mode that is based on the following two commands:

- The QoS shaping mode you set with the **qos-shaping-mode** command on port 0 and on the specific port
- The port queuing mode you set with the **qos-mode-port** command on port 0 and on the specific port

The E120 and E320 routers use the following rules to determine the operational shaping mode used for a port:

1. If the specific port has a QoS shaping mode configured, the operational shaping mode for that port is the same as the QoS shaping mode.
2. If the specific port has no QoS shaping mode and a port queuing mode of low-cdv configured, the operational shaping mode is cell.
3. If the specific port has no QoS shaping mode and no queuing mode configured, the operational shaping mode for that port is the same as the port 0 QoS shaping mode.
4. If both the specific port and port 0 have no QoS shaping mode configured, the operational shaping mode is based on the port 0 queuing mode. If the port 0 queuing mode (set by the **qos-mode-port** command) is low-cdv, the operational shaping mode is cell; otherwise the operational shaping mode is frame.

Table 18 on page 166 lists the possible combinations of the two commands and the resultant operational shaping mode.

Table 18: Operational Shaping Modes for the E120 Router and E320 Router

Rule	qos-shaping-mode for specific port	qos-mode-port for Specific Port	qos-shaping-mode for Port 0	qos-mode-port for Port 0	Operational Shaping Mode for Specific Port
Rule 1	Cell	low-cdv	Any	Any	Cell
	Frame	low-latency or none	Any	Any	Frame
Rule 2	No shaping mode	low-cdv	Any	Any	Cell
Rule 3	No shaping mode	low-latency or none	Frame	Any	Frame
	No shaping mode	low-latency or none	Cell	Any	Cell
Rule 4	No shaping mode	low-latency or none	No shaping mode	low-cdv	Cell
	No shaping mode	low-latency or none	No shaping mode	low-latency or none	Frame

- Related Topics**
- Guidelines for Configuring QoS over ATM on page 166
 - Configuring Default Integrated Mode for ATM Interface on page 168
 - Configuring Low-Latency Mode for Per-Port Queuing on ATM Interfaces on page 169
 - Configuring Low-CDV Mode for Per-Port Queuing on ATM Interfaces on page 171
 - Configuring the QoS Shaping Mode for ATM Interfaces on page 174

Guidelines for Configuring QoS over ATM

This section provides general QoS configuration guidelines for ATM line modules. These guidelines are applicable to all JUNOS releases.

The SAR scheduler generates VC backpressure as a way to control the flow of packets from the HRR scheduler to the SAR scheduler. The QoS port modes control integration of the two schedulers.

In default integrated mode, each VC queue in the SAR generates backpressure for the ATM VC node in the default traffic class group in the HRR. The backpressure throttles the dequeue rate of the ATM VC node and the nodes and queues stacked above it in the scheduler hierarchy. VC backpressure is disabled in low-latency QoS port mode and low-cdv QoS port mode.

You can configure queues in default integrated mode in the HRR that are immune to VC backpressure so that you can run voice and video applications. Queues and nodes in any named traffic class group are not subject to VC backpressure.

In addition, ATM VP and ATM (port level) queues are not stacked above ATM VC nodes, so queues are not subject to backpressure, regardless of the traffic class group.

Take care not to saturate SAR queues with too much traffic from the HRR, especially when shaping VP tunnels or VCs in the SAR. You can accomplish this in several ways:



NOTE: These rules apply only to the default integrated mode. VC backpressure is disabled in low-latency or low-cdv modes. You must account for cell tax; to do this, use the **qos-shaping-mode cell** command for the line module.

- Use external admission control to guarantee that the sum of non-backpressured traffic into the VC is less than the SAR shaping rate for the VC.
- Shape the non-backpressured queues or nodes in the HRR, making the aggregate of the non-backpressured traffic for a VC less than the VC rate.
- In JUNOS Release 6.0 and later, you can configure a shared shaper on the ATM VC node in the default traffic class group. Configure the shared-shaping rate to be less than or equal to the VC shaping rate in the SAR.
- Special rules apply for VP tunnels shaped in the SAR. When shaping in the SAR, configure ATM VP nodes in the HRR, and arrange that the aggregate traffic dequeued from the HRR for that vp-tunnel is less than or equal to the VP tunnel shaping rate in the SAR.

Use one of the following two techniques for VP tunnels shaped in the SAR:

- Partition the SAR VP tunnel bandwidth across the ATM VP nodes in the different traffic class groups in the HRR. For example, using a 4 Mbps VP tunnel, allocate 1 Mbps for the ATM VP node in the default traffic class group, 2 Mbps for the ATM VP node in the video traffic class group, and 1 Mbps for the ATM VP node in the voice traffic class group.

When using this technique, keep in mind that the different traffic classes cannot share bandwidth.

- In JUNOS Release 6.1 and later, using the EFA2 ASIC, you can configure shared shaping on the ATM VP nodes in the HRR to perform bandwidth sharing.

Related Topics

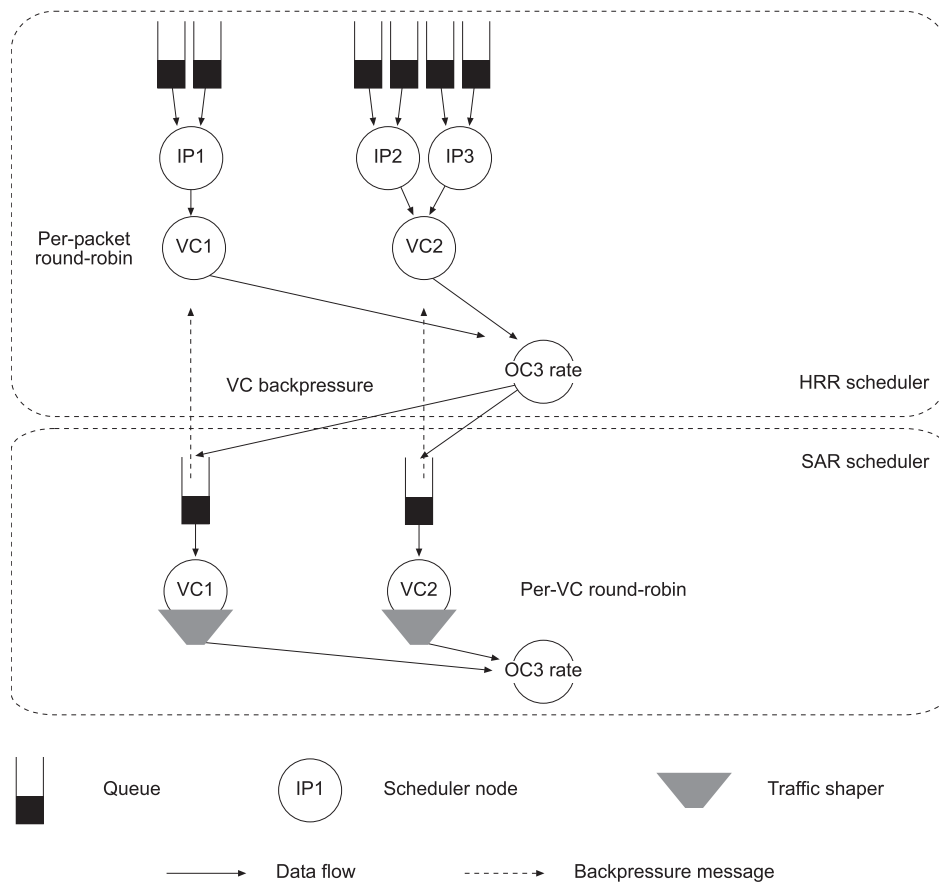
- Integrating the HRR Scheduler and SAR Scheduler on page 161
- Shared Shaping and Low-CDV Mode on page 75
- Configuring Default Integrated Mode for ATM Interface on page 168
- Configuring Low-Latency Mode for Per-Port Queuing on ATM Interfaces on page 169
- Configuring Low-CDV Mode for Per-Port Queuing on ATM Interfaces on page 171
- Configuring the QoS Shaping Mode for ATM Interfaces on page 174

Configuring Default Integrated Mode for ATM Interface

In the default integrated mode, the SAR scheduler is the dominant scheduler, and it backpressures the first-stage (HRR) scheduler per VC. Each VC buffers only a few hundred bytes.

Figure 44 on page 168 shows the default integrated mode.

Figure 44: Default Integrated Mode



To configure default integrated mode:

1. From the desired port, set the QoS port mode to default integrated mode.

```
host1(config)#interface atm 2/0
host1(config-if)#no qos-mode-port
```



TIP: For ATM interfaces on ERX7xx models, ERX14xx models, and the ERX310 router, you must specify port 0.

2. Specify the VP shaping rate.

```
host1(config-if)#atm vp-tunnel 0 2000
```



TIP: Configuring an ATM VP tunnel sets a shaping rate in the SAR scheduler. Before configuring an ATM VP tunnel, there must be no PVCs with the same VPI that you are about to configure. Before using the **atm vp-tunnel** command, remove any PVCs from the configuration. You can reconfigure the PVCs after configuring the shapeless VP tunnel.

3. Specify the shaping rate for the ATM subinterface.

```
host1(config-if)#interface atm 2/0.5  
host1(config-subif)#atm-pvc 5 0 5 aal5snap 768
```

- Related Topics**
- Per-Packet Queuing on the SAR Scheduler Overview on page 163
 - Guidelines for Configuring QoS over ATM on page 166
 - atm vp-tunnel
 - interface atm
 - qos-mode-port

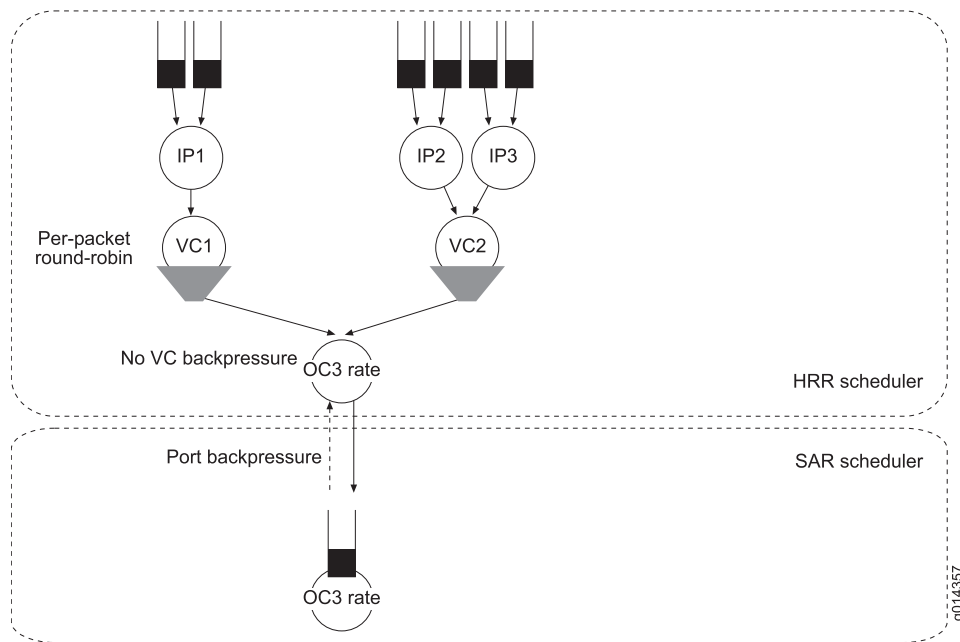
Configuring Low-Latency Mode for Per-Port Queuing on ATM Interfaces

In low-latency mode, the SAR scheduler backpressures the HRR scheduler per physical port; each physical port buffers only a few kilobytes.

When you configure low-latency mode:

- VC backpressure is disabled.
- Port backpressure is enabled as aggressive.
- SAR scheduler performs minimal buffering.
- HRR scheduler is dominant.

This procedure creates the low-latency mode configuration shown in Figure 45 on page 170.

Figure 45: Low-Latency Mode

To configure low-latency mode with a strict-priority queue and a best-effort queue:

1. Configure the traffic class.

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

2. Set the traffic class in the traffic-class group.

```
host1(config)#traffic-class-group strict
host1(config-traffic-class-group)#traffic-class strict
host1(config-traffic-class-group)#exit
```

3. Define the scheduler profile for the traffic-class group.

```
host1(config)#scheduler-profile strict
host1(config-scheduler-profile)#strict-priority
host1(config-scheduler-profile)#exit
```

4. Configure the QoS profile with two ATM VC queues.

```
host1(config)#qos-profile low-latency-q-p
host1(config-qos-profile)#atm-vc node
host1(config-qos-profile)#atm-vc queue traffic-class best-effort
host1(config-qos-profile)#atm group strict scheduler-profile strict
host1(config-qos-profile)#atm-vc queue traffic-class strict
host1(config-qos-profile)#exit
```

5. From the desired port, set the QoS port mode to low latency.

```
host1(config)#interface atm 2/0
host1(config-if)#qos-mode-port low-latency
```

host1(config-if)#**qos-profile low-latency-q-p**



TIP: For ATM interfaces on ERX7xx models, ERX14xx models, and the ERX310 router, you must specify port 0.

The **qos-mode-port** command:

- Excludes non-UBR ATM QoS services on any VC on the ATM module; for example, PCR, nrtVBR, and CBR
- Cannot be used if shaping is currently configured on the SAR scheduler
- Cannot be used with ATM VP tunnels with nonzero rates; however, can be used with tunnels with rates of zero (shapeless tunnels).

Related Topics

- Per-Packet Queuing on the SAR Scheduler Overview on page 163
- Guidelines for Configuring QoS over ATM on page 166
- interface atm
- qos-mode-port
- qos-profile



Configuring Low-CDV Mode for Per-Port Queuing on ATM Interfaces

In low-CDV mode, the HRR scheduler and the SAR scheduler operate in concert. In low-CDV QoS port mode, QoS automatically configures the shaping rate of the VPs, VCs, or both based on the QoS profile and the scheduler profile. Therefore, the QoS shaping mode must be set to the cell mode. In low-CDV mode, the SAR scheduler converts frame-atomic bursts of cells to CDVT-conformant streams of interleaved cells. There is no VC backpressure, and the port backpressure is loose, so several megabytes of cells can reside in the SAR buffer pool.

When you configure low-CDV mode:

- QoS synchronizes the shaping rates for VPs and VCs in the HRR and SAR schedulers.
- VC backpressure is disabled.
- Port backpressure is set to default thresholds of 6 MB per OC3 port and 24 MB per OC12 port.
- SAR scheduler performs more buffering than in low-latency mode.
- Use cell QoS shaping mode.

This procedure creates the low-CDV mode with per-VP CDVT configuration shown in Figure 46 on page 172. Figure 47 on page 172 shows low-CDV mode with per-VC CDVT.

Figure 46: Low-CDV Mode (per-VP CDVT)**Figure 47: Low-CDV Mode (per-VC CDVT)**

To configure low-CDV mode with a strict-priority queue and a best-effort queue:

1. Configure the traffic class.

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

2. Set the traffic class in the traffic-class group.

```
host1(config)#traffic-class-group strict
host1(config-traffic-class-group)#traffic-class strict
host1(config-traffic-class-group)#exit
```

3. Define the scheduler profiles for the traffic-class group.

```
host1(config)#scheduler-profile strict
host1(config-scheduler-profile)#strict-priority
host1(config-scheduler-profile)#exit
host1(config)#scheduler-profile 500k
host1(config-scheduler-profile)#shaping-rate 500000
host1(config-scheduler-profile)#exit
host1(config)#scheduler-profile 1m
host1(config-scheduler-profile)#shaping-rate 1000000
host1(config-scheduler-profile)#exit
host1(config)#scheduler-profile 2m
host1(config-scheduler-profile)#shaping-rate 2000000
host1(config-scheduler-profile)#exit
```

4. Configure per-VC CDVT by configuring QoS profile with ATM VC queues.

```
host1(config)#qos-profile low-cdv-q-p
host1(config-qos-profile)#atm-vc node scheduler-profile 1m
host1(config-qos-profile)#atm-vp node scheduler-profile 2m
host1(config-qos-profile)#atm-vc queue traffic-class best-effort
host1(config-qos-profile)#atm group strict scheduler-profile strict
host1(config-qos-profile)#atm-vc queue traffic-class strict scheduler-profile 500k
host1(config-qos-profile)#exit
```

5. Configure per-VP CDVT using shapeless VP tunnels that are used when the QoS application controls SAR scheduler shaping and set the QoS port mode to low CDV.

```
host1(config)#interface atm 2/0
host1(config-if)#atm vp-tunnel 0 0
host1(config-if)#atm vp-tunnel 1 0
host1(config-if)#qos-mode-port low-cdv
host1(config-if)#qos-profile low-cdv-q-p
host1(config-subif)#interface atm 2/0.5
host1(config-subif)#atm pvc 5 0 5 aal5snap
host1(config-subif)#interface atm 2/0.6
host1(config-subif)#atm pvc 6 0 6 aal5snap
host1(config-subif)#interface atm 2/0.7
host1(config-subif)#atm pvc 7 1 7 aal5snap
host1(config-subif)#interface atm 2/0.8
host1(config-subif)#atm pvc 8 1 8 aal5snap
```



TIP: For ATM interfaces on ERX7xx models, ERX14xx models, and the ERX310 router, you must specify port 0.

Configuring an ATM VP tunnel sets a shaping rate in the SAR scheduler. Before configuring an ATM VP tunnel, there must be no PVCs with the same VPI that you are about to configure. Before using the **atm vp-tunnel** command, remove any PVCs from the configuration. You can reconfigure the PVCs after configuring the shapeless VP tunnel.

The **qos-mode-port** command:

- Excludes non-UBR ATM QoS services on any VC on the ATM module; for example, PCR, nrtVBR, and CBR
- Cannot be used if shaping is currently configured on the SAR scheduler
- Cannot be used with ATM VP tunnels with nonzero rates; however, can be used with tunnels with rates of zero (shapeless tunnels)

- Related Topics**
- Per-Packet Queuing on the SAR Scheduler Overview on page 163
 - Guidelines for Configuring QoS over ATM on page 166
 - atm vp-tunnel
 - interface atm
 - qos-mode-port

Configuring the QoS Shaping Mode for ATM Interfaces

In frame mode, SAR shaping is controlled by the ATM application. Shaping is based on the number of bytes in the frame, without regard to cell encapsulation or padding overhead; this is the default mode.

In cell mode, SAR shaping is controlled by the QoS application. Shaping is based on the number of bytes in cells, and accounts for the ATM cell encapsulation and padding overhead.

To configure the operational shaping mode for ATM interfaces:

1. Configure the ATM interface.

```
host1(config)#interface atm 5/1
```

For ATM interfaces on ERX7xx models, ERX14xx models, and the ERX310 router, you must use port 0.

2. Configure the shaping mode and specify either frame or cell.

```
host1(config-if)#qos-shaping-mode cell
```



BEST PRACTICE: We recommend that you clear the statistics counters whenever you change the QoS shaping mode. Otherwise, the statistics contain a mixture of frame-based and cell-based values.

- Related Topics**
- Per-Packet Queuing on the SAR Scheduler Overview on page 163
 - interface atm
 - qos-mode-port
 - qos-shaping-mode

Disabling Per-Port Queuing on ATM Interfaces

You can remove per-port queuing on ATM interfaces and restore the default integrated mode setting.

When per-port queuing is disabled, both the VC and port backpressure are enabled. The SAR scheduler performs significant buffering, and the HRR scheduler does minimal scheduling. The **atm-vc node** command must appear in the QoS profile attached to the ATM port.

To disable per-port queuing:

1. Specify the ATM interface for which you want to disable per-port queuing.

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

2. Disable per-port queuing on that interface.

```
host1(config-if)#no qos-mode-port
```

- Related Topics**
- Configuring Default Integrated Mode for ATM Interface on page 168
 - interface atm
 - qos-mode-port

Monitoring QoS Configurations for ATM

To monitor QoS configurations for ATM:

- Monitoring the QoS Configuration of ATM Interfaces on page 343
- Monitoring the QoS Configuration of IP Interfaces on page 345
- Monitoring the QoS Profiles Attached to an Interface on page 339
- Monitoring the Configuration of QoS Port-Type Profiles on page 340

- Monitoring the Configuration of QoS Profiles on page 341
- Monitoring the QoS Scheduler Hierarchy on page 322
- Monitoring Shared Shapers on page 329

Chapter 20

Configuring QoS for Gigabit Ethernet Interfaces and VLAN Subinterfaces

This chapter provides information for configuring QoS for Gigabit Ethernet interfaces and VLAN subinterfaces.

QoS topics are discussed in the following sections:

- Providing QoS for Ethernet Overview on page 177
- QoS Shaping Mode for Ethernet Interfaces Overview on page 178
- Configuring the QoS Shaping Mode for Ethernet Interfaces on page 179
- Creating a QoS Interface Hierarchy for Bulk-Configured VLAN Subinterfaces with RADIUS on page 180
- Monitoring QoS Configurations for Ethernet on page 182

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 Topics**
- QoS Shaping Mode for Ethernet Interfaces Overview on page 178
 - Creating a QoS Interface Hierarchy for Bulk-Configured VLAN Subinterfaces with RADIUS on page 180
 - To configure subscriber-based QoS for 802.3ad link aggregation interfaces, see QoS for 802.3ad Link Aggregation Interfaces Overview on page 183

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 19 on page 178 lists the possible combinations of the **qos-shaping-mode** command and the resultant operational shaping mode.

Table 19: 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 Topics**
- Configuring the QoS Shaping Mode for Ethernet Interfaces on page 179
 - Byte Adjustment for ADSL and VDSL Traffic Overview on page 293
 - Cell Shaping Mode Using QoS Parameters Overview on page 281

Configuring the QoS Shaping Mode for Ethernet Interfaces

You can configure the shaping mode for an Ethernet interface.

In frame mode, traffic shaping is controlled by the system. Shaping is based on the number of bytes in the frame, without regard to cell encapsulation or padding overhead; this is the default mode for all E Series routers.

In cell mode, shaping is controlled by the QoS application. Shaping is based on the number of bytes in cells, and accounts for the ATM cell encapsulation and padding overhead. This option is available only for Gigabit Ethernet interfaces configured on the GE-2 line module, the GE-HDE line module, and the ES2 4G LM, and 10-Gigabit Ethernet interfaces configured on the ES2 4G LM.

1. Configure the Ethernet interface.

```
host1(config)#interface gigabitEthernet 6/0/0
```

2. Configure the shaping mode and specify frame or cell.

```
host1(config)#qos-shaping-mode cell
```



BEST PRACTICE: We recommend that you clear the statistics counters whenever you change the QoS shaping mode. Otherwise, the statistics contain a mixture of frame-based and cell-based values.

