



JUNOS® Software

Network Interfaces Configuration Guide

Release 9.6

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, California 94089
USA

408-745-2000

www.juniper.net

Revision 1
Published: 2009-07-13

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

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.

JUNOS® Software Network Interfaces Configuration Guide

Release 9.6

Copyright © 2009, Juniper Networks, Inc.

All rights reserved. Printed in USA.

Writing: Roy Spencer, Mark Barnard, Fran Singer, Stephen Meiers, Donna Ono

Editing: Stella Hackell, Nancy Kurahashi, and Sonia Saruba

Illustration: Faith Bradford and Nathaniel Woodward

Cover Design: Edmonds Design

Revision History

July 2009—R1 JUNOS 9.6

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 NONINFRINGEMENT. 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 This Guide

Ixi

Part 1	Network Interfaces Configuration Statements Overview	
Chapter 1	Network Interfaces Configuration Statements and Hierarchy	3
Part 2	Router Interfaces Configuration Concepts	
Chapter 2	Router Interfaces Overview	31
Chapter 3	Configuring Physical Interface Properties	67
Chapter 4	Configuring Logical Interface Properties	143
Chapter 5	Configuring Protocol Family and Interface Address Properties	169
Chapter 6	Configuring Circuit and Translational Cross-Connects	223
Chapter 7	Tracing Interface Operations	241
Part 3	Configuring Special Router Interfaces	
Chapter 8	Displaying the Internal Ethernet Interface	245
Chapter 9	Configuring Discard Interfaces	249
Chapter 10	Configuring IP Demultiplexing Interfaces	251
Chapter 11	Configuring the Loopback Interface	257
Part 4	Configuring Serial Interfaces	
Chapter 12	Configuring Serial Interfaces	263
Part 5	Configuring ATM Interfaces	
Chapter 13	Configuring ATM Interfaces	281
Chapter 14	Configuring ATM-over-ADSL Interfaces	355
Chapter 15	Configuring ATM-over-SHDSL Interfaces	361
Part 6	Configuring Frame Relay	
Chapter 16	Configuring Frame Relay	371
Part 7	Configuring Channelized Interfaces	
Chapter 17	Channelized Interfaces	385

Chapter 18	Configuring Channelized OC48/STM16 IQE Interfaces	405
Chapter 19	Configuring Channelized OC12/STM4 Interfaces	423
Chapter 20	Configuring Channelized OC3 IQ and IQE Interfaces	455
Chapter 21	Configuring Channelized STM1 Interfaces	465
Chapter 22	Configuring Channelized T3 Interfaces	479
Chapter 23	Configuring Channelized T1 Interfaces	495
Chapter 24	Configuring Channelized E1 Interfaces	501
Chapter 25	Configuring Channelized E1 PRI and T1 PRI Interfaces	509
Part 8	Configuring Circuit Emulation PICs	
Chapter 26	Circuit Emulation PICs Overview	519
Chapter 27	Configuring SAToP Support on Circuit Emulation PICs	523
Chapter 28	Configuring ATM Support on Circuit Emulation PICs	529
Part 9	Configuring E1, E3, T1, and T3 Interfaces	
Chapter 29	Configuring E1 Interfaces	543
Chapter 30	Configuring E3 Interfaces	551
Chapter 31	Configuring T1 Interfaces	559
Chapter 32	Configuring T3 Interfaces	569
Part 10	Configuring Ethernet Interfaces	
Chapter 33	Configuring Ethernet Interfaces	585
Chapter 34	Configuring 802.1Q VLANs	599
Chapter 35	Configuring Aggregated Ethernet Interfaces	623
Chapter 36	Stacking and Rewriting Gigabit Ethernet VLAN Tags	641
Chapter 37	Configuring Layer 2 Bridging Interfaces	663
Chapter 38	Configuring TCC and Layer 2.5 Switching	665
Chapter 39	Configuring Static ARP Table Entries	669
Chapter 40	Configuring Unrestricted Proxy ARP	671
Chapter 41	Configuring MAC Address Validation on Static Ethernet Interfaces	675
Chapter 42	Enabling Passive Monitoring on Ethernet Interfaces	677
Chapter 43	Configuring IEEE 802.1