- Related Topics**
- QoS Shaping Mode for Ethernet Interfaces Overview on page 178
 - interface gigabitEthernet
 - qos-shaping-mode

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: erx GigabitEthernet
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
```

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

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

Interface	IP-Address	Status	Protocol	Description
-----	-----	-----	-----	-----
null0	255.255.255.255/32	up	up	
loopback0	10.1.0.1/24	up	up	
GigabitEthernet6/0.100	Unnumbered	up	up	

- 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.

- 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 Topics**
- For information about bulk-configured VLAN subinterfaces, see *JUNOS Link Layer Configuration Guide*
 - For information about service profiles, see *JUNOS Broadband Access Configuration Guide*
 - For information about RADIUS VSAs, see *JUNOS Broadband Access Configuration Guide*
 - 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

Monitoring QoS Configurations for Ethernet

To monitor Ethernet configurations for QoS:

- Monitoring the QoS Configuration of Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet Interfaces on page 347
- Monitoring the QoS Configuration of IP Interfaces on page 345
- Monitoring the QoS Profiles Attached to an Interface on page 339
- Monitoring the Configuration of QoS Port-Type Profiles on page 340
- Monitoring the Configuration of QoS Profiles on page 341
- Monitoring the QoS Scheduler Hierarchy on page 322
- Monitoring Shared Shapers on page 329

Chapter 21

Configuring QoS for 802.3ad Link Aggregation Groups

This chapter provides information for configuring QoS for 802.3ad link aggregation groups.

QoS topics are discussed in the following sections:

- QoS for 802.3ad Link Aggregation Interfaces Overview on page 183
- Hashed Load Balancing for 802.3ad Link Aggregation Groups Overview on page 186
- Subscriber Load Balancing for 802.3ad Link Aggregation Groups Overview on page 186
- Guidelines for Configuring QoS over 802.3ad Link Aggregation Groups on page 190
- Configuring the Scheduler Hierarchy for Hashed Load Balancing in 802.3ad Link Aggregation Groups on page 191
- Enabling Default Subscriber Load Balancing for 802.3ad Link Aggregation Groups on page 191
- Configuring the Scheduler Hierarchy for Subscriber Load Balancing in 802.3ad Link Aggregation Groups on page 192
- Configuring Load Rebalancing for 802.3ad Link Aggregation Groups on page 193
- Monitoring QoS Configurations for 802.3ad Link Aggregation Groups on page 195

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 Topics**
- Hashed Load Balancing for 802.3ad Link Aggregation Groups Overview on page 186
 - Subscriber Load Balancing for 802.3ad Link Aggregation Groups Overview on page 186
 - For more information about configuring the lag-default QoS profile for default subscriber load balancing, see *Enabling Default Subscriber Load Balancing for 802.3ad Link Aggregation Groups* on page 191
 - For more information about Ethernet link redundancy, see *JUNOS Link Layer Configuration Guide*
 - For more information about configuring QoS parameters, see *Parameter Definition Attributes for QoS Administrators Overview* on page 229

- For more information about the munge algorithm, see Munged QoS Profile Overview on page 136
- For a list of modules that support 802.3ad link aggregation, see the *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.

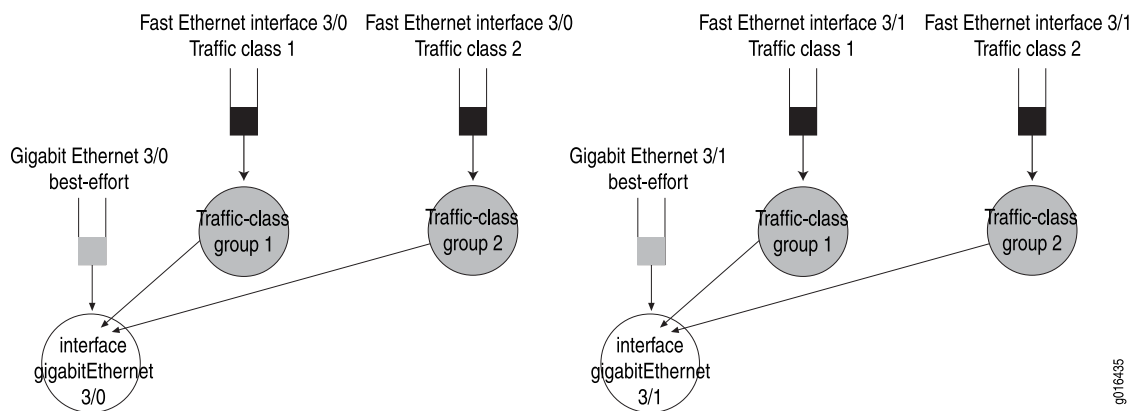
Because all traffic is carried in Ethernet queues, per-subscriber QoS features such as shared shaping for VLANs are not available.

Sample Scheduler Hierarchy for Hashed Load Balancing

Figure 48 on page 186 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 48: 802.3ad Link Aggregation Scheduler Hierarchy



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 49 on page 188 displays the scheduler hierarchy for the Gigabit Ethernet interface in slot 3, port 0. Figure 50 on page 189 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 49: Subscriber LoadBalanced Scheduler Hierarchy for Port 0

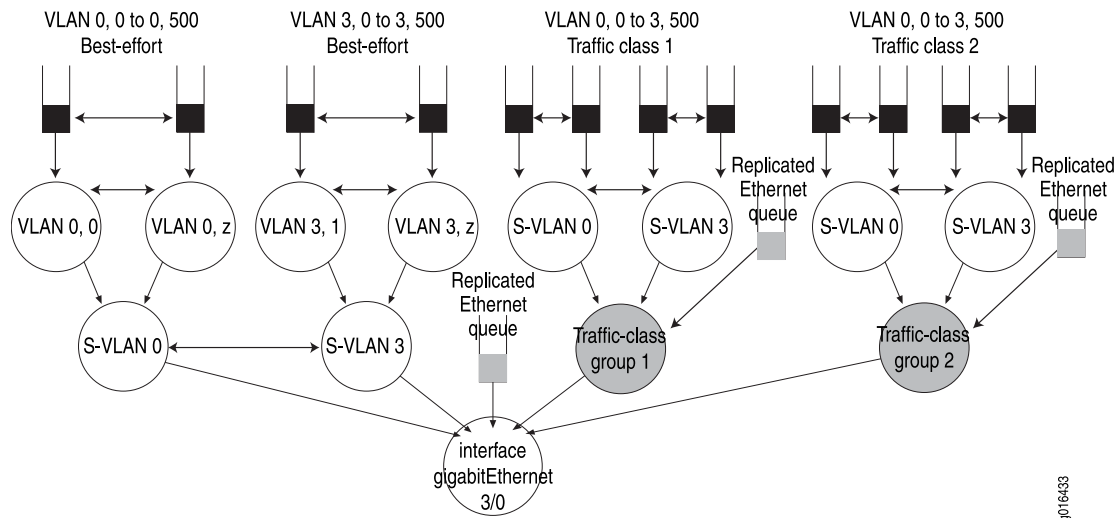
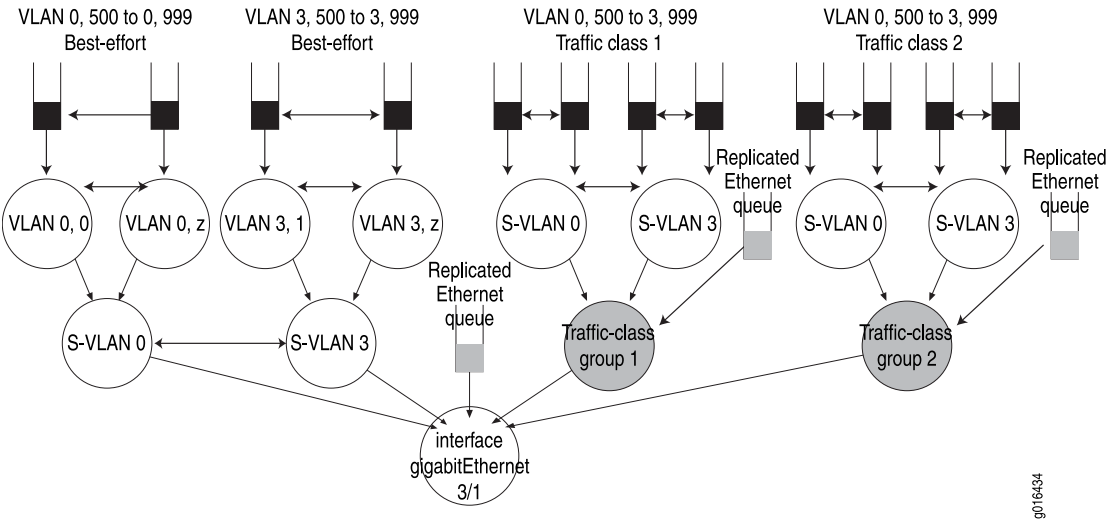


Figure 50: Subscriber LoadBalanced Scheduler Hierarchy for Port 1



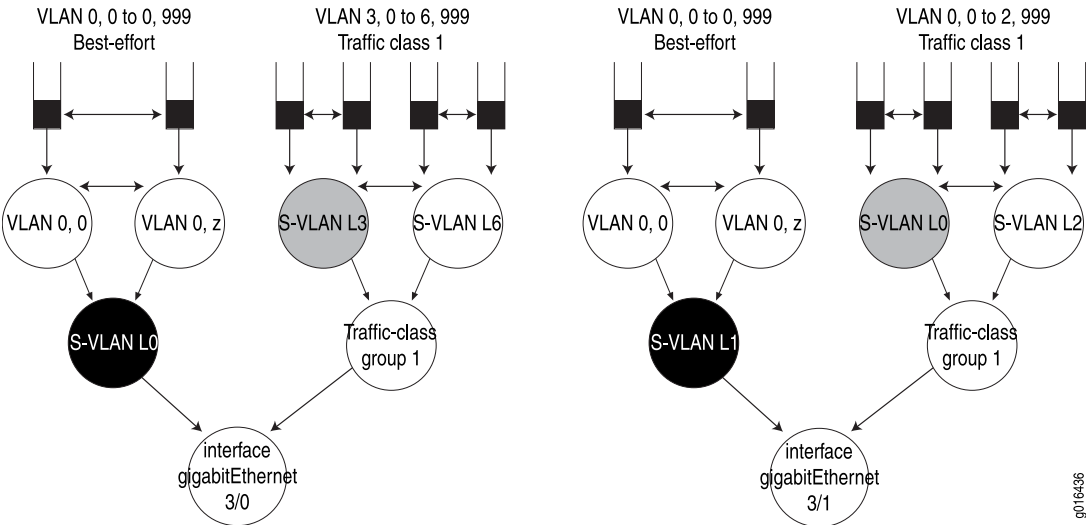
Subscriber Allocation in 802.3ad Link Aggregation Groups

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 51 on page 189 displays a sample allocation of subscribers.

Figure 51: 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 Topics** ■ For more information about load rebalancing, see *Configuring Load Rebalancing for 802.3ad Link Aggregation Groups* on page 193

Guidelines for Configuring QoS over 802.3ad Link Aggregation Groups

When you configure QoS over 802.3ad LAGs, be sure to consider the following behaviors:

- QoS profiles cannot be attached to Ethernet ports if the port is a member of a LAG. In typical QoS configurations, the Ethernet interface is considered the root of the interface hierarchy. When you configure QoS for 802.3ad link aggregation, the LAG interface is considered the root of the interface hierarchy.
- You cannot configure hierarchical QoS for IP configured directly over a LAG interface.
- You cannot obtain QoS information or statistics for IP interfaces stacked over a LAG interface using any of the **show** commands for QoS. Instead, the **show qos scheduler-hierarchy** command is designed to find the interface hierarchy rooted at the specified interface and report all scheduler nodes and queues managed by those interfaces. The typical defaults in QoS profiles such as `ethernet-default` and `atm-default` specify the "ip queue traffic-class best-effort" rule, so those queues are reported in the interface hierarchy. The `lag-default` QoS profile does not specify this rule by default.
- Do not attach QoS profiles to IP or VLAN subinterfaces in a LAG that contain downreferences (that is, rules for S-VLAN or Ethernet nodes or queues). QoS profiles attached at subinterfaces above a LAG that also include downreference create an asymmetric scheduler hierarchy. For example, one Ethernet port might be shaped and not another.

Also, if the QoS profile specifies only Ethernet, then the traffic sent to the subinterface might be only partially affected by the QoS profile, or not at all. The traffic can be allocated to another port entirely.

- Related Topics** ■ *Subscriber Load Balancing for 802.3ad Link Aggregation Groups Overview* on page 186

- Configuring the Scheduler Hierarchy for Subscriber Load Balancing in 802.3ad Link Aggregation Groups on page 192
- Configuring Load Rebalancing for 802.3ad Link Aggregation Groups on page 193

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 Topics

- QoS for 802.3ad Link Aggregation Interfaces Overview on page 183
- Hashed Load Balancing for 802.3ad Link Aggregation Groups Overview on page 186
- interface lag
- node
- qos-profile
- queue

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 Topics**
- QoS for 802.3ad Link Aggregation Interfaces Overview on page 183
 - qos-port-type-profile

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 Topics**
- QoS for 802.3ad Link Aggregation Interfaces Overview on page 183
 - Subscriber Load Balancing for 802.3ad Link Aggregation Groups Overview on page 186
 - Enabling Default Subscriber Load Balancing for 802.3ad Link Aggregation Groups on page 191
 - interface lag
 - member-interface
 - node
 - qos-profile
 - queue

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 193
- Configuring the System to Dynamically Rebalance the LAG on page 195

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 20 on page 193 describes the load balancing algorithm parameters that you can configure.

Table 20: 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.

Table 20: Load Balancing Algorithm Parameters *(continued)*

Keyword	Description
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.
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 Topics**
- Configuring Load Rebalancing for 802.3ad Link Aggregation Groups on page 193
 - QoS for 802.3ad Link Aggregation Interfaces Overview on page 183
 - Enabling Default Subscriber Load Balancing for 802.3ad Link Aggregation Groups on page 191
 - Configuring the Scheduler Hierarchy for Subscriber Load Balancing in 802.3ad Link Aggregation Groups on page 192
 - interface lag
 - load-rebalance

Monitoring QoS Configurations for 802.3ad Link Aggregation Groups

To monitor Ethernet configurations for QoS:

- Monitoring the QoS Configuration of Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet Interfaces on page 347
- Monitoring the QoS Configuration of IEEE 802.3ad Link Aggregation Group Bundles on page 349
- Monitoring the QoS Configuration of IP Interfaces on page 345
- Monitoring the QoS Profiles Attached to an Interface on page 339
- Monitoring the Configuration of QoS Port-Type Profiles on page 340
- Monitoring the Configuration of QoS Profiles on page 341
- Monitoring the QoS Scheduler Hierarchy on page 322
- Monitoring Shared Shapers on page 329

Chapter 22

Configuring QoS for L2TP Sessions

This chapter provides information for configuring QoS for L2TP sessions.

QoS topics are discussed in the following sections:

- Providing QoS for L2TP Overview on page 197
- Sample Scheduler Hierarchies for L2TP on page 197
- Configuring QoS for an L2TP Session on page 199
- Configuring QoS for Tunnel-Server Ports for L2TP LNS Sessions on page 202
- QoS and L2TP TX Speed AVP 24 Overview on page 203
- Monitoring QoS Configurations for L2TP on page 204

Providing QoS for L2TP Overview

The JUNOS software supports QoS queues and scheduler nodes for L2TP session interfaces. L2TP QoS provides per-L2TP session queuing and allows QoS profiles to be dynamically attached to L2TP session interfaces on E Series Broadband Services Routers. The routers can be configured as either an LAC or LNS.

The dynamic attachment process uses RADIUS and AAA, enabling a QoS profile to be attached to a dynamic L2TP session interface when the newly created interface has the QoS-Profile-Name [26-26] RADIUS VSA associated with it. L2TP QoS support gives you the ability to shape tunneled users through L2TP interfaces.

L2TP QoS profiles are attached at the L2TP session interface, except on the LNS with nonmultilink interfaces. On the LNS with nonmultilink interfaces, L2TP QoS profiles are attached at the IP interface. The queues and scheduler node are built at the L2TP client interface on the line module.

Sample Scheduler Hierarchies for L2TP

The figures in this section show the different scheduler hierarchies that you can build for QoS over L2TP. The type of networking architecture in which the QoS profile is used determines the actual hierarchy that is built.

Figure 52 on page 198 through Figure 56 on page 199 show scheduler hierarchies for different networking architectures.




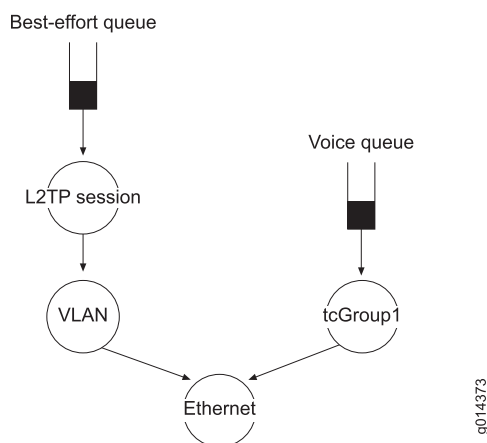
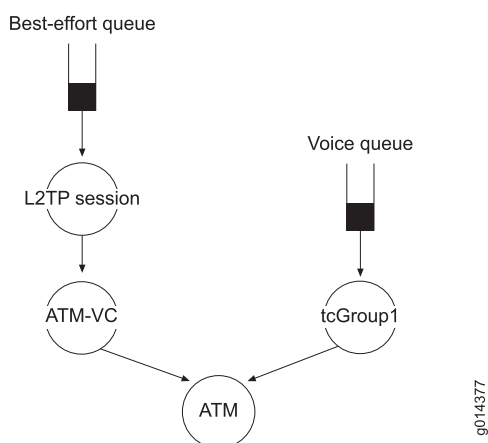
Figure 52: LNS (Non-MLPPP) Scheduler Hierarchy**Figure 53: LNS (MLPPP) QoS Scheduler Hierarchy****Figure 54: LAC over Ethernet (Without VLANs) Scheduler Hierarchy**

Figure 55: LAC over Ethernet (With LANs) Scheduler Hierarchy**Figure 56: LAC over ATM**

Related Topics ■ Configuring QoS for an L2TP Session

Configuring QoS for an L2TP Session

This section provides general procedures for configure QoS for an L2TP LNS session or a LAC L2TP session.

For both procedures, the resulting scheduler hierarchy depends on the type of network architecture that you use.

L2TP session interfaces have default QoS profiles and scheduler nodes. The default configuration includes the following settings:

```
host1(config)# show qos-profile l2tp-session-default
t-class  interface  rule  traffic  scheduler  queue  drop  statistics
group    type          type  class    profile    profile profile profile
```

```
l2tp-session queue best-effort default default default default
```

1. Configuring QoS for an L2TP LNS Session on page 200
2. Configuring QoS for an L2TP LAC Session on page 201

Configuring QoS for an L2TP LNS Session

When you configure QoS for an LNS, you must modify the server-default QoS profile to remove the **best-effort** traffic class rule from the IP interface type. This enables you to create L2TP session queues, and is not required to provide QoS on an LAC.

Before you configure QoS for an L2TP LNS session:

- Configure the traffic classes.
See “Configuring Traffic Classes That Define Service Levels” on page 14.
- Configure the queuing hierarchy.
See “Configuring Queue Profiles to Manage Buffers and Thresholds” on page 22.
- Configure the scheduler hierarchy and shaping with scheduler profiles.
See “Configuring a Scheduler Hierarchy” on page 47.

To configure QoS for an L2TP LNS session:

1. Remove the **best-effort** traffic class rule from the IP interface type of the server-default QoS profile.

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

2. Create a traffic-class group, and enter Traffic Class Group Configuration mode. Add the traffic class voice to the new group.

```
host1(config)#traffic-class-group tcGroup1
host1(config-traffic-class-group)#traffic-class voice
host1(config-traffic-class-group)#exit
```

3. Configure the QoS profile.

- Create the QoS profile, and enter QoS Profile Configuration mode.

```
host1(config)#qos-profile l2tpQpro25
host1(config-qos-profile)#
```

- Add queues for L2TP session interfaces to the QoS profile.

```
host1(config-qos-profile)#l2tp-session queue traffic-class best-effort
scheduler-profile 400k
host1(config-qos-profile)#l2tp-session queue traffic-class voice
scheduler-profile 100k
host1(config-qos-profile)#exit
host1(config)#
```


4. Attach the QoS profile to the interface on which you have configured L2TP.

```
host1(config)#interface gigabitEthernet 6/0
host1(config-if)#qos-profile
```

5. (Optional) Verify the new QoS profile configuration.

```
host1(config)#show qos-profile l2tpQpro25
```

```
qos-profile l2tpQpro25:
t-class interface rule traffic scheduler queue drop statistics
group type type class profile profile profile profile
-----
12tp-session queue best-effort 400k default default default
tcGroup1 12tp-session queue voice 100k default default default
```

Configuring QoS for an L2TP LAC Session

Before you configure QoS for an L2TP LAC session:

- Configure traffic classes.
See “Configuring Traffic Classes That Define Service Levels” on page 14.
- Configure the queuing hierarchy.
See “Configuring Queue Profiles to Manage Buffers and Thresholds” on page 22.
- Configure the scheduler hierarchy and shaping with scheduler profiles.
See “Configuring a Scheduler Hierarchy” on page 47.

To configure QoS for an L2TP LAC session:

1. Configure the QoS profile.

- Create the QoS profile, and enter QoS Profile Configuration mode.

```
host1(config)#qos-profile l2tpQpro25
host1(config-qos-profile)#
```

- Add queues for L2TP session interfaces to the QoS profile.

```
host1(config-qos-profile)#lt2p-session queue traffic-class best-effort
scheduler-profile 400k
host1(config-qos-profile)#lt2p-session queue traffic-class voice
scheduler-profile 100k
host1(config-qos-profile)#exit
host1(config)#
```

2. Attach the QoS profile to the interface on which you have configured L2TP.

```
host1(config)#interface gigabitEthernet 6/0
host1(config-if)#qos-profile l2tpQpro25
```

3. (Optional) Verify the new QoS profile configuration.

```
host1(config)#show qos-profile l2tpQpro25
```

```
qos-profile l2tpQpro25:
t-class interface rule traffic scheduler queue drop statistics
group type type class profile profile profile profile
-----
l2tp-session queue best-effort 400k default default default
tcGroup1 l2tp-session queue voice 100k default default default
```

- Related Topics**
- Supported Interface Types for QoS Profiles on page 131
 - Sample Scheduler Hierarchies for L2TP on page 197
 - group
 - interface
 - qos-profile
 - queue
 - scheduler-profile
 - show qos-profile
 - traffic-class

Configuring QoS for Tunnel-Server Ports for L2TP LNS Sessions

You can configure QoS for a tunnel-service port that can be used as a dynamic interface associated with an L2TP LNS session.

Before you configure QoS for a tunnel-server port:

- Configure the dedicated or shared tunnel-server port.
See *JUNOS Physical Layer Configuration Guide*.
- Configure the traffic classes.
See “Configuring Traffic Classes That Define Service Levels” on page 14.
- Configure the queuing hierarchy.
See “Configuring Queue Profiles to Manage Buffers and Thresholds” on page 22.
- Configure the scheduler hierarchy and shaping with scheduler profiles.
See “Configuring a Scheduler Hierarchy” on page 47.

To configure QoS for the tunnel-server port:

1. Create the QoS profile.

```
host1(config)#qos-profile lns-tsport
```

2. Configure group nodes for the tunnel-server ports.

```
host1(config-qos-profile)#ip queue traffic-class best-effort scheduler-profile
business-data queue-profile data
host1(config-qos-profile)#ip queue traffic-class video scheduler-profile video
queue-profile video
host1(config-qos-profile)#ip queue traffic-class voice scheduler-profile voice
queue-profile voice
host1(config-qos-profile)#server-port group video
host1(config-qos-profile)#server-port group data
host1(config-qos-profile)#server-port group voice scheduler-profile strict-priority
```

3. Create and attach the QoS port-type profile for server ports.

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

- Related Topics**
- For more information about tunnel-server ports, see *JUNOS Physical Layer Configuration Guide*
 - group
 - interface
 - node
 - qos-port-type-profile
 - qos-profile
 - queue
 - scheduler-profile
 - traffic-class
 - tunnel-server

QoS and L2TP TX Speed AVP 24 Overview

You can configure the router to use QoS settings to calculate the transmit connect speed of the subscriber's access interface reported for an L2TP tunneled session. The router reports the transmit connect speed in L2TP Transmit (TX) Speed AVP 24. During the establishment of an L2TP tunneled session, the LAC sends AVP 24 to the LNS to convey the transmit speed of the subscriber's access interface.

Logical Interfaces and Shared-Shaping Rates

You can configure QoS to control the rate for any of the logical interfaces of the following interface columns:

- ATM 1483 subinterface over ATM VP over ATM interface
- PPPoE subinterface over Ethernet interface

- PPPoE subinterface over VLAN subinterface over Ethernet interface

For those logical interfaces with a rate controlled by QoS, QoS reports this configured rate as the transmit connect speed for that interface. For the logical interfaces that do not have a QoS-configured rate, QoS reports the speed of the underlying physical port as the transmit connect speed.

For each logical interface, QoS determines the rate of the interface using either the shaping rate or the shared-shaping rate, if one is configured. The numeric value of the shaping rate or shared-shaping rate is determined as the result of a provider-specified arithmetic expression in a scheduler profile. This expression can either be a constant value, such as 1,000,000, or an expression using QoS parameters, with values supplied by RADIUS or statically in non-volatile storage (NVS).

If the QoS profile or the QoS parameters are configured in RADIUS, these values are used in computing the rate at the time of login. The system can subsequently modify the value of parameters through change of authorization (CoA), Service Manager, or L2C. Modifications are not reflected in the rate QoS reports because they might take place after the LAC has sent the message that contains AVP 24.

Shaping Mode

When the QoS shaping mode is set to cell for an interface, QoS reports the ATM rate. In cell mode, user-specified rates account for cell headers and trailers, which are ATM native rates; therefore, QoS does not convert the rates for AVP 24.

- Related Topics**
- For information about how to configure the transmit connect speed, see the *JUNOS Broadband Access Configuration Guide*
 - For information about shared-shaping rates, see Simple Shared Shaping Overview on page 79 and Compound Shared Shaping Overview on page 99
 - For information about QoS parameters that are configured in RADIUS, see QoS Parameter Overview on page 225

Monitoring QoS Configurations for L2TP

To monitor QoS configurations for L2TP:

- Monitoring the QoS Configuration of Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet Interfaces on page 347
- Monitoring the QoS Configuration of IEEE 802.3ad Link Aggregation Group Bundles on page 349
- Monitoring the QoS Configuration of IP Interfaces on page 345
- Monitoring the QoS Profiles Attached to an Interface on page 339
- Monitoring the Configuration of QoS Port-Type Profiles on page 340
- Monitoring the Configuration of QoS Profiles on page 341

- Monitoring the QoS Scheduler Hierarchy on page 322
- Monitoring Shared Shapers on page 329

Chapter 23

Configuring Interface Sets for QoS

This chapter describes how to configure a set of logical interfaces with the same scheduling and queuing properties using interface sets.

QoS topics are discussed in the following sections:

- Interface Sets for QoS Overview on page 207
- Architecture of Interface Sets for QoS on page 208
- Configuring Interface Sets for Scheduling and Queuing on page 211
- Configuring Interface Supersets for QoS on page 212
- Configuring Interface Sets for QoS on page 213
- Adding Member Interfaces to an Interface Set on page 214
- Creating a QoS Parameter on an Interface Superset or Interface Set on page 216
- Attaching a QoS Profile to an Interface Superset or an Interface Set on page 218
- Deleting an Interface Superset or an Interface Set on page 219
- Example: Configuring Interface Sets for 802.3ad Link Aggregation Groups on page 220

Interface Sets for QoS Overview

This topic describes how to use interface sets to configure a set of logical interfaces with the same scheduling and queuing properties.

Interface sets are supported for VLANs, ATM VCs, and IP.

You can use interface sets for various scenarios in a broadband access network. For example, you can use an interface set to configure a local loop with a small number of subscribers. Interface sets are also useful for grouping a large number of subscribers into a particular service class or for defining traffic engineering aggregates for DSLAMs.

Interface Set Terms

Table 21 on page 208 lists the terminology used for interface sets

Table 21: Interface Set Terms

Term	Description
Interface set	Set of logical interfaces of the same type: VLAN, ATM VC, and IP. An interface set shares a common parent interface.
Interface superset	Set of QoS interface sets and logical interfaces of the same type. A superset shares a common parent interface.
Parent interface	Logical interface associated with a set. All members of an interface set must be configured under this parent interface, or they cannot join an interface set.
VLAN set	Interface set of VLAN subinterfaces.
VLAN superset	Interface superset that contains VLAN interface sets and VLAN subinterfaces, stacked over a common Ethernet major interface.
ATM VC set	Interface set of ATM 1483 subinterfaces.
Restricted VLAN set	Interface set that is restricted to VLAN subinterfaces sharing a common VLAN ID. The parent of the interface set is an S-VLAN. The VLAN subinterfaces that are members of the interface set must have the same S-VLAN ID as the parent.
Spanning VLAN set	Interface set of VLAN subinterfaces that span S-VLANs. The parent of the interface set is an Ethernet major interface. The member VLAN subinterfaces can have different S-VLAN IDs.
Restricted ATM VC set	Interface set that is restricted to ATM VC subinterfaces sharing a common VPI. The parent of the interface set is a VPI. The ATM VC subinterfaces that are members of the interface set must have the same VPI as the parent.
Spanning ATM VC set	Interface set of ATM VC subinterfaces that span VPIs. The parent of the interface set is an ATM major interface. The member ATM VC subinterfaces can have different VPIs.

- Related Topics**
- Architecture of Interface Sets for QoS on page 208
 - Configuring Interface Sets for Scheduling and Queuing on page 211

Architecture of Interface Sets for QoS

To configure groups of logical interfaces, you must configure both interface sets and interface supersets.

When an interface is grouped in an interface set, the logical interface column is modified, and interface set appears below the interface in the column. The interface superset appears below the interface set.

Although interface sets enable you to configure more types of scheduler nodes, the number of node and queue resources supported in the current scheduler hierarchy are the same.

Interface Set Parents and Types

When configuring an interface set, you must assign a parent and the types of member interfaces allowed in the set.

The parent of an interface set is an interface superset. The parent of the interface superset can be any type of interface over which IP can be configured, including ATM VP, Gigabit Ethernet, or 802.3ad LAG.

The parent of the interface superset controls the type of member interfaces you can have in an interface set. Currently, member interface types include VLAN, ATM-VC, and IP. For example, a interface superset with a Gigabit Ethernet or LAG parent interface can only be the parent of interface set that contains VLAN and IP member interfaces. In addition, all members of the interface set must have the same port.

Sample Interface Columns and Scheduler Hierarchies

Figure 57 on page 210 shows a sample interface column using interface sets and interface supersets for VLANs.



Figure 57: VLAN Interface Column with Interface Sets

Figure 58 on page 210 shows a scheduler hierarchy with VLAN nodes at the interface set.

Figure 58: Scheduler Hierarchy with Nodes at Interface Set and Superset

Scheduling and Shaping Interface Sets

You can apply QoS to interface sets and interface supersets in the same way as a logical interface.

Each interface set or interface superset can have a shared shaper applied to it. The constituents of the shared shaper are the scheduler nodes and queues associated with the interface set.

You can use QoS profiles and QoS parameters to manage the scheduling and shaping in the interface set. When you attach a QoS profile to an interface set or an interface superset, the QoS profile applies to all of the interfaces in the set and the superset.

You can create parameter instances for an interface set or a superset by specifying the set or superset as a controlled-interface type and instance-interface type.

- Related Topics**
- Interface Sets for QoS Overview on page 207
 - Configuring Interface Sets for Scheduling and Queuing on page 211
 - For more information about scheduler resources, see Managing System Resources for Nodes and Queues on page 127

Configuring Interface Sets for Scheduling and Queuing

There are a variety of tasks that you need to complete to configure a set of logical interfaces for scheduling and queuing.

To configure a set of logical interfaces:

1. Configure an interface superset.
See “Configuring Interface Supersets for QoS” on page 212.
2. Configure an interface set as members of the superset.
See “Configuring Interface Sets for QoS” on page 213.
3. Add interfaces as members of the interface set.
See “Adding Member Interfaces to an Interface Set” on page 214.
4. Configure the scheduler hierarchy on the interface superset or the interface set.

You can configure the scheduler hierarchy using one of the following methods:

- Attach a QoS profile to an interface superset or an interface set. QoS profiles reference queue, drop, statistics, and scheduler profiles.
See “Attaching a QoS Profile to an Interface Superset or an Interface Set” on page 218.
 - Create a QoS parameter instance on an interface superset or an interface set. QoS parameter instances enable you to configure shaping rates independent of the QoS profile and scheduler profile.
See “Creating a QoS Parameter on an Interface Superset or Interface Set” on page 216.
5. (Optional) Monitor the configuration of interface sets and supersets.
See “Monitoring the Configuration of QoS Interface Sets” on page 350 and “Monitoring the Configuration of QoS Interface Supersets” on page 351.

Related Topics ■ Interface Sets for QoS Overview on page 207

Configuring Interface Supersets for QoS

Tasks to configure an interface superset for QoS include:

- Configuring an Interface Superset on page 212
- Restricting an Interface Superset to an S-VLAN ID or an ATM VP on page 212

Configuring an Interface Superset

To configure an interface superset that contains interface sets and logical interfaces:

1. Create the interface superset.

```
host1(config)#qos-interface-superset vlan-superset-shaping
```

2. Specify the parent interface for the superset.

```
host1(config-interface-superset)#qos-interface-parent interface tenGigabitEthernet 4/0/0
```

You can configure an Ethernet major interface, an ATM major interface, or a LAG for the parent interface. You must define the parent before you add interface sets or subinterfaces to the interface superset.

Restricting an Interface Superset to an S-VLAN ID or an ATM VP

When you configure interface supersets for VLANs or ATM VCs, the member interfaces normally span different S-VLAN IDs or ATM VPs.

Optionally, you can restrict all interface members of a VLAN superset to a specific S-VLAN ID, or all members of an ATM VC to a specific ATM VP.

To restrict the interface members of an interface superset:

1. Specify the interface superset.

```
host1(config)#qos-interface-superset residential-customers
```

2. Restrict the interfaces in the interface superset.

For VLAN supersets, restrict the VLAN subinterfaces to an S-VLAN ID:

```
host1(config-qos-interface-superset)#restricted interface tenGigabitEthernet 4/0/0 svlan id 2
```

For ATM VC supersets, restrict an ATM VC to an ATM VP:

```
host1(config-qos-interface-superset)#restricted interface atm 2/0/0 atm-vp 2
```

- Related Topics**
- To add members to the interface superset, see *Configuring Interface Sets for QoS* on page 213
 - *Monitoring the Configuration of QoS Interface Supersets* on page 351
 - qos-interface-superset
 - qos-interface-parent
 - restricted

Configuring Interface Sets for QoS

Interface sets are members of interface supersets. An interface set joins a superset when it is defined with the superset as a parent.

Tasks to configure an interface set for QoS include:

- *Configuring an Interface Set* on page 213
- *Deleting an Interface Set from an Interface Superset* on page 213

Configuring an Interface Set

To configure an interface set:

1. Configure the interface set.

```
host1(config)#qos-interface-set vlan-business
```

2. Specify the parent interface.

```
host1(config-interface-set)#qos-interface-parent vlan-superset
```

The parent must be an interface superset.

3. Set the member-interface type.

```
host1(config)#member-interface-type vlan
host1(config)#member-interface-type ip
```

You can specify vlan, atm-vc, or ip for the member-interface type.

If the parent interface superset is attached to an Ethernet major interface, the valid interface types are vlan and ip.

If the parent interface superset is attached to an ATM major interface, the valid interface types are atm-vc or ip.

Deleting an Interface Set from an Interface Superset

To delete an interface set from an interface superset:

1. Specify the interface superset.

```
host1(config)#qos-interface-superset business-data
```

2. Delete the interface set from the interface superset.

```
host1(config-qos-interface-superset)#no qos-interface-set atm-vc-data
```

- Related Topics**
- For more information about configuring interface supersets, see *Configuring Interface Supersets for QoS* on page 212
 - *Monitoring the Configuration of QoS Interface Sets* on page 350
 - `member-interface-type`
 - `qos-interface-parent`
 - `qos-interface-set`
 - `restricted`

Adding Member Interfaces to an Interface Set

You can add interfaces as members of an interface set using the CLI or RADIUS.

Tasks to add members to an interface set include:

- *Adding Interfaces to an Interface Set with the CLI* on page 214
- *Adding Interfaces to an Interface Set with RADIUS* on page 215
- *Changing and Deleting Interface Members in an Interface Set* on page 215
- *Changing Interface Members with Upper-Layer Protocols in an Interface Set* on page 215

Adding Interfaces to an Interface Set with the CLI

To add subscriber interfaces to the interface set:

1. Specify the VLAN or ATM-VC subinterface you want to add to the interface set.

For VLAN subinterfaces:

```
host1(config)#interface gigabitEthernet 4/0/0.1
```

For ATM subinterfaces:

```
host1(config)#interface atm 2/0/0.1
host1(config-sub-if)#atm pvc 21 0 21 aal5snap
```

2. Configure the interface set as the parent of this interface.

```
host1(config-sub-if)#qos-interface-parent residential-customers
```

The interface type must match the member-interface type specified in the interface set.

Adding Interfaces to an Interface Set with RADIUS

You can add interfaces to an interface set using the QoS-Interfaceset-Name RADIUS VSA attribute [26-130].

This VSA is useful when configuring local loop topologies of interface sets in the network. When the subscriber interface is created, the VSA supplies the interface name and the subscriber interface. The system matches the subscriber interface with the member-interface type that is specified in the interface set. Note that the VSA cannot specify an interface superset.

When multiple subscribers exist in one interface set, such as PPPoE subscribers over the same VLAN, they are joined with the first subscriber's VSA. The system generates an error message if one of the subscribers attempts to join another interface set using another VSA.

Leaving an interface set is not supported from RADIUS. The interface leaves an interface set when it is deleted or manually removed from the interface set through the CLI.

Changing and Deleting Interface Members in an Interface Set

If you want to move interface members in an interface set, we recommend that you delete the interface member from the interface set before associating it with another interface.

To move an interface member from an interface set to another interface set:

1. Specify the subinterface associated with the member interface.

```
host1(config)#interface gigabitEthernet 4/0/0.1
```

2. Delete the member interface from the subinterface using the **no** version of the command.

```
host1(sub-if)#no qos-interface-parent vlan-business
```

3. Configure the new interface set for the member interface.

```
host1(sub-if)#qos-interface-parent vlan-residential
```

Changing Interface Members with Upper-Layer Protocols in an Interface Set

When upper-layer protocols such as IP are configured on an interface set, moving an interface member from one interface set to another interface set can cause problems with the interface column in munged QoS profiles.

For example, moving an interface set to an upper layer binding causes the interface set to appear below the subinterface level. If a QoS profile were attached to the VLAN subinterface in this example, the munged QoS profile for all IP interfaces stacked above the subinterface would change.

```

host1(config)#interface gigabitEthernet 4/0/0.1
host1(config-sub-if)#svlan id 3 1
host1(config-sub-if)#ip address 1.2.3.4/24
host1(config-sub-if)#qos-interface-parent vlan-business

```

Instead of moving the member interface, we recommend that you add an interface member to an interface set at the subinterface rather than at the upper-layer binding. For example:

```

host1(config)#interface gigabitEthernet 4/0/0.1
host1(config-sub-if)#svlan id 3 1
host1(config-sub-if)#qos-interface-parent vlan-business
host1(config-sub-if)#ip address 1.2.3.4/24

```

- Related Topics**
- For more information about RADIUS VSAs, see the *JUNOS Broadband Access Configuration Guide*
 - `qos-interface-parent`

Creating a QoS Parameter on an Interface Superset or Interface Set

Tasks to create a QoS parameter on an interface superset or an interface set include:

- Configuring a QoS Parameter Definition for an Interface Superset or an Interface Set on page 216
- Creating a QoS Parameter Instance for an Interface Superset on page 217
- Creating a QoS Parameter Instance for an Interface Set on page 217

Configuring a QoS Parameter Definition for an Interface Superset or an Interface Set

You can configure a parameter definition to recognize an interface superset or an interface set as controlled-interface types and instance interface-types.

Controlled-interface types specify resources that the parameter instance can control. Instance-interface types are the interface types to which the QoS client can apply a parameter instance.

To specify an interface set or an interface superset in a QoS parameter definition:

1. Specify the QoS parameter definition.

```

host1(config)#qos-parameter-define business

```

2. Configure the controlled-interface type and specify the interface set or the interface superset.

```

host1(config-qos-parameter-define)#controlled-interface-type set
host1(config-qos-parameter-define)#controlled-interface-type superset

```

You can specify up to four controlled-interface types for each parameter definition.

3. Configure the instance-interface type and specify the interface set or the interface superset.

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

You can specify up to eight instance-interface types for each parameter definition.

Creating a QoS Parameter Instance for an Interface Superset

To create a QoS parameter instance for an interface superset:

1. Specify the QoS interface superset.

```
host1(config)#qos-interface-superset vlan-superset
```

2. Create the QoS parameter instance.

```
host1(config-qos-interface-superset)#qos-parameter business-data
```

3. Attach the QoS profile.

```
host1(config-qos-interface-superset)#qos-profile vlan
```

Creating a QoS Parameter Instance for an Interface Set

To create a QoS parameter instance for an interface set:

1. Specify the QoS interface set.

```
host1(config)#qos-interface-set vlan-set
```

2. Create the QoS parameter instance.

```
host1(config-qos-interface-set)#qos-parameter business-data
```

3. Attach the QoS profile.

```
host1(config-qos-interface-set)#qos-profile vlan
```

- Related Topics**
- Monitoring the QoS Profiles Attached to an Interface on page 339
 - Monitoring the QoS Scheduler Hierarchy on page 322
 - Monitoring Shared Shapers on page 329
 - controlled-interface-type
 - instance-interface-type
 - qos-interface-set
 - qos-interface-superset

- qos-parameter
- qos-profile

Attaching a QoS Profile to an Interface Superset or an Interface Set

You can configure a QoS profile to manage the scheduler resources for an interface set or an interface superset.

Tasks to attach a QoS profile to an interface set or an interface set include:

- Configuring a QoS Profile for an Interface Superset or an Interface Set on page 218
- Attaching a QoS Profile to an Interface Superset on page 218
- Attaching a QoS Profile to an Interface Set on page 219

Configuring a QoS Profile for an Interface Superset or an Interface Set

To configure a QoS profile to manage the scheduler resources for an interface superset or an interface set:

1. Configure the QoS profile.

```
host1(config)#qos-profile vlan
```

2. Specify an interface superset or an interface set as a queue.

For interface supersets:

```
host1(config-qos-profile)#superset queue traffic-class video
```

For interface sets:

```
host1(config-qos-profile)#set queue traffic-class video
```

3. Specify an interface superset or an interface set as a node.

For interface supersets:

```
host1(config-qos-profile)#superset node scheduler-profile video
```

For interface sets:

```
host1(config-qos-profile)#set node scheduler-profile video
```

Attaching a QoS Profile to an Interface Superset

To attach a QoS profile to an interface superset:

1. Specify the QoS interface superset.

```
host1(config)#qos-interface-superset vlan-superset
```

2. Attach the QoS profile.

```
host1(config-qos-interface-superset)#qos-profile business-data
```

Attaching a QoS Profile to an Interface Set

To attach a QoS profile to an interface set:

1. Specify the QoS interface set.

```
host1(config)#qos-interface-superset vlan-set
```

2. Attach the QoS profile.

```
host1(config-qos-interface-set)#qos-profile business-data
```

- Related Topics**
- Monitoring the QoS Profiles Attached to an Interface on page 339
 - Monitoring the QoS Scheduler Hierarchy on page 322
 - Monitoring Shared Shapers on page 329
 - node
 - qos-interface-set
 - qos-interface-superset
 - qos-profile
 - queue

Deleting an Interface Superset or an Interface Set

You can delete an interface superset or an interface set after configuration.

Tasks to delete an interface superset or an interface set include:

- Deleting an Interface Superset on page 219
- Deleting an Interface Set on page 220

Deleting an Interface Superset

You must remove the interface sets from an interface superset before you can delete the interface superset. If the interface superset has a QoS profile attached or a QoS parameter instance, all of these attachments are also deleted.

To delete an interface superset:

1. Remove the interface sets from the interface superset.

```
host1(config)#qos-interface-superset business-data
host1(config-qos-interface-superset)#no qos-interface-set atm-vc-data
```

2. Delete the interface superset.

```
host1(config)#no qos-interface-superset business-data
```

Deleting an Interface Set

You must remove the members of an interface set before you can delete the interface set. If the interface set has a QoS profile attached or a QoS parameter instance, all of these attachments are also deleted.

To delete an interface set:

1. Remove the interface set from the interface member.

```
host1(config)#interface gigabitEthernet 4/0/0.1
host1(sub-if)#no qos-interface-parent vlan-business
```

2. Delete the interface set.

```
host1(config)#no qos-interface-set vlan-business
```

- Related Topics**
- For more information about deleting an interface set from an interface superset, see *Configuring Interface Supersets for QoS* on page 212
 - For more information about deleting interface members of an interface set, see *Configuring Interface Sets for QoS* on page 213
 - qos-interface-parent
 - qos-interface-set
 - qos-interface-superset

Example: Configuring Interface Sets for 802.3ad Link Aggregation Groups

This example shows how to configure interface sets to restrict a VLAN interface to a specific link in an 802.3ad link aggregation group (LAG).

When the parent interface of an interface superset is a LAG, the children of the superset are distributed to different links of the LAG using the load balancing scheme.

You can, however, specify the Ethernet physical port as the anchor of the interface superset. When the link is up, the interface superset is attached to the specified Ethernet port. When the link is down, the system chooses an available link. When the link comes back up, the system moves the interface superset and its member back to the primary link.

To restrict a VLAN interface to a specific link in the LAG:

1. Configure the interface superset.

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

2. Assign the parent interface as LAG.

```
host1(config)#qos-interface-parent interface lag
```

3. Restrict the interface superset to the Ethernet parent interface.

```
host1(config-interface-superset)#restricted interface gigabitEthernet 4/0/0
```

- Related Topics**
- qos-interface-parent
 - qos-interface-superset
 - restricted

Part 6

Managing Queuing and Scheduling with QoS Parameters

- QoS Parameter Overview on page 225
- Configuring a QoS Parameter on page 229
- Configuring Hierarchical QoS Parameters on page 261
- Configuring IP Multicast Bandwidth Adjustment with QoS Parameters on page 269
- Configuring the Shaping Mode for Ethernet with QoS Parameters on page 281
- Configuring Byte Adjustment for Shaping Rates with QoS Parameters on page 293
- Configuring the Downstream Rate Using QoS Parameters on page 301

Chapter 24

QoS Parameter Overview

This chapter provides information about quality of service (QoS) parameters.

QoS parameters are discussed in the following sections:

- QoS Parameter Overview on page 225
- QoS Parameter Audience on page 225
- QoS Parameter Terms on page 226
- Relationship Among QoS Parameters, Scheduler Profiles, and QoS Profiles on page 227

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 Topics**
- QoS Parameter Audience on page 225
 - QoS Parameter Terms on page 226

QoS Parameter Audience

This topic collection contains QoS parameter 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 the scheduler profiles and referencing them from QoS profiles. QoS administrators also configure parameter definitions that control the parameters, interfaces, and ranges of values that QoS clients, using QoS parameters, can assign.

QoS clients are responsible for configuring services for individual subscribers by creating parameter instances. The parameter instances that QoS clients create depend on the settings that the QoS administrator defined in parameter definitions. QoS clients can use the CLI, Session and Resource Control (SRC), IP multicast bandwidth adjustment, RADIUS, or Service Manager to manage these services.

- Related Topics**
- QoS Parameter Overview on page 225
 - Relationship Among QoS Parameters, Scheduler Profiles, and QoS Profiles on page 227

QoS Parameter Terms

Table 22 on page 226 defines terms used in this discussion of QoS parameters.

Table 22: QoS Parameter Terminology Used in This Chapter

Term	Description
Downreference	QoS feature that controls a node or queue lower in the scheduler hierarchy. For example, a QoS profile that is attached to an ATM virtual circuit (ATM VC) modifies QoS settings on ATM virtual path (VP) nodes. You cannot configure downreferences for QoS parameters. We also recommend that you do not configure downreferences for QoS profiles.
Explicit parameter instance	Hierarchical parameter instance whose value is explicitly specified by a client. This term is meaningful only when referring to hierarchical parameter instances; non-hierarchical parameter instances are always explicit.
Hierarchical parameter	Parameter with both explicit instances that are configured by a QoS client, and with implicit instances that are automatically generated for all controlled interfaces. The value for the implicit instance is the sum of the explicit instances for interfaces stacked above the controlled interface.
Implicit parameter instance	Hierarchical parameter instance where the value is the sum of explicit parameter instances on scheduler nodes and queues stacked above them in the scheduler hierarchy.
Parameter definition	Definition of a parameter name and attributes that a QoS administrator creates.
Parameter expression	Parameters used in conjunction with operators. Scheduler profiles reference a parameter definition name within a parameter expression.

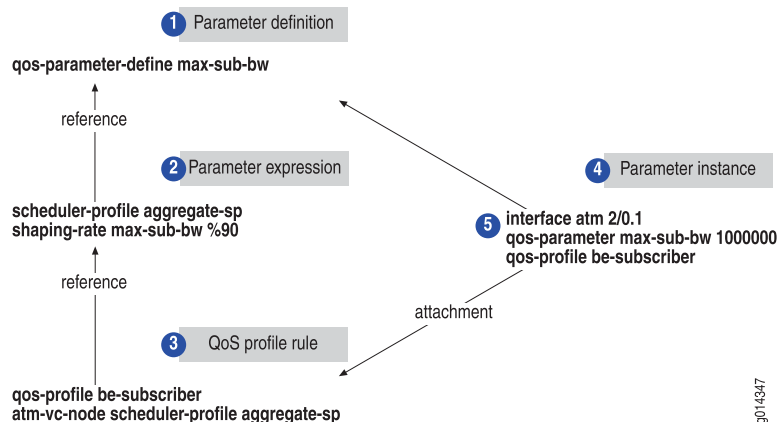
Table 22: QoS Parameter Terminology Used in This Chapter *(continued)*

Term	Description
Parameter instance	Parameter name and value that a QoS client associates with a logical interface.
Parameter value	32-byte unsigned integer value associated with a parameter instance.
QoS administrator	Person responsible for implementing a QoS queuing architecture by configuring QoS profiles, scheduler profiles, and parameter definitions.
QoS client	Person responsible for configuring services for individual subscribers and setting rates for those services by using the parameter definitions and QoS profiles that the QoS administrator configures. QoS clients can use the CLI, SRC, Service Manager, IP multicast bandwidth adjustment, or RADIUS.

Related Topics ■ For definitions of other common QoS terms, see QoS Terms on page 5

Relationship Among QoS Parameters, Scheduler Profiles, and QoS Profiles

Figure 59 on page 227 shows the relationship among the parameter definitions, scheduler profiles, and QoS profiles that QoS administrators create. It also indicates how these profiles control the parameter instances that QoS clients create.

Figure 59: Relationship of Parameter Definitions, Scheduler Profiles, and QoS Profiles

The following sections describe the steps displayed in Figure 59 on page 227, based on the tasks that the QoS administrator performs and those the QoS client performs.

QoS Administrator Tasks

Before the QoS client can specify settings for subscribers by using the QoS parameters feature:

1. The QoS administrator defines the attributes that the QoS client can modify by configuring a parameter definition.
2. The QoS administrator specifies the parameter definition name in a scheduler profile.
3. The QoS administrator references the scheduler profile in a QoS profile rule.

QoS Client Tasks

After the QoS administrator defines parameter definitions:

1. The QoS client creates a parameter instance and associates it with a logical interface.
2. The QoS client attaches a QoS profile to the logical interface.

Related Topics ■ QoS Parameter Audience on page 225

Chapter 25

Configuring a QoS Parameter

This chapter provides information for configuring quality of service (QoS) parameters on E Series routers.

QoS parameters are discussed in the following sections:

- Parameter Definition Attributes for QoS Administrators Overview on page 229
- Scheduler Profiles and Parameter Expressions for QoS Administrators on page 235
- Configuring a Basic Parameter Definition for QoS Administrators on page 238
- Parameter Instances for QoS Clients Overview on page 240
- Creating Parameter Instances on page 242
- Example: QoS Parameter Configuration for Controlling Subscriber Bandwidth on page 243

Parameter Definition Attributes for QoS Administrators Overview

As the QoS administrator, you can create a parameter definition that constrains how a QoS client can create a parameter instance. When QoS clients create a parameter instance, they work within the attributes that you have defined.

Table 23 on page 229 lists the parameter attributes that you can define for a parameter definition.

Table 23: Attributes in Parameter Definitions

Parameter Data Setting	Description
Name	Name for the parameter.
Instance-interface type	Interface types to which the QoS client can apply a parameter instance. The QoS administrator can specify up to eight instance-interface types for each parameter definition.
Controlled-interface type	Interface types that specify resources that the parameter instance can control. The QoS administrator can specify up to four controlled-interface types for each parameter definition.

Table 23: Attributes in Parameter Definitions *(continued)*

Parameter Data Setting	Description
Subscriber-interface type	Subscriber interfaces to which QoS clients can apply parameters obtained through RADIUS or profiles. The QoS administrator can specify up to four subscriber-interface types for each parameter definition.
Range	Valid range of values that a QoS client can specify.
Expression	Boolean that indicates whether the parameter uses implicit parameter instances, which are the sum of explicit instances of the parameter on all scheduler nodes or queues above them in the scheduler hierarchy.
Application	Application that binds parameter instance to a specific application, such as IP multicast bandwidth adjustment.

Naming Guidelines for QoS Parameters

You define the parameter name by issuing the **qos-parameter-define** command to enter QoS Parameter Definition Configuration mode.

The naming guidelines for parameters differ from other QoS features such as QoS profiles and scheduler profiles.

Parameter names must begin with a letter to avoid confusion with numbers and operators. Because QoS clients reference this parameter name to create a parameter instance, we recommend that you use a name that is descriptive.

Table 24 on page 230 lists some sample parameter names and descriptions.

Table 24: Sample Parameter Names

Parameter Name	Description
max-subscriber-bandwidth	Total bandwidth for a subscriber (average of all services)
max-voice-bandwidth	Shaping rate for a subscriber voice queue
min-data-bw	Assured rate for a priority-data service queue
max-data-bw	Shaping rate for the same priority data service queue as min-data-sw

In addition, parameter names cannot be the same as an arithmetic operator. Table 25 on page 231 lists examples of valid and invalid parameter names that use operators.

Table 25: Valid and Invalid Parameter Names

Valid Names	Invalid Names
n1	1
f+	1n
-	+
-	+ foo
-	min
-	max

Parameter names are case-sensitive. For example, max-subscriber-bw and max-Subscriber-bw are different parameter names.

Because the shaping rate and shared-shaping rates determine the maximum scheduler rates, and the assured rate determines minimum scheduler rates, we recommend that you use min or max operands in the parameter name.

Interface Types and QoS Parameters

You can specify the following attributes in a parameter definition to control the scope of a parameter on interfaces:

- Controlled-interface types
- Instance-interface types
- Subscriber-interface types

Controlled-Interface Types

Controlled-interface types specify interface types for queues and scheduler nodes that a parameter instance can control. You can define up to four controlled-interface types for each parameter definition by issuing the **controlled-interface-type** command in QoS Parameter Definition Configuration mode. Examples of controlled interface types include atm-vp (ATM virtual paths), atm-vc (ATM virtual circuits), and VLAN (virtual LANs).

For example, if you specify controlled-interface types of atm-vc and vlan, then you can use the parameter instance to shape or weight an ATM VC or VLAN node. However, because you did not specify ip, the system does not allow this parameter in a scheduler profile that was referenced in a QoS profile with an ip node (for example, ip node scheduler-profile test1).

Controlled-Interface Type Example

In this example, you configure a parameter definition for a scheduler hierarchy in which a VLAN represents a subscriber. The parameter definition specifies that the

parameter controls VLAN nodes and queues and sets the maximum rate for any parameter instance.

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

Then you reference the parameter definition within a scheduler profile.

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

This scheduler profile can be referenced only by QoS profile VLAN rules. When a user attempts to reference the scheduler profile using rules other than VLAN, an error message is displayed. For example, a QoS profile rule cannot associate the scheduler profile with an atm-vc rule, as shown in the following example:

```
host1(config-qos-profile)#atm-vc queue traffic-class best-effort scheduler-profile
subscriber
% scheduler-profile parameter's controlled-interface-type(s) do not control this
atm-vc qos-profile rule type
```

After you reference the parameter in a scheduler profile, you can reference the scheduler profile from a QoS profile. In this example, you configure a vlan node for each subscriber with a shared-shaping rate specified by the parameter max-subscriber-bandwidth.

```
host1(config)#qos-profile subscriber-triple-play
host1(config-qos-profile)#vlan queue traffic-class best-effort scheduler-profile
subscriber
```

Instance-Interface Types

After you configure at least one controlled-interface type, you configure one or more instance-interface types that specify the types of logical interfaces to which the QoS client can apply the parameter. You can define up to eight instance-interface types for each parameter definition by issuing the **instance-interface-type** command in QoS Parameter Definition Configuration mode.

QoS clients cannot create a *downreference* for a parameter instance for instance-interface types that is above the lowest controlled-interface type of the same family in the interface stack.



NOTE: The guidelines are different for using instance-interface types with hierarchical parameters. For more information, see “Scheduler Profiles and Parameter Expressions for QoS Administrators” on page 235.

Instance-Interface Type Example

In the following example, you configure a parameter definition with a controlled-interface type of VLAN. You then enable QoS clients to create a parameter

instances at VLAN, SVLAN, and Ethernet interfaces by configuring instance-interface types of vlan, svlan, and ethernet.

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

In the scheduler hierarchy, IP is above VLANs. If you attempt to configure an instance-interface type for ip, an error message indicates that you cannot downreference IP from VLANs.

```
host1(config-qos-parameter-define)#instance-interface-type ip
% instance-interface-type ip cannot stack above the lowest controlled-interface-type
```

Subscriber-Interface Types

Subscriber-interface types represent subscriber interfaces to which you can apply QoS parameters obtained through RADIUS or SRC. You can define up to four subscriber-interface types for each parameter definition by issuing the **subscriber-interface-type** command in QoS Parameter Definition Configuration mode.

The following interface types are supported:

- ip
- l2tp-session
- atm-vc
- vlan

QoS clients cannot create a parameter instance for subscriber-interface types that is above the lowest controlled-interface type of the same family in the interface stack.

If an interface profile contains a QoS parameter instance rule of max-subscriber-bandwidth 1000000, the system searches the logical interface column, starting at the top, and associates the parameter instance with the first interface with the subscriber-interface type that it locates.

A RADIUS administrator can enter multiple QoS parameter name and value pairs when configuring the RADIUS server with the Juniper Networks VSA [26-82]. This means that the RADIUS can return multiple instances of the same VSA in a single request. For more information about Juniper Networks VSA [26-82], see Juniper Networks VSAs.

Subscriber-Interface Type Example

In the following example, you configure a parameter definition with a controlled-interface type and a subscriber-interface type of IP. These settings enable you to create QoS parameter VSAs on an IP interface.

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

```

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

```

Range of QoS Parameters

You can specify the range of values that the QoS client can enter for a parameter instance by issuing the **range** command in QoS Parameter Definition Configuration mode.

In the following example, you specify that a QoS client can enter a value for the parameter from 512 Kbps to 8 Mbps. The system does not accept values outside of this range.

```

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

```

If the QoS client attempts to configure values outside of this range, a message is displayed.

```

host1(config)#interface fastEthernet 9/0.1
host1(config-subif)#qos-parameter max-subscriber-bandwidth 1000000
host1(config-subif)#exit
host1(config)#interface fastEthernet 9/0.1
host1(config-subif)#qos-parameter max-subscriber-bandwidth 200000
% parameter instance is out of range

```

You cannot create or modify an existing range if the change causes any explicit parameter instance values to be outside the valid range. For example:

```

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
host1(config)#interface fastEthernet 9/0.1
host1(config-subif)# ! This parameter instance is within the range of 512Kbps to
8Mbps.
host1(config-subif)#qos-parameter max-subscriber-bandwidth 1000000
host1(config-subif)#exit
host1(config)#qos-parameter-define max-subscriber-bandwidth
host1(config-qos-parameter-define)#range 2048000 8192000
% cannot modify a range when parameter instances exist with values outside the new
range

```

However, you can remove ranges by using the **no range** command.



NOTE: You can also define a range in parameter expressions when referencing a parameter within a scheduler profile. For more information, see “Specifying a Range in Expressions” on page 237.

Applications and QoS Parameters

You can associate a parameter definition with an application in the system by issuing the **application** keyword with the **qos-parameter-define** command. The applications that you can configure include:

- IP Multicast Bandwidth Adjustment
- QoS Cell Mode
- Byte Adjustment (Cell and Frame)
- QoS Downstream Rate

- Related Topics**
- Configuring a Basic Parameter Definition for QoS Administrators on page 238
 - IP Multicast Bandwidth Adjustment for QoS Overview on page 269
 - Cell Shaping Mode Using QoS Parameters Overview on page 281
 - Byte Adjustment for ADSL and VDSL Traffic Overview on page 293
 - QoS Downstream Rate Application Overview on page 301

Scheduler Profiles and Parameter Expressions for QoS Administrators

After you have created the parameter definition, you reference the parameter within a scheduler profile. You can choose to use parameter expressions in the scheduler profile.

Referencing a Parameter Definition in a Scheduler Profile

You can reference a parameter in a scheduler profile as long as all parameters in the scheduler profile share at least one controlled-interface type. Otherwise, a QoS profile rule cannot reference the scheduler profile.

For example:

```
host1(config)#qos-parameter-define max-subscriber-bandwidth
host1(config-qos-parameter-define)#controlled-interface-type vlan
host1(config-qos-parameter-define)#exit
host1(config)#scheduler-profile subscriber
host1(config-scheduler-profile)#shared-shaping-rate max-subscriber-bandwidth auto
```

When a scheduler profile references a parameter, the system implicitly assigns controlled-interface types to the scheduler profile that are the same as the controlled-interface types of all referenced parameters. The system validates scheduler profile types using the QoS profile rules that refer to those scheduler profiles. For example, if the parameter definition max-sub-bw has the controlled-interface types atm-vc and ip, the scheduler profile cannot be referenced in QoS profile rules that have a type other than atm-vc or ip.

Removing or Modifying a Scheduler Profile

You can modify a scheduler profile as long as the QoS profile rules that use the scheduler profile are of the same type. All nodes and queues controlled by the scheduler profile are adjusted to the new rate.

You can also remove a parameter reference from a scheduler profile. The system modifies the nodes and queues that are controlled by the scheduler profile with the new rate.

Using Expressions for QoS Parameters

Expressions are combinations of parameter names, constants, and operators. You can specify some scheduler profile attributes using an expression, such as the shaping rate. All operations within expressions are performed using 64 bit unsigned math, resulting in a 32 bit, signed integer value.

Expressions consist of both operators and operand values. Operators are arithmetic functions, and operand values are the inputs for the mathematical function. Operand values can be a parameter name or an integer. You specify an expression consisting of an operand, followed by zero or more [operator, operand] pairs.

Simple parameter expressions are displayed in the following example. Simple parameter expressions usually contain a constant rate or a single parameter name.

```
host1(config-scheduler-profile)#shaping-rate 10000000
host1(config-scheduler-profile)#shared-shaping-rate max-sub-bw auto
host1(config-scheduler-profile)#shaping-rate max-sub-be-bw
host1(config-scheduler-profile)#assured-rate assured-bw
```

More complicated parameter expressions are displayed in the following example. Complicated parameter expressions contain combinations of constant rates, parameter names, and operators.

```
host1(config-scheduler-profile)#shaping-rate max-sub-bw % 90
host1(config-scheduler-profile)#shared-shaping-rate max-data-bw + max-voice-bw +
max-video-bw auto
host1(config-scheduler-profile)#assured-rate min-data-bw % oversubscription-rate +
min-video-bw % oversubscription-rate
host1(config-scheduler-profile)#shared-shaping-rate 400000 - multicast-adjustment
burst 100 milliseconds auto
```

Operators and Precedence

Table 26 on page 237 lists the operators that QoS parameters support and the precedence of the operator within the expression.

Table 26: Operators for Parameter Expressions

Operator	Description	Precedence	Examples
%	Percent in the range 1–100	1	<code>max-subscriber-bw % 100</code> <code>max-subscriber-bw % 10</code>
*	Multiplication	1	<code>5 * maxBandwidth</code>
/	Division	1	<code>maxBandwidth / 64000</code>
+	Addition	2	<code>max-subscriber-bw + 50000</code> <code>max-subscriber-bw + l2c-rate</code>
-	Subtraction	2	<code>max-subscriber-bw - 50000</code> <code>max-subscriber-bw - l2c-rate</code>
min	Minimum	3	<code>max-subscriber-bw min 50000</code> <code>max-subscriber-bw min l2c-rate</code>
max	Maximum	3	<code>max-subscriber-bw max 50000</code> <code>max-subscriber-bw max l2c-rate</code>

Specifying a Range in Expressions

You can use the min and max operators to specify the allowable range of an expression result.

For example, to specify a shaping rate at a minimum of 1 Mbps and a maximum of 5 Mbps, use the following expression:

```
host1(config)#scheduler-profile subscriber-rate
host1(config-scheduler-profile)#shaping-rate (( subscriber-rate max 1000000 ) min
5000000 )
```

Operations Using This Expression

1. Take the max of the subscriber-rate scheduler profile, or 1 Mbps, and name it x.
2. Take the min of x and 5 Mbps.

Example 1 The value of the subscriber-rate scheduler profile is less than 1 Mbps, specifically 500,000.

- The max of 500K and 1 Mbps is 1 Mbps
- The min of 1 Mbps and 5 Mbps is 1 Mbps

Result—Made the subscriber-rate a minimum of 1 Mbps.

Example 2 The value of the subscriber-rate scheduler profile is greater than 5 Mbps, specifically 6 Mbps.

- The max of 6 Mbps and 1 Mbps is 6 Mbps
- The min of 6 Mbps and 5 Mbps is 5 Mbps

Result—Made the subscriber-rate a maximum of 5 Mbps.

Example 3 The value of the subscriber-rate scheduler profile is within the range of 1–5 Mbps, specifically 3 Mbps.

- The max of 3 Mbps and 1 Mbps is 3 Mbps
- The min of 3 Mbps and 5 Mbps is 3 Mbps

Result—Maintained the subscriber-rate within the range of 1–5 Mbps.

- Related Topics**
- Using Expressions for Bandwidth and Burst Values in a Scheduler Profile on page 48
 - Configuring a Basic Parameter Definition for QoS Administrators on page 238

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 Topics**
- Parameter Definition Attributes for QoS Administrators Overview on page 229
 - Example: QoS Parameter Configuration for Controlling Subscriber Bandwidth on page 243
 - For more information about configuring a scheduler hierarchy with rates and weights, see Configuring a Scheduler Hierarchy on page 47
 - For more information about configuring a QoS profile, see Configuring a QoS Profile on page 132

- assured-rate
- controlled-interface-type
- instance-interface-type
- node
- qos-parameter-define
- qos-profile
- queue
- range
- scheduler-profile
- shaping-rate
- shared-shaping-rate
- subscriber-interface-type
- traffic-class
- weight

Parameter Instances for QoS Clients Overview

The QoS administrator implements a QoS architecture for the provider based on QoS profiles and parameter definitions. The QoS client creates the parameter instances and attaches QoS profiles to logical interfaces. The QoS client can be a user accessing parameters through CLI or through client software such as RADIUS or SRC.

As a QoS client, you can use QoS parameter instances to set the following attributes of a node or queue:

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

Global QoS Parameter Instance Overview

In the following example, a parameter instance is created in Global Configuration mode.

```
host1(config)#qos-parameter max-subscriber-bandwidth 8000000
```

When you create a parameter instance in Global Configuration mode, the value that you set for a rate becomes the default value for the router. We recommend that you create a global default value for a parameter instance to provide a minimal level of service by default for the router.

QoS Parameters for Interfaces Overview

When you attach a parameter instance to an interface in Interface Configuration mode, the default value for the chassis overrides the default value for the router. When attached to subinterfaces, parameter instances override both interface and global configurations.

In the following example, a parameter instance is created on a Fast Ethernet interface in Interface Configuration mode.

```
host1(config)#interface fastEthernet 9/0.2
host1(config-if)#qos-parameter max-subscriber-bandwidth 8000000
```

Parameter instances have hierarchical scope. The scope of a parameter instance is the set of logical interfaces stacked above the interface upon which you create it. Any interface stacked above the instance that is one of the controlled-interface types that are configured in the parameter definition can have its nodes or queues controlled by that instance. For example, a parameter named max-sub-bw might have logical interface types of IP and l2tp-session; therefore, it controls rates only for nodes and queues associated with those interface types.

For example, the scope of a parameter instance at a S-VLAN can be all VLANs stacked above that particular S-VLAN. Scopes can overlap, for example, if a parameter instance is created for both an S-VLAN and a VLAN. The most specific instance overrides the other instances.

However, you cannot configure QoS parameter instances to downreference through the interface stack. For example, you cannot create a parameter instance with an interface type of ATM VP on an ATM1483 subinterface.

When you attach the parameter instance to an interface, it provides a default subscriber bandwidth for terminated and tunneled subscribers that terminate over that interface. To set parameter instances for a subscriber, a parameter instance is attached to a subscriber interface such as a vlan or atm-vc. The QoS administrator defines the available subscriber-interface types in the parameter definition. The parameter instance overrides the QoS profile attachment lower down the interface stack, providing a subscriber-specific value.

You can attach QoS profiles and QoS parameters to a logical interface in either order. If a scheduler profile calls for a parameter and no parameter instance is defined, the system behaves as if there is no shaping rate, shared-shaping rate, or assured rate for that node or queue.

- Related Topics**
- Creating Parameter Instances
 - For more information about using global parameter instances for IP multicast bandwidth adjustment, see IP Multicast Bandwidth Adjustment for QoS Overview on page 269

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 242
- Creating a Parameter Instance for an Interface on page 242
- Creating a Parameter Instance for an ATM VP on page 242
- Creating a Parameter Instance for an S-VLAN on page 243

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 Topics**
- Parameter Instances for QoS Clients Overview on page 240
 - For information about creating QoS parameter instances for Service Manager, see *JUNOS Broadband Access Configuration Guide*
 - atm-vp qos-parameter
 - atm vp-tunnel
 - encapsulation vlan
 - interface
 - qos-parameter
 - svlan id
 - svlan qos-parameter

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 60 on page 244.

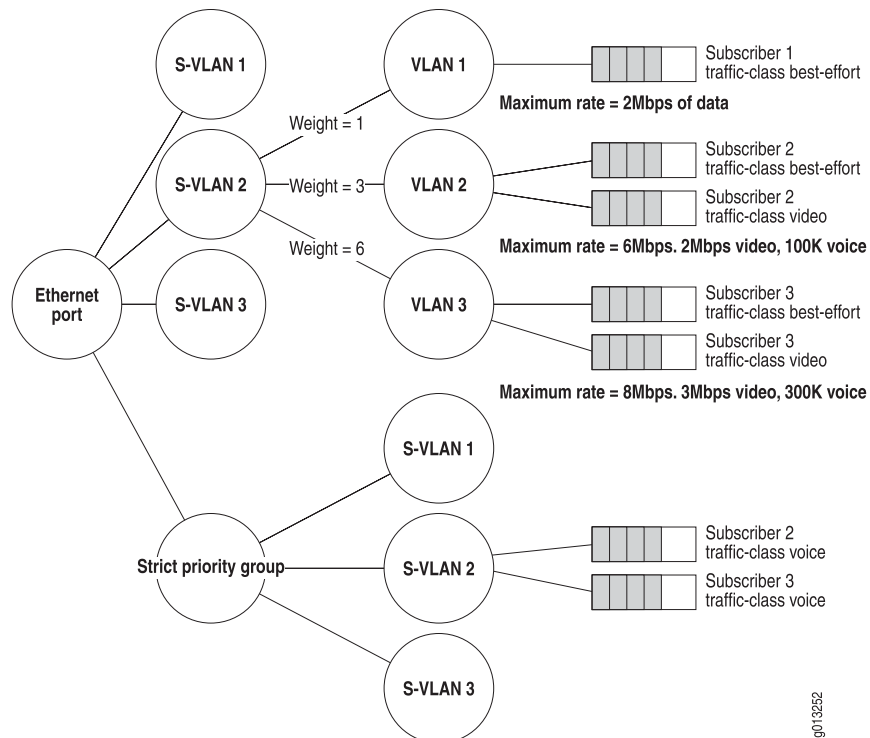
Figure 60: Physical Network Topology



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

Figure 61 on page 244 shows the QoS scheduler hierarchy that the QoS client creates when configuring a different service for each subscriber.

Figure 61: 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 61 on page 244 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)#vlan 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 61 on page 244. 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)#vlan 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 a 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
```

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

2. 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

```

3. 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>

4. 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
-----
vlan node subscriber-weight
svlan node default
vlan queue best-effort subscriber-best-effort default default

vlan queue video subscriber-video default default
EF svlan node default
EF vlan queue voice subscriber-voice default default
statistics
profile
-----

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
```

interface	class	traffic rate	forwarded rate	aggregate drop rate	minimum rate	maximum 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
vlan Eth9/0.1	vlan node			
	A vlan queue best-effort	2000000		rate 2000000
vlan Eth9/0.2	vlan node			
	A vlan queue best-effort	6000000		rate 6000000
	A vlan queue video		2000000	
	A vlan queue EF voice		100000	
vlan Eth9/0.3	vlan node			
	A vlan queue best-effort	8000000		rate 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 *JUNOS 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
```


Chapter 26

Configuring Hierarchical QoS Parameters

This chapter provides information for configuring hierarchical quality of service (QoS) parameters on E Series routers.

QoS parameters are discussed in the following sections:

- Hierarchical QoS Parameters Overview on page 261
- Guidelines for Configuring Hierarchical Parameters on page 261
- Configuring a Parameter Definition to Calculate Hierarchical Instances on page 262
- Example: QoS Parameter Configuration for Hierarchical Parameters on page 263

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 Topics**
- Configuring a Parameter Definition to Calculate Hierarchical Instances on page 262
 - For information about the IP multicast bandwidth adjustment application, see IP Multicast Bandwidth Adjustment for QoS Overview on page 269

Guidelines for Configuring Hierarchical Parameters

Use the following guidelines when specifying a hierarchical parameter:

- You can specify only a subset of the instance-interface types that are supported for non-hierarchical parameters. The following output lists the instance-interface types that are supported:

```
host1(config)#qos-parameter-define hierarchical-parameter hierarchical
host1(config-qos-parameter-define)#instance-interface-type ?
atm-vc ATM Virtual Circuit (VC)
ip IP interface
ipv6 IP version 6 interface
l2tp-session L2tp session interface
vlan VLAN subinterface
```

- You can specify only one instance-interface type per hierarchical parameter. For example:

```
host1(config)#qos-parameter-define hierarchical-parameter hierarchical
host1(config-qos-parameter-define)#instance-interface-type ip
host1(config-qos-parameter-define)#instance-interface-type vlan
% only one instance-interface-type can be specified for a hierarchical parameter
```

- Hierarchical instance-interface types cannot stack above the *highest* controlled-interface type. For example:

```
host1(config)#qos-parameter-define hierarchical-parameter hierarchical
host1(config-qos-parameter-define)#controlled-interface-type ip
host1(config-qos-parameter-define)#instance-interface-type vlan
% hierarchical instance-interface-type vlan cannot stack above
controlled-interface-type ip
```

In contrast, a non-hierarchical instance-interface type cannot stack above the *lowest* controlled-interface type (vlan). For example:

```
host1(config)#qos-parameter-define non-hierarchical-parameter
host1(config-qos-parameter-define)#controlled-interface-type vlan
host1(config-qos-parameter-define)#instance-interface-type ip
% instance-interface-type ip cannot stack above the lowest controlled-interface-type
```

- You must specify a subscriber-interface type that is identical to the instanceinterface type that you specified.

Related Topics ■ Configuring a Parameter Definition to Calculate Hierarchical Instances on page 262

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 Topics**
- Hierarchical QoS Parameters Overview on page 261
 - Configuring a Basic Parameter Definition for QoS Administrators on page 238
 - Configuring a Parameter Definition for IP Multicast Bandwidth Adjustment on page 271
 - Example: QoS Parameter Configuration for Hierarchical Parameters on page 263
 - qos-parameter-define

Example: QoS Parameter Configuration for Hierarchical Parameters

The example in this section illustrates how to configure hierarchical parameters for VLANs and S-VLANs.

Figure 62 on page 264 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 62: Hierarchical Parameters Scheduler Hierarchy

Procedure for QoS Administrators

This section describes the procedures to configure the scheduler hierarchy shown in Figure 62 on page 264 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 *JUNOS 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
vlan 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
vlan 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
vlan 10 qos-profile qp-shape-svlan
```

Related Topics ■ Hierarchical QoS Parameters Overview on page 261

Chapter 27

Configuring IP Multicast Bandwidth Adjustment with QoS Parameters

This chapter provides information for configuring quality of service (QoS) parameters on E Series routers.

QoS parameters are discussed in the following sections:

- IP Multicast Bandwidth Adjustment for QoS Overview on page 269
- Guidelines for Configuring IP Multicast Adjustment for QoS on page 271
- Configuring a Parameter Definition for IP Multicast Bandwidth Adjustment on page 271
- Example: QoS Parameter Configuration for IP Multicast Bandwidth Adjustment on page 273

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 1000000 - 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 Topics

- Guidelines for Configuring IP Multicast Adjustment for QoS on page 271
- For more information about multicast bandwidth maps and QoS adjustment, see *JUNOS Multicast Routing Configuration Guide*

- For more information about configuring scheduler rates for QoS parameters, see Scheduler Profiles and Parameter Expressions for QoS Administrators on page 235

Guidelines for Configuring IP Multicast Adjustment for QoS

When you specify the IP multicast bandwidth adjustment application, the following considerations apply:

- You must specify a controlled-interface type.
- You cannot specify any instance-interface types or subscriber-interface types. By default, the system assigns a default instance-interface type of ip.
- When you specify the IP multicast bandwidth adjustment application, the parameter definition is hierarchical. You must specify the **hierarchical** keyword with the **application** keyword.
- The system prevents you from defining more than one parameter definition with the **ip-multicast** application specified. For example:

```
host1(config)#qos-parameter-define vpShaper application ip-multicast hierarchical
host1(config-qos-parameter-define)#controlled-interface-type ip
host1(config-qos-parameter-define)#exit
host1(config)#qos-parameter-define bar application ip-multicast hierarchical
% there cannot be more than one parameter defined with this property
```

- Parameter instances associated with the IP multicast bandwidth adjustment application are not stored in non-volatile storage (NVS). (Parameter definitions are stored in NVS.) Because the application is activated based on IGMP joins and leaves received on an interface, the system removes the instances when you turn off or reset the router, then re-creates it based on new messages received on an interface.

- Related Topics**
- Configuring a Parameter Definition for IP Multicast Bandwidth Adjustment on page 271

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 *JUNOS 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 Topics**
- IP Multicast Bandwidth Adjustment for QoS Overview on page 269
 - Example: QoS Parameter Configuration for IP Multicast Bandwidth Adjustment on page 273
 - controlled-interface-type
 - encapsulation vlan

- interface gigabitEthernet
- node
- qos-parameter-define
- qos-profile
- queue
- scheduler-profile
- shared-shaping-rate
- traffic-class
- vlan id

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 63 on page 274 shows the scheduler hierarchy built in this configuration.

Figure 63: 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 1.1.1.1 255.255.255.0
```

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
           traffic forwarded aggregate minimum maximum
interface class rate drop rate rate 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
atm-vc ATM2/0.1 A	atm-vc node	10000000		rate 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 *JUNOS 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 Topics ■ IP Multicast Bandwidth Adjustment for QoS Overview on page 269

Chapter 28

Configuring the Shaping Mode for Ethernet with QoS Parameters

This chapter provides information for configuring the shaping mode for Ethernet using quality of service (QoS) parameters on E Series routers.

QoS parameters are discussed in the following sections:

- Cell Shaping Mode Using QoS Parameters Overview on page 281
- Guidelines for Configuring the Cell Shaping Mode with QoS Parameters on page 283
- Configuring a Parameter Definition to Shape Ethernet Traffic Using Cell Mode on page 284
- Example: QoS Parameter Configuration for QoS Cell Mode and Byte Adjustment for Cell Shaping on page 286

Cell Shaping Mode Using QoS Parameters Overview

You can associate the QoS cell mode application (**qos-cell-mode**) with a parameter definition for Ethernet interfaces configured on any E Series Broadband Services Routers.

Creating a parameter instance with the QoS cell mode application on a VLAN subinterface enables the scheduler to perform cell mode shaping and scheduling for queues and nodes associated with the controlled-interface types above the logical interface on which you create the parameter instance.

Overriding the QoS Shaping Mode

The QoS cell mode application overrides the shaping mode specified at the port using the **qos-shaping-mode** command.

The QoS cell mode application applies the shaping mode to all logical interfaces specified in the controlled-interface type list above the logical interface on which you created the parameter instance.

For example, all of the interfaces stacked above the Gigabit Ethernet interface configured on slot 6, adapter 0, port 2 have cell shaping mode:

```
host1(config)#interface gigabitEthernet 6/0/2  
host1(config-if)#qos-shaping-mode cell
```

The QoS administrator then applies frame shaping mode to the Gigabit Ethernet interface configured on slot 6, adapter 0, port 2, subinterface 1 using the QoS cell mode application. This parameter instance overrides the shaping mode configured at the port.

```
host1(config-if)#interface gigabitEthernet 6/0/2.1
host1(config-if)#qos-parameter cell-mode 0
```

Module Types and Capabilities for QoS Cell Mode Application

The QoS cell mode application is supported by all E Series routers. However, different module types support the application.

Table 27 on page 282 lists the supported modules for the **qos-shaping-mode cell** command and the **qos-cell-mode** application for parameters. It also describes how the cell mode adjustment is performed by each module type.

Table 27: Supported Interfaces for qos-shaping-mode and qos-cell-mode Commands

Module Type	qos-shaping-mode cell Command	qos-cell-mode Application	Adjustment Performed By
Ethernet interfaces on ES2 4G LM and ES2 10G LM (E120 and E320 Broadband Services routers)	✓	✓	Internal cell-taxing mechanism
Ethernet interfaces on GE-2 and GE-HDE line modules (ERX7xx models, ERX14xx models, and ERX310 routers)	✓	✓	Internal cell-taxing mechanism
Ethernet interfaces on ERX7xx models, ERX14xx models, and ERX310 routers	–	✓	Parameter expression associated with qos-cell-mode application (See “Cell Tax Adjustment Using QoS Cell Mode” on page 282.)
ATM interfaces on all E Series routers	✓	–	Internal cell-taxing mechanism
All other interface types on all E Series routers	–	–	–

Cell Tax Adjustment Using QoS Cell Mode

The internal cell-taxing mechanism does not perform the cell mode adjustment on certain interface types. On these interfaces, the system uses a parameter expression

associated with the **qos-cell-mode** application to determine whether the cell adjustment is required.



NOTE: Do not use the parameter expression on Ethernet interfaces configured on the ES2 4G LM, GE-2 line module, or the GE-HDE line module.

For example, the subscriber-rate parameter represents the bandwidth of a subscriber. The shaping rate for the parameter is calculated by referencing an expression that represents the cell mode adjustment in a scheduler profile:

```
(config-scheduler-profile)# shaping-rate subscriber-rate - subscriber-rate * cell-mode  
% 25
```

The subscriber-rate - subscriber-rate * cell-mode % 25 expression provides for an explicit cell-tax factor of 25 percent when the subscriber local loop is transmitting cells. In cases where the local loop is very-high-bit-rate digital subscriber line (VDSL), the second term in the expression drops to 0.

Relationship with QoS Downstream Rate Application

ANCP dynamically controls the QoS cell mode application when you create parameter instances for VLANs using both the QoS downstream rate application and the QoS cell mode application.

ANCP controls QoS cell mode parameter instances at the VLAN subinterface only; the protocol does not control parameter instances at the major Ethernet interface or S-VLAN subinterface.

- Related Topics**
- Configuring a Parameter Definition to Shape Ethernet Traffic Using Cell Mode on page 284
 - Example: QoS Parameter Configuration for QoS Cell Mode and Byte Adjustment for Cell Shaping on page 286
 - For more information about configuring the QoS shaping mode, see QoS Shaping Mode for Ethernet Interfaces Overview on page 178
 - For more information about configuring scheduler rates for QoS parameters, see Scheduler Profiles and Parameter Expressions for QoS Administrators on page 235
 - For more information about shaping the downstream rate using QoS parameter instances that were created dynamically by ANCP, see QoS Downstream Rate Application Overview on page 301

Guidelines for Configuring the Cell Shaping Mode with QoS Parameters

When you specify the QoS cell mode application, the following considerations apply:

- You can have only one parameter definition with the QoS cell mode application configured.
- You must specify a controlled-interface type.
- You can specify only instance-interface types of atm, atm-vp, atm-vc, ethernet, svlan, and vlan.
- You can specify only the subscriber-interface type of vlan when you configure QoS cell mode application on its own or with the byte adjustment application. When you configure the QoS cell mode application with the QoS downstream rate application, you must specify a subscriber-interface type. ANCP uses the subscriber-interface type to determine the instance-interface type on which to dynamically create the parameter.
- You can specify only 0 or 1 as the values for a parameter instance with the QoS cell mode application configured. 0 indicates frame mode, and 1 indicates cell mode. You cannot configure another range for the parameter definition using the **range** command.

Related Topics

- Cell Shaping Mode Using QoS Parameters Overview on page 281
- Configuring a Parameter Definition to Shape Ethernet Traffic Using Cell Mode on page 284
- Example: QoS Parameter Configuration for QoS Cell Mode and Byte Adjustment for Cell Shaping on page 286
- For more information about configuring scheduler rates for QoS parameters, see Scheduler Profiles and Parameter Expressions for QoS Administrators on page 235
- For more information about the QoS downstream rate application, see QoS Downstream Rate Application Overview on page 301

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 Topics**
- Configuring a Parameter Definition to Shape Ethernet Traffic Using Cell Mode on page 284
 - Example: QoS Parameter Configuration for QoS Cell Mode and Byte Adjustment for Cell Shaping on page 286
 - For more information about configuring scheduler rates for QoS parameters, see Scheduler Profiles and Parameter Expressions for QoS Administrators on page 235
 - For more information about the QoS downstream rate application, see QoS Downstream Rate Application Overview on page 301
 - controlled-interface-type
 - instance-interface-type
 - interface gigabitEthernet
 - ip address
 - qos-parameter
 - qos-parameter-define
 - scheduler-profile
 - svlan qos-parameter
 - traffic-class
 - vlan id


Example: QoS Parameter Configuration for QoS Cell Mode and Byte Adjustment for Cell Shaping

The example in this section illustrates how to configure the byte adjustment application to adjust the shaping rate for downstream ATM traffic from the customer premise equipment (CPE) to Ethernet interfaces configured on an E320 router.

In this example, the QoS administrator manages the shaping rate using a combination of the byte adjustment application and cell shaping mode to account for different layer 2 encapsulations and the ATM cell pad, header, and trailer.

Figure 64 on page 286 displays the Ethernet network to which the QoS administrator applies the byte adjustment.

Figure 64: Byte Adjustment for VC1 and VC2



In Figure 64 on page 286, VLAN 1 and VLAN 2 map to the subscribers at VC1 and VC2.

The QoS administrator allocates a total of 10 Mbps of bandwidth for voice, video, and data services to VC1, and 2 Mbps of bandwidth of data traffic for VC2.

Table 28 on page 286 lists the shaping rate and byte adjustment for both subscribers.

Table 28: Byte Adjustment for Subscribers VC1 and VC2

	VC1	VC2
Protocol	A3 encapsulation	A1 encapsulation
Byte Adjustment	-28	-2
Voice Bandwidth	1000000 bps	1000000 bps
Video Bandwidth	10000 bps	–
Data Bandwidth	8000000 bps	–
Total Bandwidth	–	1000000 bps

Configuring Traffic Classes

The QoS administrator configures the traffic classes and traffic-class groups for video and voice services.

1. Configure the traffic class named voice.

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

2. Configure the traffic class named video.

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

Configuring the QoS Parameter Definition

The QoS administrator configures a parameter definition and the byte adjustment application. The QoS administrator then enables the QoS client to create a parameter instance of the byte adjustment from VLAN interfaces. All interfaces above the VLAN use the same byte adjustment value.

1. Configure a parameter definition named byte-adjustment.

```
host1(config)#qos-parameter-define byte-adjustment application
qos-byte-adjustment
```

2. Define the controlled-interface types for vlan and ip to adjust the shaping rate for the VLAN and IP queues.

- a. Configure the controlled-interface type for VLAN.
- b. Configure the controlled-interface type for IP.

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

Configuring the QoS Shaping Mode

The QoS administrator then configures the QoS shaping mode using the QoS cell mode application. When you configure the QoS shaping mode to cell mode on port 0 of the IOA, all ports on the IOA use the same value.

1. Configure a parameter definition named cell-mode.

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

2. Define the controlled-interface types for vlan and ip for the shaping mode.

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

Reference the Parameter Definition Within a Scheduler Profile

The QoS administrator configures the shaping rate and the shared-shaping rate within scheduler profiles for the subscribers at VC1 and VC2.

1. Configure the scheduler profile for the subscriber VC1.
 - a. Configure the scheduler profile named vc1.

- b. Configure the shared-shaping rate of 10000000 with a burst of 10 milliseconds.

```
host1(config)#scheduler-profile vc1
host1(config-scheduler-profile)#shared-shaping-rate 10000000 burst 10
milliseconds
host1(config-scheduler-profile)#exit
```

2. Configure the scheduler profile for the voice service.
 - a. Configure the scheduler profile named voice.
 - b. Configure the shared-shaping rate of 100000 with a burst of 10 milliseconds.

```
host1(config)#scheduler-profile voice
host1(config-scheduler-profile)#shaping-rate 100000 burst 10 milliseconds
host1(config-scheduler-profile)#exit
```

3. Configure the scheduler profile for the video service.
 - a. Configure the scheduler profile named voice.
 - b. Configure the shared-shaping rate of 8000000 with a burst of 10 milliseconds.

```
host1(config)#scheduler-profile video
host1(config-scheduler-profile)#shaping-rate 8000000 burst 10 milliseconds
host1(config-scheduler-profile)#exit
```

4. Configure the scheduler profile for the subscriber VC2.
 - a. Configure the scheduler profile named vc2.
 - b. Configure the shared-shaping rate of 1000000 with a burst of 10 milliseconds.

```
host1(config)#scheduler-profile vc2
host1(config-scheduler-profile)#shaping-rate 1000000 burst 10 m
host1(config-scheduler-profile)#exit
```

Adding the Scheduler Profiles to a QoS Profile

After configuring the scheduler profiles, the QoS administrator then configures QoS profiles for subscribers VC1 and VC2.

1. Configure the vc1 QoS profile with a shared-shaping rate of 10 Mbps.
 - a. Configure the QoS profile vc1.
 - b. Configure the vlan node and reference the scheduler profile vc1.
 - 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 vc1
host1(config-qos-profile)#vlan node scheduler-profile vc1
host1(config-qos-profile)#vlan queue traffic-class voice schedule-profile voice
host1(config-qos-profile)#vlan queue traffic-class video schedule-profile video
host1(config-qos-profile)#exit
```


2. Configure the vc2 QoS profile with a shaping rate of 1 Mbps.
 - a. Configure the QoS profile vc2.
 - b. Configure the vlan node and reference the scheduler profile vc2.

```
host1(config)#qos-profile vc2
host1(config-qos-profile)#vlan node scheduler-profile vc2
host1(config-qos-profile)#exit
```

Attaching the Parameter Definition to an Interface

The QoS administrator creates logical interfaces for VLAN1 and VLAN2 and attaches the parameter definitions to them.

1. Attach the parameter definition 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. Create a parameter instance for byte-adjustment with a value of -28.
 - f. Create a parameter instance for cell-mode with a value of 1 (cell shaping mode).
 - g. Attach the QoS profile vc1 to the Gigabit Ethernet 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-parameter byte-adjustment -28
host1(config-if)#qos-parameter cell-mode 1
host1(config-if)#qos-profile vc1
host1(config-if)#exit
```

2. Attach the parameter definition to VLAN2.
 - a. Specify the Gigabit Ethernet interface in slot 6, adapter 0, port 1.
 - b. Assign a VLAN ID of 2.
 - c. Create a parameter instance for byte-adjustment with a value of -2.
 - d. Create a parameter instance for cell-mode with a value of 1 (cell shaping mode).
 - e. Attach the QoS profile vc2 to the Gigabit Ethernet interface.

```
host1(config-if)#interface gigabitEthernet 6/0/1.1
host1(config-if)#vlan id 2
host1(config-if)#qos-parameter byte-adjustment -2
host1(config-if)#qos-parameter cell-mode 1
host1(config-if)#qos-profile vc2
```

```
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 *JUNOS 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 byte-adjustment QoS parameter definition.
qos-parameter-define byte-adjustment application qos-byte-adjustment
controlled-interface-type vlan
controlled-interface-type ip
exit
! Create the cell-mode QoS parameter definition.
qos-parameter-define cell-mode application qos-cell-mode
controlled-interface-type vlan
controlled-interface-type ip
exit
! Configure the vc1 and vc2 scheduler profiles.
scheduler-profile vc1
shared-shaping-rate 10000000 burst 10 milliseconds
exit
scheduler-profile voice
shaping-rate 100000 burst 10 milliseconds
exit
scheduler-profile video
shaping-rate 8000000 burst 10 milliseconds
exit
scheduler-profile vc2
shaping-rate 1000000 burst 10 m
exit
! Add the scheduler profiles to the vc1 QoS profile.
qos-profile vc1
vlan node scheduler-profile vc1
vlan queue traffic-class voice schedule-profile voice
vlan queue traffic-class video schedule-profile video
exit
qos-profile vc2
vlan node scheduler-profile vc2
```

```
! Configure the byte adjustment for VLAN1 and VLAN2.  
interface gigabitEthernet 6/0/0  
encapsulation vlan  
interface gigabitEthernet 6/0/0.1  
vlan id 1  
qos-parameter byte-adjustment -28  
qos-parameter cell-mode 1  
qos-profile vc1  
interface gigabitEthernet 6/0/1.1  
vlan id 2  
qos-parameter byte-adjustment -2  
qos-parameter cell-mode 1  
qos-profile vc2
```

Related Topics ■ Cell Shaping Mode Using QoS Parameters Overview on page 281

Chapter 29

Configuring Byte Adjustment for Shaping Rates with QoS Parameters

This chapter provides information for configuring byte adjustment with quality of service (QoS) parameters on E Series routers.

QoS parameters are discussed in the following sections:

- Byte Adjustment for ADSL and VDSL Traffic Overview on page 293
- Guidelines for Configuring Byte Adjustment of Cell and Frame Shaping Rates Using QoS Parameters on page 296
- Configuring a Parameter Definition to Adjust Cell Shaping Rates for ADSL Traffic on page 297
- Configuring a Parameter Definition to Adjust Frame Shaping Rates for VDSL Traffic on page 299

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, JUNOS software supports a byte adjustment application (**qos-byte-adjustment**) to adjust rates for cell shaping mode. For VDSL traffic, JUNOS 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 65 on page 294 displays an example of an Ethernet encapsulation and an ATM encapsulation.

Figure 65: Byte Adjustment Calculation for Ethernet and ATM Encapsulations

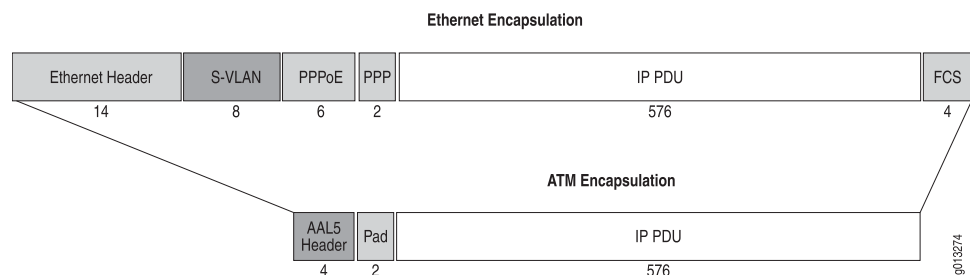


Table 29 on page 294 lists the header lengths for the Ethernet encapsulation, which represents the CPE protocol overhead. The hierarchy is PPPoE over S-VLAN over Ethernet.

Table 29: 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

Table 29: Header Lengths for Ethernet Encapsulation *(continued)*

Header	Number of Bytes
Total	34 bytes

Table 30 on page 295 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 30: 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 31 on page 296 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 31: 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 Topics**
- Configuring a Parameter Definition to Adjust Cell Shaping Rates for ADSL Traffic on page 297
 - Configuring a Parameter Definition to Adjust Frame Shaping Rates for VDSL Traffic on page 299
 - Example: QoS Parameter Configuration for QoS Cell Mode and Byte Adjustment for Cell Shaping on page 286
 - For more information about configuring shaping modes for Ethernet, see QoS Shaping Mode for Ethernet Interfaces Overview on page 178 and Cell Shaping Mode Using QoS Parameters Overview on page 281
 - For more information about shaping the downstream rate using QoS parameter instances that were created dynamically by ANCP, see QoS Downstream Rate Application Overview on page 301

Guidelines for Configuring Byte Adjustment of Cell and Frame Shaping Rates Using QoS Parameters

When you specify the cell or frame byte-adjustment application, the following considerations apply:

- You can have only one QoS parameter definition with the cell byte-adjustment application (**qos-byte-adjustment**) configured.
- You can only have one QoS parameter definition with the frame byte-adjustment application (**qos-frame-byte-adjustment**) configured.
- You can specify only instance-interface types of lag, ethernet, svlan, and vlan.
- You can specify only an subscriber-interface type of vlan.
- The available range for parameters with the byte adjustment application is -32–63. You cannot configure another range using the **range** command.

- We recommend that you apply the byte adjustment parameter at the lowest interface column so that upper interfaces automatically have the parameter.
- On the ES2 10G LM, the shaping rate adjustment is performed more efficiently by the TFA ASIC than ASICS on other modules. The TFA ASIC performs an internal adjustment of 4 bytes. The maximum byte adjustment value that you can configure is 59. When you configure a byte adjustment value greater than 59 in a QoS parameter, the system automatically resets the value to 59.

Related Topics

- Byte Adjustment for ADSL and VDSL Traffic Overview on page 293
- Configuring a Parameter Definition to Adjust Cell Shaping Rates for ADSL Traffic on page 297
- Configuring a Parameter Definition to Adjust Frame Shaping Rates for VDSL Traffic on page 299
- Example: QoS Parameter Configuration for QoS Cell Mode and Byte Adjustment for Cell Shaping on page 286
- For more information about configuring shaping modes for Ethernet, see QoS Shaping Mode for Ethernet Interfaces Overview on page 178 and Cell Shaping Mode Using QoS Parameters Overview on page 281

Configuring a Parameter Definition to Adjust Cell Shaping Rates for ADSL Traffic

You can adjust shaping rates to account for different layer 2 encapsulations as well as the ATM cell pad, header, and trailer on interfaces with cell shaping mode using the **qos-byte-adjustment** application.



NOTE: 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.

To associate a parameter instance with the byte adjustment application:

1. Configure the 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 byteadjust1 application
qos-byte-adjustment
```

- 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. Do one of the following:

- Configure the shaping mode by issuing the **qos-shaping-mode** command.
Frame shaping mode is the default for Ethernet interfaces on all E Series routers. You can only set the cell shaping mode for Gigabit Ethernet and 10-Gigabit Ethernet interfaces configured on the GE-2 line module, the GE-HDE line module, and the ES2 4G LM.
- Configure the shaping mode by specifying the QoS cell mode application with a parameter definition.

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

4. Attach the parameter definition to a logical Ethernet interface.

In this example, parameter instances are created for both the byte adjustment and QoS cell mode applications.

```
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 1
host1(config-if)#qos-parameter byteadjustment -16
host1(config-if)#qos-parameter cell-mode 1
host1(config-if)#ip address 1.1.1.1 255.255.255.0
```

- Related Topics**
- Byte Adjustment for ADSL and VDSL Traffic Overview on page 293
 - Guidelines for Configuring Byte Adjustment of Cell and Frame Shaping Rates Using QoS Parameters on page 296
 - Example: QoS Parameter Configuration for QoS Cell Mode and Byte Adjustment for Cell Shaping on page 286
 - For information about managing packet fragmentation for traffic with frame shaping mode, see Configuring a Parameter Definition to Adjust Frame Shaping Rates for VDSL Traffic on page 299
 - For more information about configuring shaping modes for Ethernet, see QoS Shaping Mode for Ethernet Interfaces Overview on page 178 and Cell Shaping Mode Using QoS Parameters Overview on page 281
 - controlled-interface-type
 - encapsulation vlan
 - instance-interface-type
 - ip address

- node
- qos-parameter
- qos-parameter-define
- qos-profile
- queue
- traffic-class
- vlan id

Configuring a Parameter Definition to Adjust Frame Shaping Rates for VDSL Traffic

Packet fragmentation can occur at a DSLAM because of the associated segment header that is added for VDSL2 in frame shaping 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 VDSL downstream rate at the DSLAM. You can also configure the cell mode application to account for ADSL downstream traffic that is also being received.

To configure a QoS parameter definition to adjust frame shaping rates and manage packet fragmentation:

1. Configure the QoS parameter definition to accept downstream shaping rate instantiation from ANCP.

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

2. Configure the QoS parameter definition for the frame byte-adjustment application to adjust the packet header.

```
host1(config)#qos-parameter-define frame-byte application  
qos-frame-byte-adjustment
```

You can also configure the qos-byte-adjustment application with a different value.

3. Create the QoS parameter definition for the cell mode application to track the subscriber DSL type.

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

The ADSL type corresponds to cell mode and VDSL corresponds to frame mode.

4. Configure the parameter expression to reduce the shaping rate to account for packet fragmentation.

In the following expression, the adjustment is applied to traffic with frame shaping mode only. The byte adjustment value is 8 and the shaping rate is reduced by 2 percent.

```
host1(config)#scheduler-profile service-provider-business
host1(config-scheduler-profile)# shaping-rate ancp-downstream -
(ancp-downstream % 2 * (1 - sp-qos-cell-mode))
```



TIP: To determine the expression value and the byte adjustment required, you must account for the actual segmentation header overhead added by the DSLAM. DSLAMs have different segmentation header overheads.

If the user packet size changes, you must change the expression value and the byte adjustment value.

5. To ensure that the router handles the byte adjustment value consistently for VDSL and ADSL networks, apply the QoS parameter for frame shaping mode globally.

```
host1(config)#qos-parameter frame-byte 8
```



NOTE: The ancp-downstream rate and sp-qos-cell-mode QoS parameters are dynamically applied to QoS by ANCP.

- Related Topics**
- Byte Adjustment for ADSL and VDSL Traffic Overview on page 293
 - qos-parameter
 - qos-parameter-define
 - qos-profile
 - scheduler-profile
 - shaping-rate

Chapter 30

Configuring the Downstream Rate Using QoS Parameters

This chapter provides information for configuring quality of service (QoS) parameters on E Series routers.

QoS parameters are discussed in the following sections:

- QoS Downstream Rate Application Overview on page 301
- Guidelines for Configuring QoS Downstream Rate on page 303
- Configuring a Parameter Definition for QoS Downstream Rate on page 303
- Example: QoS Parameter Configuration for QoS Downstream Rate on page 305

QoS Downstream Rate Application Overview

You can associate the QoS downstream rate (**qos-downstream-rate**) application with a parameter definition. The QoS downstream rate application enables you to shape the downstream rate of VLANs and ATM VCs based on parameter instances that are created dynamically by the Access Node Control Protocol (ANCP), also known as the layer 2 control (L2C) protocol, or the Actual-Data-Rate-Downstream [26-130] DSL Forum vendor-specific attribute (VSA). The values of the parameter instances track the bandwidth of the local loop that is communicated by ANCP or the [26-130] VSA.

Downstream Rate and the Shaping Mode

After you configure a parameter definition with the QoS downstream rate application, you can configure the shaping mode for the VLAN or ATM VC. For ATM VCs, use the `qos-shaping-mode` command.

For VLANs, you can use the QoS cell mode application with QoS parameters to perform a cell mode adjustment. ANCP creates instances of the parameter based on the DSL type of the local loop associated with the VLAN.

VLANs configured on the ES2 4G LM on the E120 and E320 Broadband Services routers use an internal cell-taxing mechanism to perform the cell mode adjustment. For VLANs configured on all other E Series Broadband Services Routers, you must also configure a parameter expression to configure the cell mode adjustment.

QoS Adaptive Mode and Downstream Rate

After you create the parameter definition, you must enable QoS adaptive mode for ANCP by issuing the **qos-adaptive-mode** command. ANCP uses this setting to dynamically create the parameter instances for the QoS downstream rate application and, if applicable, the QoS cell mode application. It also uses the setting to determine the value that the system uses when recalculating the shaping rate.

For example, if you created a parameter definition with the QoS cell mode application, ANCP configures parameter instances associated with a value of 0 to indicate a frame-oriented DSL types such as VDSL2. ANCP configures cell-oriented DSL types such as ADSL with a value of 1.

Table 32 on page 302 lists the DSL types, interface type, and resultant shaping modes that ANCP configures when creating a parameter instance for the QoS cell mode application.

Table 32: Access Loop Types and Resultant Shaping Mode

Access Loop Type	Access Loop Interface Type	Shaping Mode
ADSL1	ATM	Cell
ADSL2	ATM	Cell
ADSL2 +	ATM	Cell
VDSL1	ATM	Cell
VDSL2	Ethernet	Frame
SDSL/SHDSL	ATM	Cell

Obtaining Downstream Rates from a DSL Forum VSA

You can configure the QoS downstream rate application to shape VLANs or ATM VCs based on downstream rates obtained from the Actual-Data-Rate-Downstream [26-130] DSL Forum vendor-specific attribute (VSA).

- Related Topics**
- Configuring a Parameter Definition for QoS Downstream Rate on page 303
 - Example: QoS Parameter Configuration for QoS Downstream Rate on page 305
 - For more information about configuring the shaping mode for Ethernet interfaces, see QoS Shaping Mode for Ethernet Interfaces Overview on page 178 and Cell Shaping Mode Using QoS Parameters Overview on page 281
 - For more information about configuring byte adjustment for downstream rates, see Byte Adjustment for ADSL and VDSL Traffic Overview on page 293

- For information about configuring the shaping mode for ATM interfaces, see *Configuring the QoS Shaping Mode for ATM Interfaces* on page 174
- For more information about DSL Forum VSAs, see *JUNOS Broadband Access Configuration Guide*

Guidelines for Configuring QoS Downstream Rate

When you specify the QoS downstream rate application, the following considerations apply:

- You can have only one parameter definition with the QoS downstream rate configured.
- You must specify a controlled-interface type.
- You must configure a subscriber-interface-type. ANCP uses the subscriber-interface type to determine the instance-interface type on which to dynamically create the parameter.
- Access loops can synchronize after the user has logged in. The business logic depends on the rate that is reported in the Access-Request message. We recommend that service providers use RADIUS Connect-Info attribute [77] as the default value for their business logic. When the ANCP rate information is not present, the system uses the default QoS parameter instance (which can be defined globally or per VLAN). The advisory transmit speed configurable per VLAN is reported to the RADIUS Connect-Info attribute [77]. Ensure that the value of the default QoS parameter is aligned with the value in RADIUS Connect-Info attribute 77.

- Related Topics**
- QoS Downstream Rate Application Overview on page 301
 - Configuring a Parameter Definition for QoS Downstream Rate on page 303
 - Example: QoS Parameter Configuration for QoS Downstream Rate on page 305
 - For more information about the RADIUS Connect-Info attribute, see *JUNOS Broadband Access Configuration Guide*

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 284.

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

See “Configuring the QoS Shaping Mode for ATM Interfaces” on page 174.

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
```

- Related Topics**
- Example: QoS Parameter Configuration for QoS Downstream Rate on page 305
 - For information about downstream rate and RADIUS, see *JUNOS Broadband Access Configuration Guide*
 - For more information about configuring ANCP (L2C) parameters, see *JUNOS IP Services Configuration Guide*
 - `aaa qos downstream-rate`
 - `controlled-interface-type`
 - `encapsulation vlan`
 - `instance-interface-type`
 - `ip address`
 - `node`
 - `qos-parameter`
 - `qos-adaptive-mode`
 - `qos-parameter-define`
 - `qos-profile`
 - `queue`
 - `shared-shaping-rate`
 - `subscriber-interface-type`
 - `traffic-class`
 - `vlan id`

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 33 on page 306 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 33: Shaping Rate and Shaping Mode

	VLAN1	VLAN2
Shaping mode	Cell	Cell
Shaping rate	10000000 bps	100000 bps

Configuring Traffic Classes

The QoS administrator configures the traffic classes for voice and video services.

1. Configure the traffic class named voice.

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

2. Configure the traffic class named video.

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

Configuring the QoS Parameter Definition for QoS Downstream Rate

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.

1. Configure a parameter definition named ancpVlan.

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

2. Define the controlled-interface types for vlan and ip to adjust the shaping rate for the VLAN and IP queues.

- a. Configure the controlled-interface type for VLAN.
- b. Configure the controlled-interface type for IP.

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

3. Define the subscriber-interface types for vlan and ethernet.

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

Configuring the QoS Parameter Definition for QoS Cell Mode

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.

1. Configure a parameter definition named cellmodeVlan.

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

2. Define the controlled-interface types for vlan and ip for the shaping mode.

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

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

3. Define the subscriber-interface types for vlan and ethernet.

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

Enabling QoS Adaptive Mode

The QoS administrator enables QoS adaptive mode for ANCP.

1. Enter Layer 2 Control Configuration mode.

```
host1(config)#l2c
```

2. Enable QoS adaptive mode for the system.

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

Reference the Parameter Definition Within a Scheduler Profile

The QoS administrator configures the shaping rate and the shared-shaping rate within scheduler profiles for the VLAN1 and VLAN2 subscribers.

1. Configure the scheduler profile for the subscriber vlan1.
 - a. Configure the scheduler profile named vlan1.
 - b. Configure the shared-shaping rate by referencing the ancpVlan parameter with a burst of 10 milliseconds.

```
host1(config)#scheduler-profile vlan1
host1(config-scheduler-profile)#shared-shaping-rate ancpVlan burst 10
milliseconds auto
host1(config-scheduler-profile)#exit
```

2. Configure the scheduler profile for the voice service.
 - a. Configure the scheduler profile named voice.
 - b. Configure the shaping rate of 100000 with a burst of 10 milliseconds.

```
host1(config)#scheduler-profile voice
host1(config-scheduler-profile)#shaping-rate 100000 burst 10 milliseconds
host1(config-scheduler-profile)#exit
```

3. Configure the scheduler profile for the video service.
 - a. Configure the scheduler profile named video.
 - b. Configure the shaping rate of 8000000 with a burst of 10 milliseconds.

```
host1(config)#scheduler-profile video
host1(config-scheduler-profile)#shaping-rate 8000000 burst 10 milliseconds
host1(config-scheduler-profile)#exit
```

4. Configure the scheduler profile for the subscriber vlan2.

- a. Configure the scheduler profile named vlan2.
- b. Configure the shaping rate by referencing the ancpVlan parameter with a burst of 10 milliseconds.

```
host1(config)#scheduler-profile vlan2
host1(config-scheduler-profile)#shaping-rate ancpVlan burst 10 milliseconds
host1(config-scheduler-profile)#exit
```

Adding the Scheduler Profiles to a QoS Profile

After configuring the scheduler profiles, the QoS administrator then configures QoS profiles for the VLAN1 and VLAN2 subscribers.

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 *JUNOS 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 Topics ■ [QoS Downstream Rate Application Overview on page 301](#)

Part 7

Monitoring and Troubleshooting QoS

- Monitoring QoS on E Series Routers on page 313
- Troubleshooting QoS on page 357

Chapter 31

Monitoring QoS on E Series Routers

This chapter provides information for monitoring specific QoS configurations.



NOTE: The E120 and E320 Broadband Services Routers output for monitor and show commands is identical to output from other E Series routers, except that the E120 and E320 router output also includes information about the adapter identifier in the interface specifier (slot/adapter/port).

QoS topics are discussed in the following sections:

- Monitoring Service Levels with Traffic Classes on page 314
- Monitoring Service Levels with Traffic-Class Groups on page 315
- Monitoring Queue Thresholds on page 316
- Monitoring Queue Profiles on page 319
- Monitoring Drop Profiles for RED and WRED on page 320
- Monitoring the QoS Scheduler Hierarchy on page 322
- Monitoring the Configuration of Scheduler Profiles on page 327
- Monitoring Shared Shapers on page 329
- Monitoring Shared Shaper Algorithm Variables on page 330
- Monitoring Forwarding and Drop Events on the Egress Queue on page 331
- Monitoring Forwarding and Drop Rates on the Egress Queue on page 333
- Monitoring Queue Statistics for the Fabric on page 337
- Monitoring the Configuration of Statistics Profiles on page 338
- Monitoring the QoS Profiles Attached to an Interface on page 339
- Monitoring the Configuration of QoS Port-Type Profiles on page 340
- Monitoring the Configuration of QoS Profiles on page 341
- Monitoring the QoS Configuration of ATM Interfaces on page 343
- Monitoring the QoS Configuration of IP Interfaces on page 345
- Monitoring the QoS Configuration of Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet Interfaces on page 347
- Monitoring the QoS Configuration of IEEE 802.3ad Link Aggregation Group Bundles on page 349

- Monitoring the Configuration of QoS Interface Sets on page 350
- Monitoring the Configuration of QoS Interface Supersets on page 351
- Monitoring the AAA Downstream Rate for QoS on page 352
- Monitoring QoS Parameter Instances on page 352
- Monitoring QoS Parameter Definitions on page 355

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      fabric      fabric
class        weight     strict
-----
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 34 on page 314 lists the **show traffic-class** command output fields.

Table 34: 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

Table 34: show traffic-class Output Fields (continued)

Field Name	Field Description
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 Topics**
- Configuring Traffic Classes That Define Service Levels on page 14
 - show traffic-class

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 35 on page 315 lists the **show traffic-class-group** command output fields.

Table 35: show traffic-class-group Output Fields

Field Name	Field Description
traffic-class group	Name of the traffic-class group

Table 35: show traffic-class-group Output Fields (continued)

Field Name	Field Description
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 Topics**
- Configuring Traffic-Class Groups That Define Service Levels on page 15
 - show traffic-class-group

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 (OMB - 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 (0MB - 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 36 on page 319 lists the **show qos queue-thresholds** command output fields.

Table 36: 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 Topics**
- Configuring Queue Profiles to Manage Buffers and Thresholds on page 22
 - show qos queue-thresholds

Monitoring Queue Profiles

Purpose Display information about queue profiles and references to queue profiles.

Action To display information about all queue profiles:

```
host1#show queue-profile
      committed  conformed  exceeded  fraction:
queue  length:    length:    length:    conformed,  buffer
profile min, max  min, max  min, max  exceeded    weight
-----
default 0, <none>  0, <none>  0, <none>  50, 25      8
```

To display the number of times that a QoS profile references a queue profile:

```
host1#show queue-profile brief
queue-profile default referenced 31 times in qos-profiles
```

To display a list of QoS profiles that reference the queue profile:

```
host1#show queue-profile references
queue-profile default
  Referenced by QoS profiles:
    atm-default
    serial-default
    ethernet-default
    server-default
```

Meaning Table 37 on page 320 lists the **show queue-profile** command output fields.

Table 37: show queue-profile Output Fields

Field Name	Field Description
queue profile	Name of the queue profile
committed length	Greater queue length than the length of the conformed or exceeded length
conformed length	A queue length that is less than the committed length but greater than the exceeded length
exceeded length	A queue length less than the conformed length which is less than the committed length
conformed fraction	Percentage of the total queue that can be occupied before conformed packets are dropped
exceeded fraction	Percentage of the total queue that can be occupied before exceeded packets are dropped
buffer weight	Weight of the queue

- Related Topics**
- Configuring Queue Profiles to Manage Buffers and Thresholds on page 22
 - show queue-profile

Monitoring Drop Profiles for RED and WRED

Purpose Display information about drop profiles and references to drop profiles.

Action To display information about all drop profiles:

```
host1#show drop-profile
               committed      conformed      exceeded
               threshold:      threshold:      threshold:
```


drop profile	Average length exponent	min, max, max drop prob	min, max, max drop prob	min, max, max drop prob
default	0	0, <none>, <none>	0, <none>, <none>	0, <none>, <none>
drop1	10	0, 750000, 80%	0, <none>, <none>	0, <none>, <none>
drop2	10	0, 750000, 80%	0, <none>, <none>	0, <none>, <none>
drop3	10	0, 750000, 80%	0, <none>, <none>	0, <none>, <none>
drop4	10	0, 750000, 80%	0, <none>, <none>	0, <none>, <none>
drop5	0	0, 750000, 80%	0, <none>, <none>	0, <none>, <none>
drop6	10	0, <none>, <none>	0, <none>, <none>	0, <none>, <none>
drop7	10	10%, 90%, 5%	0, <none>, <none>	0, <none>, <none>
drop8	10	0, 750000, 80%	0, <none>, <none>	0, <none>, <none>
drop9	10	0, 750000, 80%	0, <none>, <none>	0, <none>, <none>
drop10	10	0, 750000, 80%	0, <none>, <none>	0, <none>, <none>
drop11	10	0, 750000, 80%	0, <none>, <none>	0, <none>, <none>
drop12	10	0, 750000, 80%	0, <none>, <none>	0, <none>, <none>
drop13	10	0, 750000, 80%	0, <none>, <none>	0, <none>, <none>
drop14	10	0, 750000, 80%	0, <none>, <none>	0, <none>, <none>
drop15	10	0, 750000, 80%	0, <none>, <none>	0, <none>, <none>

To display information about drop profiles in condensed format:

```
host1#show drop-profile brief
```

To display the QoS profiles that reference the drop profile:

```
host1#show drop-profile rates1 references
```

Meaning Table 38 on page 321 lists the **show drop-profile** command output fields.

Table 38: show drop-profile Output Fields

Field Name	Field Description
drop profile	Name of the drop profile
Average length exponent	Exponent used to weight the average queue length over time, controlling WRED responsiveness
committed threshold	Minimum and maximum committed queue thresholds and maximum drop probability
conformed threshold	Minimum and maximum conformed queue thresholds and maximum drop probability
exceeded threshold	Minimum and maximum exceeded queue thresholds and maximum drop probability

- Related Topics**
- Configuring RED on page 27
 - Configuring WRED on page 31
 - show drop-profile

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

assured

interface	resource	shaping rate	shared shaping rate	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 vlnset1          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

```

                                     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 vlnset1          set queue EF EF      wgt 8

```

Scheduler hierarchy for traffic-class group AF

```

                                     assured
                                     shared  rate
                                     shaping  shaping  or
                                     rate    rate    weight
-----
ethernet Eth1/0/0 ethernet port      wgt 8
superset cluster      superset node AF      100000000      wgt 8
set vlnset1          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	assured		
		shared		rate
		shaping	shaping	or
rate		rate	rate	weight
-----	-----	-----	-----	-----
ethernet Eth1/0/0	ethernet port			wgt 8
superset cluster	superset node		800000000	wgt 8
set v1anset1	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 v1anset2	set node			wgt 8
vlan Eth1/0/0.3	vlan queue best-effort			wgt 8

Scheduler hierarchy for traffic-class group EF

interface	resource	assured		
		shared		rate
		shaping	shaping	or
rate		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 v1anset1	set queue EF EF			wgt 8
set v1anset2	set queue EF EF			wgt 8

Scheduler hierarchy for traffic-class group AF

interface	resource	assured		
		shared		rate
		shaping	shaping	or
rate		rate	rate	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
set vlnaset2            set node AF                                wgt 8
vlan Eth1/0/0.3          vlan queue AF AF                        wgt 8

```

Meaning Table 39 on page 327 lists the **show qos scheduler-hierarchy** command output fields.

Table 39: 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
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 Topics**
- Configuring a Scheduler Hierarchy on page 47
 - Configuring Simple Shared Shaping on page 81
 - Configuring Compound Shared Shaping on page 100
 - Configuring Interface Sets for QoS on page 213
 - Configuring Interface Supersets for QoS on page 212
 - show qos scheduler-hierarchy

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
scheduler  rate      burst  weight  strict  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 40 on page 328 lists the **show scheduler-profile** command output fields.

Table 40: 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 Topics**
- Configuring a Scheduler Hierarchy on page 47
 - Configuring Simple Shared Shaping on page 81
 - Configuring Compound Shared Shaping on page 100
 - `show scheduler-profile`

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
atm-vc ATM11/0.10	A atm-vc node	500000		rate 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		rate 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 41 on page 330 lists the **show qos shared-shaper** command output fields.

Table 41: 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	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
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 Topics**
- Configuring a Scheduler Hierarchy on page 47
 - Configuring Simple Shared Shaping on page 81
 - Configuring Compound Shared Shaping on page 100
 - Configuring Interface Sets for QoS on page 213
 - Configuring Interface Supersets for QoS on page 212
 - show qos shared-shaper

Monitoring Shared Shaper Algorithm Variables

Purpose Display information about the user-configurable variables for controlling the simple shared shaper algorithm.

Action To display information about all variables:

```
host1#show qos-shared-shaper-control
control
```

control name	value	units
maximum voql	400	milliseconds
reaction factor	75	percent
convergence factor	50	percent
minimum dynamic rate	0	percent

Meaning Table 42 on page 331 lists the **show qos shared-shaper-control** command output fields.

Table 42: show qos shared-shaper-control Output Fields

Field Name	Field Description
control name	Name of the simple shared shaper control
control value	Value of the simple shared shaper control; default values are displayed if none specified
units	Expressed units for the value of the simple shared shaper control

- Related Topics**
- Configuring Simple Shared Shaper Algorithm Variables on page 93
 - show qos-shared-shaper-control

Monitoring Forwarding and Drop Events on the Egress Queue

Purpose Display information about forwarding and drop event counts on the egress queue.

Action To display events for a specific interface:

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

interface	traffic class	forwarded events	committed drop events	conformed drop events	exceeded drop events	rate period count
ip GigabitEthernet1/0	tc1	132	0	0	0	132
	tc2	132	132	0	0	132
	tc3	6	0	132	0	132
	tc4	0	0	0	132	132

To display events for an L2TP session:

```
host1#show egress-queue events l2tp-session session1
```

To display events 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 egress-queue events tunnel-server 6/0
```

To display events for queues only on the specified interface and not stacked above the interface:

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

To display the sum of events for the queues bound to interfaces that are stacked above the specified interface:

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

To display events for queues belonging to a specific traffic class:

```
host1#show egress-queue events gigabitEthernet 1/0 traffic-class voice
```

To filter output based on the number of events that exceed the specified value.

```
host1#show egress-queue events gigabitEthernet 1/0 event-exceeding committed
host1#show egress-queue events gigabitEthernet 1/0 event-exceeding conformed
host1#show egress-queue events gigabitEthernet 1/0 event-exceeding exceeded
host1#show egress-queue events gigabitEthernet 1/0 event-exceeding forwarded
```

Meaning Table 43 on page 332 lists the **show egress-queue events** command output fields.

Table 43: show egress-queue events Output Fields

Field Name	Field Description
interface	Name of the interface
traffic class	Name of the traffic class
forwarded events	Number of forwarded rate events
committed drop events	Number of committed drop events
conformed drop events	Number of conformed drop events
exceeded drop events	Number of exceeded drop events
rate period count	Time frame during which events are counted (in seconds)

- Related Topics**
- Configuring Event Statistics on page 40
 - show egress-queue events

Monitoring Forwarding and Drop Rates on the Egress Queue

Purpose Display information about forwarding and drop rates on the egress queue. The **show egress-queue rates** command is useful even if no statistics profiles are configured. You can view information about all of the queues even if statistics gathering has not been enabled.

The minimum rate for the queue is the minimum rate at which a node or queue can transmit when all other nodes and queues compete for bandwidth. The system determines the minimum rates by the weight and assured rate configured in a scheduler profile, and are subject to shaping rate and shared-shaping rate configured.

The maximum rate is the maximum rate at which a node or queue can transmit when there are no other nodes or queues competing for bandwidth. The system calculates the maximum rate as the minimum of all shaping rates, shared-shaping rates, and the port rate from the node or queue down to the port.

For example, if a scheduler column configured over a Fast Ethernet port consists of a VLAN queue that has been shaped to 5 Mbps over a VLAN node that has been shaped to 8 Mbps, over an S-VLAN node which is not shaped, then:

- The VLAN queue that is shared-shaped to 5 Mbps has a maximum rate of 5 Mbps.
- The VLAN node that is shaped to 8 Mbps has a maximum rate of 8 Mbps.
- The S-VLAN node which is not shaped has a maximum rate of 100 Mbps.
- The Fast Ethernet port with a bandwidth of 100 Mbps has a maximum rate of 100 Mbps.

Action To display rate statistics only for queues that have queue rate statistics enabled:

```
host1# show egress-queue rates brief interface fastEthernet 9/0.2
```

interface	traffic class	forwarded rate	aggregate drop rate	minimum rate	maximum rate
ip FastEthernet9/0.2	best-effort	0	0	25000	1000000
	videoTrafficClass	0	0	375000	1000000
	multicastTrafficClass	0	0	925000	1000000
	internetTrafficClass	0	0	50000	1000000
Total:		0	0		

```

Queues reported: 4
Queues filtered (under threshold): 0
Queues disabled (no rate period): 0
Queues disabled (no resources): 0
Total queues: 4

```

To display rate statistics by color rather than as an aggregate of all colors:

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

interface	traffic class	forwarded rate	committed drop rate	conformed drop rate	exceeded drop rate
-----	-----	-----	-----	-----	-----

```

ip GigabitEthernet1/0  tc1      14645184      0      0      0
                        tc2      11950400      2706400      0      0
                        tc3      9960792      0      4707200      0
                        tc4      7967200      0      0      6705600
Queues reported:      4
Queues filtered (under threshold): 0
Queues disabled (no rate period): 1
Queues disabled (no resources): 0
Total queues:      5

```

To display rate statistics all of the configured queues, along with the minimum and maximum rates for the queues, even when statistics gathering has not been enabled:

```

host1#show egress-queue rates full interface atm 11/0

```

interface	traffic class	forwarded rate	aggregate drop rate	minimum rate	maximum rate
ip ATM11/0.1	best-effort	*	*	24979	30000000
	tc1	0	0	14987510	30000000
	tc2	0	0	9991673	30000000
	tc3	0	0	4995836	30000000
ip ATM11/0.2	best-effort	*	*	19980	20000000
	tc1	0	0	11988011	20000000
	tc2	0	0	7992007	20000000

```

Queues reported:      5
Queues filtered (under threshold): 0
* Queues disabled (no rate period): 2
**Queues disabled (no resources): 0
Total queues:      7

```

To display rate statistics based on an S-VLAN:

```

host1# show egress-queue rates interface gigabitEthernet 11/0 svlan 0

```

interface	traffic class	forwarded rate	aggregate drop rate	minimum rate
svlan GigabitEthernet 11/0 svlan 0 tc1	0	0		166666666
vlan GigabitEthernet 11/0.1	tc1	0	0	166666666
ip GigabitEthernet 11/0.1	best-effort	0	0	0
vlan GigabitEthernet 11/0.2	tc2	0	0	0
ip GigabitEthernet 11/0.2	best-effort	0	0	0

interface	maximum rate
svlan GigabitEthernet 11/0 svlan 0	1000000000
vlan GigabitEthernet 11/0.1	1000000000
ip GigabitEthernet 11/0.1	1000000000
vlan GigabitEthernet 11/0.2	1000000000
ip GigabitEthernet 11/0.2	1000000000

```

Queues reported:      5
Queues filtered (under threshold): 0
* Queues disabled (no rate period): 0
**Queues disabled (no resources): 0
Total queues: 5

```

In the output of this command, the aggregate of all drop rates—WRED, tail, and forwarding events—is displayed in the aggregate drop rate field. You cannot

distinguish among the counters used for different drop rates from the output of this command. As a result, for ES2 10G ADV LMs, you cannot identify the counters used for committed, conformed, and exceeded packet dropping by WRED functionality from the value displayed in this field. View the value displayed for the Dropped by WRED committed field in the output of the **show ip interface** command to know the cumulative number of committed, conformed, and exceeded packets dropped by WRED for ES2 10G ADV LMs.

To display rate statistics for the previous or current rate period:

```
host1#show egress-queue rates previous interface gigabitEthernet 11/0 svlan
0
host1#show egress-queue rates current interface gigabitEthernet 11/0 svlan
0
```

To display rate statistics for an L2TP session:

```
host1#show egress-queue rates l2tp session session1
```

To display rate statistics 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 egress-queue rates tunnel-server 6/0
```

To display rate statistics for queues bound to the specified interface:

```
host1#show egress-queue rates interface gigabitEthernet 11/0 svlan 0 explicit
```

To display the sum of all rates of queues bound to interfaces that are stacked above the specified interface.

```
host1#show egress-queue rates interface gigabitEthernet 11/0 svlan 0 summary
```

To display rate statistics for queues belonging to a specific traffic class:

```
host1#show egress-queue rates interface gigabitEthernet 11/0 svlan 0 traffic-class
voice
```

To filter output based on the number of queues with rates that exceed the specified value.

```
host1#show egress-queue rates gigabitEthernet 1/0 rate-exceeding committed
host1#show egress-queue rates gigabitEthernet 1/0 rate-exceeding conformed
host1#show egress-queue rates gigabitEthernet 1/0 rate-exceeding exceeded
host1#show egress-queue rates gigabitEthernet 1/0 rate-exceeding forwarded
```

Meaning Table 44 on page 335 lists the **show egress-queue rates** command output fields.

Table 44: show egress-queue rates Output Fields

Field Name	Field Description
interface	Name of the interface

Table 44: show egress-queue rates Output Fields (continued)

Field Name	Field Description
traffic class	Name of the traffic class
forwarded rate	<p>Statistics for the rate at which packets are enqueued.</p> <p>In some time periods, the enqueue rate might exceed the dequeue rate. This can occur when a burst of traffic arrives at a queue which might be dequeuing at a slower rate because of a shaper or congestion.</p> <p>In other time periods, the enqueue rate might be less than the dequeue rate. This can occur when a buffered burst of packets are being dequeued, and no new packets are arriving at the queue.</p>
aggregate drop rate	Total number of all drop rates
committed drop rate	Drop rate for green packets
conformed drop rate	Drop rate for yellow packets
exceeded drop rate	Drop rate for red packets
minimum rate	Minimum rate for queue
maximum rate	Maximum rate for queue
Queues reported	Number of queues reported
Queues filtered (under threshold)	Number of queues not reported because they are under the threshold
Queues disabled (no rate period)	Number of queues not displayed because statistics gathering is disabled (that is, the referenced statistics profile does not have a rate period set)
Queues disabled (no resources)	Number of queues not displayed because no resources were available
Total queues	Total number of queues within the hierarchical scope of the command

- Related Topics**
- Configuring Rate Statistics on page 39
 - For more information about the assured rate, see Configuring an Assured Rate for a Scheduler Node or Queue on page 55
 - show egress-queue rates

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    egress
      class      slot    type    forwarded    forwarded    dropped    dropped
                        packets    bytes    packets    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 45 on page 337 lists the **show fabric-queue** command output fields.

Table 45: 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
dropped packets	Number of dropped packets
dropped bytes	Number of dropped bytes

- Related Topics**
- Configuring Rate Statistics on page 39
 - Configuring Event Statistics on page 40
 - show fabric-queue

Monitoring the Configuration of Statistics Profiles

Purpose Display information about statistics profiles.

Action To display information about all statistics profiles:

```
host1#show statistics-profile
      forwarding committed conformed exceeded
statistics rate drop drop drop rate
profile threshold threshold threshold threshold period
-----
default <none> <none> <none> <none> <none>
statpro-1 10000000 2000000 4000000 6000000 30
```

To display the number of times that a QoS profile references the statistics profile:

```
host1#show statistics-profile rates1 brief
```

To display a list of QoS profiles that reference the statistics profile:

```
host1#show statistics-profile rates1 references
```

Meaning Table 46 on page 338 lists the **show statistics-profile** command output fields.

Table 46: show statistics-profile Output Fields

Field Name	Field Description
statistics profile	Name of the statistics profile
forwarding rate threshold	Threshold above which forwarded-rate-exceeded events are counted
committed drop threshold	Threshold above which committed-drop-events are counted
conformed drop threshold	Threshold above which conformed-drop-events are counted
exceeded drop threshold	Threshold above which exceeded-drop-events are counted
rate period	Time frame during which statistics are gathered

- Related Topics**
- Configuring Rate Statistics on page 39
 - Configuring Event Statistics on page 40
 - show statistics-profile

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 interface rule   traffic   scheduler   queue
      -----      -
      group          type   type   class      profile      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 47 on page 339 lists the **show qos interface-hierarchy** command output fields.

Table 47: 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

Table 47: show qos interface-hierarchy Output Fields *(continued)*

Field Name	Field Description
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 Topics**
- Configuring a QoS Profile on page 132
 - Attaching a QoS Profile to an Interface on page 134
 - Creating Parameter Instances
 - Configuring Interface Sets for QoS on page 213
 - Configuring Interface Supersets for QoS on page 212
 - show qos interface-hierarchy

Monitoring the Configuration of QoS Port-Type Profiles

Purpose Display information about QoS port-type profiles.

Action To display information about all interface types:

```
host1#show qos-port-type-profile
default-port-profile ethernet qos-profile ethernet-default
default-port-profile atm qos-profile atm-default
default-port-profile serial qos-profile serial-default
default-port-profile server-port qos-profile server-default
default-port-profile lag qos-profile lag-default
```

Meaning Displays a list of all **qos-port-type-profile** commands as they have been entered.

- Related Topics**
- Configuring a QoS Profile on page 132
 - Attaching a QoS Profile to an Interface on page 134
 - Creating Parameter Instances
 - Example: Port-Type QoS Profile Attachment on page 139
 - show qos-port-type-profile

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 48 on page 342 lists the **show qos-profile** command output fields.

Table 48: 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
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

Table 48: show qos-profile Output Fields *(continued)*

Field Name	Field Description
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 Topics**
- Configuring a QoS Profile on page 132
 - Attaching a QoS Profile to an Interface on page 134
 - Creating Parameter Instances
 - show qos-profile

Monitoring the QoS Configuration of ATM Interfaces

Purpose Display ATM port queuing mode and QoS shaping mode status for a specific ATM interface.

Action To display the QoS configuration on an ATM interface:

```

host1# show interfaces atm 2/0
ATM Interface 2/0 is up, line protocol is disabled
AAL5 operational status:      up
    time since last status change: 01:08:32
ATM operational status:      up
    time since last status change: 01:08:32
.
.
.....
InPackets:      0
InBytes:        0
InCells:        0
OutPackets:     7803262
OutBytes:       7803262000
OutCells:       163868502
InErrors:       0
OutErrors:      0
InPacketDiscards: 0
InByteDiscards: 0
InCellErrors:   0
Administrative qos-shaping-mode: frame
Operational qos-shaping-mode: frame
Administrative qos-mode-port: none
Operational qos-mode-port: none
Attached QoS profile: shaping

```

Meaning Table 49 on page 344 lists the related **show interfaces atm** command output fields.

Table 49: show interfaces atm Output Fields

Field Name	Field Description
Administrative qos-mode-port	Per-port queuing mode status: disabled, low-latency, low-cdv, none
Operational qos-mode-port	Per-port queuing mode status: disabled, low-latency, low-cdv, none
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 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 values configured for the qos-shaping-mode command or the qos-port-mode command. For more information, see “Per-Packet Queuing on the SAR Scheduler Overview” on page 163.</p> <ul style="list-style-type: none"> ■ disabled—Shaping mode is configured but disabled. ■ frame—Default shaping mode for shaping and policing rates. Reports 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	QoS profile attachment at or below the displayed interface. For example, if the interface being displayed is a VC, and the attachment is at the ATM AAL5 interface, the ATM AAL5 interface attachment is displayed.

- Related Topics**
- Configuring the QoS Shaping Mode for ATM Interfaces on page 174
 - Configuring a QoS Profile on page 132
 - Attaching a QoS Profile to an Interface on page 134
 - Creating Parameter Instances
 - For more information about other fields displayed with this command, see *JUNOS Link Layer Configuration Guide*
 - `show atm interface`
 - `show interfaces atm`

Monitoring the QoS Configuration of IP Interfaces

Purpose Display the QoS configuration on a particular IP interface.

A dynamic IP interface can have a QoS profile attached by RADIUS. For example, if configured by RADIUS, the **show ip interface** command might show the following:

Attached QoS profile: Strict-qos

However, if the profile is configured statically, the QoS profile is attached to the ATM subinterface, and the attachment is displayed by the **show atm subinterface** command rather than **show ip interface**.

Action To display the QoS configuration for an IP interface:

```
host1#show ip interface atm 2/0.1
ATM2/0.1 line protocol Atm1483 is up, ip is up
.....
Attached QoS profile: test @ ATM2/0
queue 0: traffic class best-effort, bound to ip ATM2/0.1
  Queue length 0 Bytes
  Forwarded packets 0, Bytes 0
  Dropped committed packets 0, Bytes 0
  Dropped conformed packets 0, Bytes 0
  Dropped exceeded packets 0, Bytes 0
  Dropped by WRED committed packets 0, bytes 0
  Dropped by WRED conformed packets 0, bytes 0
  Dropped by WRED exceeded packets 0, bytes 0
  Average queue length 150576 bytes
queue 1: traffic class tc1, bound to ip ATM2/0.1
  Queue length 0 Bytes
  Forwarded packets 0, Bytes 0
  Dropped committed packets 0, Bytes 0
  Dropped conformed packets 0, Bytes 0
  Dropped exceeded packets 0, Bytes 0
  Dropped by WRED committed packets 0, bytes 0
  Dropped by WRED conformed packets 0, bytes 0
  Dropped by WRED exceeded packets 0, bytes 0
  Average queue length 150576 bytes
```

Meaning Table 50 on page 346 lists the related **show ip interface** command output fields.

Table 50: show ip interface Output Fields

Field Name	Field Description
Attached QoS profile	QoS profile attachment at or below the displayed interface. For example, if the interface being displayed is an IP interface, and the attachment is at the VC, the VC interface attachment is displayed.
queue 0	Number of the queue for which statistics are being displayed and whether the queue is under traffic class control
traffic class	Name of traffic class
bound to	Interface to which queue is bound
Queue length	Size of queue in length and bytes
Forwarded	Number of forwarded packets and bytes
Dropped committed	Number of committed packets and bytes dropped
Dropped conformed	Number of conformed packets and bytes dropped
Dropped exceeded	Number of exceeded packets and bytes dropped
Dropped by WRED committed	Number of committed packets and bytes dropped by WRED Displays a cumulative number of committed, conformed, and exceeded packets dropped by WRED for ES2 10G ADV LMs.
Dropped by WRED conformed	Number of conformed packets and bytes dropped by WRED Displays a value of zero for ES2 10G ADV LMs because of the single counter used to calculate packets dropped by WRED functionality (as an aggregate of all colors) for these LMs.
Dropped by WRED exceeded	Number of exceeded packets and bytes dropped by WRED Displays a value of zero for ES2 10G ADV LMs because of the single counter used to calculate packets dropped by WRED functionality (as an aggregate of all colors) for these LMs.
Average queue length	Average length of queue in bytes

- Related Topics**
- Configuring a QoS Profile on page 132
 - Attaching a QoS Profile to an Interface on page 134

- Creating Parameter Instances
- show ip interface

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 51 on page 348 lists the related **show interfaces** command output fields.

Table 51: show interfaces Output Fields

Field Name	Field Description
Administrative qos-shaping-mode	Configured shaping mode for the interface: <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	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 178. <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	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.

- Related Topics**
- Configuring the QoS Shaping Mode for Ethernet Interfaces on page 179
 - Creating Parameter Instances
 - For more information about other fields displayed with this , see *JUNOS Physical Layer Configuration Guide*
 - show interfaces

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 1g0 members
Lag 1g0 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 52 on page 349 lists the related **show interfaces lag members** command output fields.

Table 52: 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 Topics**
- Configuring the Scheduler Hierarchy for Hashed Load Balancing in 802.3ad Link Aggregation Groups on page 191
 - Configuring the Scheduler Hierarchy for Subscriber Load Balancing in 802.3ad Link Aggregation Groups on page 192
 - Creating Parameter Instances
 - For more information about other fields displayed with this command, see the *JUNOS Physical Layer Configuration Guide*
 - **show interfaces lag members**

Monitoring the Configuration of QoS Interface Sets

Purpose Display information about configured interface sets.

Action To display information about a specific interface set:

```
host1#show qos-interface-set vlan-set
interface          member member restricted
set      parent   port    type  count  interface
-----
vlan-set  vlan-ss  lag test1  vlan    2      none
```

To display detailed information about a specific interface set:

```
host1#show qos-interface-set vlan-set detail
interface          member member restricted
set      parent   port    type  count  interface
-----
vlan-set  vlan-ss  lag test1  vlan    2      none
```

```
Children:
  vlan lag test1.1      vlan lag test1.4
```

Meaning Table 53 on page 350 lists the **show qos-interface-set** command output fields.

Table 53: show qos-interface-set Output Fields

Field Name	Field Description
interface set	Name of the interface set
parent	Name of the interface superset that is the parent of the interface set
port	Interface that is the parent of the interface superset
member type	Member-interface type defined for the interface set: <ul style="list-style-type: none"> ■ vlan ■ ip ■ vlan
restricted interface	Restricted interface configured for this interface set
Children	List of interface members associated with the interface superset

- Related Topics**
- Configuring Interface Sets for QoS on page 213
 - show qos-interface-set

Monitoring the Configuration of QoS Interface Supersets

Purpose Display information about configured interface supersets.

Action To display information about a specific interface superset:

```
host1#show qos-interface-superset vlan-ss
interface          member restricted
superset    port    count  interface
-----
vlan-ss      lag test1  1      none
```

To display detailed information about a specific interface superset:

```
host1#show qos-interface-superset vlan-ss detail
interface          member restricted
superset    port    count  interface
-----
vlan-ss      lag test1  1      none

Children:
  set set
```

Meaning Table 54 on page 351 lists the **show qos-interface-superset** command output fields.

Table 54: show qos-interface-superset Output Fields

Field Name	Field Description
interface superset	Name of the interface superset
port	Interface that is the parent of the interface superset
member count	Number of interface members associated with the interface superset
restricted interface	Restricted interface configured for this interface superset
Children	List of interface set members associated with the interface superset

- Related Topics**
- Configuring Interface Supersets for QoS on page 212
 - show qos-interface-superset

Monitoring the AAA Downstream Rate for QoS

Purpose Display whether the QoS downstream rate application is enabled to use downstream rates from the Actual-Data-Rate-Downstream [26-130] DSL Forum VSA.

Action To display the status of the QoS downstream rate application:

```
host1#show aaa qos downstream-rate
Downstream-rate reporting is disabled
```

Meaning Table 55 on page 352 **show aaa qos downstream-rate** command output fields.

Table 55: show aaa qos downstream-rate Output Fields

Field Name	Field Description
Downstream-rate reporting is	Status of the QoS downstream rate application: enabled or disabled

- Related Topics**
- Configuring a Parameter Definition for QoS Downstream Rate on page 303
 - show aaa qos downstream-rate

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
                                     instance
      interface      parameter name      value      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
                                     service
      interface      parameter      value source  manager  persistence
                  name
-----
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 56 on page 354 lists the **show qos-parameter** command output fields.

Table 56: 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
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 Topics**
- Creating Parameter Instances
 - show qos-parameter

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 57 on page 355 lists the **show qos-parameter-define** command output fields.

Table 57: 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
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 Topics**
- Configuring a Basic Parameter Definition for QoS Administrators on page 238
 - show qos-parameter-define

Chapter 32

Troubleshooting QoS

This chapter provides information for troubleshooting QoS.

QoS topics are discussed in the following sections:

- Troubleshooting Memory and Processor Use for Egress Queue Rate Statistics and Events on page 357

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 Topics**
- Monitoring Forwarding and Drop Rates on the Egress Queue on page 333
 - show egress-queue rates

Part 8

Index

- Index on page 361

Index

Symbols

10-Gigabit Ethernet	
monitoring.....	347
10-Gigabit Ethernet forwarding ASIC	
(TFA).....	73, 99, 128, 149
802.3ad link aggregation	
configuring for QoS.....	183
link redundancy and QoS.....	184

A

aaa qos downstream-rate command.....	304
Access Node Control Protocol. <i>See</i> ANCP	
ADSL traffic	
managing cell tax.....	293
ANCP (Access Node Control Protocol)	
shaping downstream rates from.....	301, 303
ASIC scheduler.....	3
assured rate.....	5
assured-rate command.....	238
ATM (Asynchronous Transfer Mode)	
cell shaping.....	164
configuration guidelines for QoS.....	166
configuring for QoS.....	161
configuring the integrated scheduler.....	159
frame shaping.....	164
monitoring.....	343
monitoring for QoS.....	175
SAR shaping.....	166
shaping for QoS.....	164
atm commands	
atm vp-tunnel	168, 172
atm-vp qos-parameter.....	242
atm-vp qos-profile.....	134
ATM modules with relative strict priority.....	63
minimizing latency on the SAR.....	63
oversubscribing.....	63
ATM SAR shaping, QoS.....	166
ATM VP	
interface attachments.....	134
audience for QoS.....	4

B

backpressure.....	160, 166
default integrated mode.....	166
low-cdv mode.....	166
low-latency mode.....	166
best effort.....	5, 13
best-effort queue.....	5
best-effort scheduler node.....	5
buffer-weight command.....	23
burst size, setting in a shaping rate.....	63
byte adjustment application	
configuration example.....	286
byte adjustment applications	
overview.....	293

C

CDV (cell delay variation).....	5
CDVT (cell delay variation tolerance).....	5, 164
cell byte-adjustment application	
configuring.....	297
cell delay variation tolerance. <i>See</i> CDVT	
cell delay variation. <i>See</i> CDV	
clear egress-queue command.....	42
clear fabric-queue command.....	42
clearing statistics.....	41
color-based thresholds.....	18
committed drop threshold.....	37
committed-drop-threshold command.....	40
committed-length command.....	23
committed-threshold command.....	28, 32
compound shared shaping. <i>See</i> shared shaping	
configuration examples	
Diffserv.....	143
QoS parameters.....	243, 263, 273
QoS profiles.....	139
configuring. <i>See</i> specific feature, product, or protocol	
conformed drop threshold.....	37
conformed-drop-threshold command.....	41
conformed-fraction command.....	23
conformed-length command.....	23
conformed-threshold command.....	28, 32
constituents, shared-shaping.....	73
controlled-interface-type command.....	238

controlling subscriber bandwidth	
configuration example.....	243, 263
conventions	
notice icons.....	xxix
text and syntax.....	xxx
convergence-factor command.....	93
customer support.....	xxxi
contacting JTAC.....	xxxi

D

DiffServ	
configuration example.....	143
networks.....	3
documentation set	
comments on.....	xxxi
drop profile.....	25
configuration examples for RED.....	28, 29
configuration examples for WRED.....	33
configuring RED.....	27
configuring WRED.....	31
RED (random early detection).....	26
dynamic shaping of traffic.....	71

E

effective weight.....	6
egress forwarding ASIC (EFA).....	99, 127
Ethernet	
802.3ad link aggregation and QoS.....	183
configuring for QoS.....	178, 183
overview for QoS.....	177
QoS shaping mode.....	178
Ethernet link aggregation commands	
member-interface.....	192
event statistics.....	40
exceeded drop threshold.....	37
exceeded-drop-threshold command.....	41
exceeded-fraction command.....	23
exceeded-length command.....	23
exceeded-threshold command.....	28, 32
expressions.....	48, 235

F

fabric-strict-priority command.....	15
fabric-weight command.....	15
Fast Ethernet	
monitoring.....	347
forwarding classes. <i>See</i> traffic classes	
forwarding rate threshold.....	37
forwarding-rate-threshold command.....	40
frame byte adjustment application	
configuring.....	299
frame forwarding ASIC (FFA).....	100, 127

G

Gigabit Ethernet	
monitoring.....	347
group command.....	132, 238
group node.....	6

H

HAR (hierarchical assured rate).....	6
hierarchical assured rate. <i>See</i> HAR	
hierarchical round-robin. <i>See</i> HRR	
hierarchy, QoS scheduler.....	6, 322
HRR (hierarchical round-robin).....	6, 59
HRR scheduler.....	159, 161
relative strict priority on.....	60, 63

I

implicit constituents	
selection for compound shared shaping.....	109
selection for simple shared shaping.....	109
instance-interface-type command.....	239
integrated scheduler	
configuring for QoS.....	159
interface profile	
attachments.....	134
IP multicast bandwidth adjustment	
configuration example.....	273
configuring.....	271
overview.....	269

L

L2TP (Layer 2 Tunneling Protocol)	
calculating the transmit connect speed.....	203
configuring for QoS.....	199
monitoring for QoS.....	204
overview for QoS.....	197
L2TP sessions	
QoS.....	197
latency.....	6
layer 2 control. <i>See</i> ANCP	
load balancing	
configuring parameters.....	193
hashed	
configuring.....	191
overview.....	183, 186
manged QoS profiles.....	184
subscriber	
configuring.....	192
enabling default configuration.....	191
overview.....	183, 186
load-rebalance command.....	193, 195

M

manuals
 comments on.....xxxi
 maximum-voql command.....94
 minimum-dynamic-rate-percent command.....94
 monitoring. *See* specific feature, product, or protocol
 multiple traffic-class groups.....14
 munged QoS profile
 attachments.....136
 Ethernet link aggregation.....184

N

node
 command.....53, 67, 133, 152, 191, 192, 203, 238
 nodes
 best-effort scheduler.....5
 group.....6
 scheduler.....6
 system resources.....127
 notice icons.....xxix

O

operational QoS shaping mode.....164

P

packet fragmentation, managing.....299
 parameters.....261, 269, 281, 293, 301
 See also QoS parameters
 phantom nodes.....149
 port shaping.....51
 port-type profile, QoS.....6
 attachments.....135
 profile
 drop.....25
 QoS
 attachment.....6
 overview.....127
 port-type.....6
 scheduler.....45
 statistics.....37
 profiles.....17

Q

QoS (quality of service)
 administrators of.....4, 225
 clients of.....4, 225
 description of.....3
 differentiated services
 assured forwarding.....3
 expedited forwarding.....3
 extends Diffserv.....3
 features.....7

 overview.....3
 parameters.....225
 terms.....5
 QoS cell mode application
 configuration example.....286
 configuring.....284
 overview.....281
 QoS commands
 qos-mode-port163
 qos-parameter242
 qos-parameter-define238
 qos-port-type-profile135, 191
 qos-profile
 39, 40, 53, 66, 132, 135, 152, 191, 192, 199
 qos-shaping-mode161, 174, 179
 qos-shared-shaper-control command.....93
 QoS downstream rate application
 configuring.....303
 overview.....301
 relationship with QoS cell mode.....282
 QoS interface sets
 terms.....208
 QoS parameters
 802.3ad link aggregation.....184
 audience.....225
 configuration examples.....243, 263, 273
 configuring for QoS administrators.....229
 configuring for QoS clients.....240
 overview.....225
 parameter definitions
 assured rate.....48, 235
 cell byte adjustment.....293, 297
 configuring.....238
 expressions.....48, 235
 frame byte adjustment.....294, 299
 IP multicast bandwidth
 adjustment.....269, 271
 overview.....229
 QoS cell mode.....281, 284
 QoS downstream rate.....301, 303
 referencing scheduler profiles.....235
 shaping rate.....48, 235
 shared-shaping rate.....48, 235
 parameter instances
 configuring for an interface.....242
 configuring for QoS clients.....240
 configuring globally.....242
 overview.....240
 relationship with other profiles.....227
 terminology.....226
 using with Service Manager.....225, 243
 QoS profile
 ATM VP attachments.....134
 attaching.....134
 configuring.....132
 interface attachments.....134

monitoring.....	155
munged.....	136, 184
munged attachments.....	136
overview.....	127
port-type attachments.....	135
rules illustrated.....	146
S-VLAN attachments.....	135
using with Service Manager.....	132
QoS shaping mode	
configuring.....	174, 179
overview.....	178
qos-profile command.....	238
queue bandwidth.....	59
queue buffers.....	20, 24
queue command.....	67, 133, 153, 191, 192, 199, 238
queue length.....	20, 24
queue profiles.....	17
color-based thresholds.....	18
configuring.....	22
monitoring.....	24
overview.....	17
queue-profile command.....	23
queues.....	6
system resources.....	127

R

random early detection. <i>See</i> RED	
range command.....	239
rate shaping.....	51
QoS.....	6
rate statistics.....	37
rate-period command.....	39, 40
reaction-factor command.....	94
RED (random early detection).....	6, 25
and dynamic queue thresholds.....	34
configuration examples.....	28, 29
configuring.....	27
configuring average queue length.....	28
configuring color blind RED.....	29
configuring colored RED.....	29
how it works.....	26
monitoring.....	36
relative strict-priority scheduling	
configuration example.....	62
configuring.....	60
on ATM modules.....	63
minimizing latency on the SAR.....	63
oversubscribing.....	63
setting burst size in shaping rate.....	63
shaping rate for nonstrict queues.....	63
tuning latency on strict-priority queues.....	66
zero-weight queues.....	63

S

S-VLAN	
interface attachments.....	135
SAR (segmentation and reassembly)	
scheduler.....	159, 161
strict-priority on.....	61
scheduler	
hierarchy.....	6, 46
HRR.....	159
node, best-effort	
best-effort scheduler.....	5
profile.....	45, 235
configuring.....	47
SAR.....	159
scheduler hierarchy	
monitoring.....	322
scheduler map. <i>See</i> QoS profile	
scheduler-profile command.....	48, 199, 238
scheduling	
monitoring.....	123
shadow nodes	
configuration examples	
different traffic-class group.....	154
same traffic-class group.....	154
VLAN and IP queues.....	153
configuring.....	152
interface types.....	132
overview.....	149
system resources.....	151
shadow-node command.....	152
shapeless tunnel.....	171, 173, 174
shaping	
using expressions for.....	48, 235
shaping mode	
configuring.....	174, 179
overview.....	178
shaping rate	
for nonstrict queues.....	63
setting burst size in.....	63
shaping, QoS ATM.....	164
cell.....	164
frame.....	164
shaping-rate command.....	238
shared shaping	
active constituents.....	73
compound.....	99
active constituents.....	107
configuration.....	100
configuration example, VC shared	
shaping.....	102
configuration example, VP shared	
shaping.....	104
configuration limitations.....	74
hardware dependency.....	99
considerations.....	74

- constituents.....73
 - active.....73
 - comparison of explicit and implicit.....107
 - inactive.....107
 - explicit constituents
 - example.....115
 - example of weighted.....117
 - selection.....107, 115
 - implicit constituents
 - example at best-effort node.....109
 - example at best-effort queue.....110
 - example for mixed interface types.....111
 - selection.....107
 - selection for compound.....109
 - selection for simple.....109
 - inactive constituents.....107
 - individual shaping and.....74
 - limiting bandwidth.....72
 - low-CDV mode.....75
 - on the SAR, limitations of.....75
 - overview.....71, 72
 - simple.....79
 - active constituents.....107
 - configuration.....81
 - configuration example, Ethernet.....86
 - configuration example, VC shared
 - shaping.....83
 - configuration example, VP shared
 - shaping.....85
 - controlling the algorithm.....89
 - example, basic.....79
 - example, on best-effort scheduler node.....79
 - traffic starvation.....76
 - types, simple versus compound.....72
 - shared-shaping-constituent command.....120
 - shared-shaping-rate command.....82, 101, 238
 - show commands
 - show aaa qos downstream-rate.....352
 - show atm interface.....343
 - show drop-profile.....320
 - show egress-queue events331
 - show egress-queue rates333
 - show fabric-queue337
 - show interfaces gigabitEthernet.....347
 - show interfaces lag members.....349
 - show interfaces tenGigabitEthernet.....347
 - show qos-parameter.....352
 - show qos-parameter-define355
 - show qos-port-type profile340
 - show qos-profile.....341
 - show qos-shared-shaper-control command.....330
 - show queue-profile319
 - show scheduler-profile327
 - show statistics-profile338
 - show traffic-class.....314
 - show traffic-class-group315
 - show qos commands
 - show qos interface-hierarchy.....339
 - show qos queue-thresholds316
 - show qos scheduler-hierarchy.....322
 - show qos shared-shaper329
 - simple shared shaping. *See* shared shaping
 - statistics
 - ATM.....164
 - statistics profile.....37
 - clearing.....41
 - committed drop threshold.....37
 - conformed drop threshold.....37
 - event statistics.....37
 - exceeded drop threshold.....37
 - failover mode.....357
 - forwarding rate threshold.....37
 - maximum.....37
 - monitoring.....42
 - overview.....37
 - rate period.....37, 39, 40
 - rate statistics.....37
 - resource use.....357
 - thresholds.....40
 - statistics-profile command.....39
 - strict-priority command.....67
 - strict-priority scheduling.....59, 61
 - true versus relative.....61
 - See also* relative strict-priority scheduling
 - subscriber-interface-type command.....239
 - support, technical *See* technical support
 - svlan commands
 - svlan qos-parameter.....243
 - svlan qos-profile.....135
- T**
- TCP friendly.....51
 - technical support
 - contacting JTAC.....xxxi
 - text and syntax conventions.....xxx
 - traffic classes
 - configuring.....14
 - monitoring.....16
 - multiple, configuration example.....143
 - overview.....13
 - traffic flow.....4
 - traffic-class command.....14, 16, 199, 238
 - traffic-class groups
 - monitoring.....16
 - multiple.....14
 - overview.....14
 - traffic-class-group command.....15
 - triple play configurations.....102, 245
 - true strict priority scheduling.....61

V

- variables
 - configuring for shared shaping.....89
- VDSL traffic
 - managing packet fragmentation for.....299

W

- weight command.....57, 238
- weight, QoS.....6
- weighted random early detection. *See* WRED
- WRED (weighted random early detection).....7, 25
 - configuration examples.....33
 - configuring.....31
 - different drop behavior for each queue.....33
 - different treatment of colored packets.....33
 - how it works.....27
 - monitoring.....36

Z

- zero-weight queues.....63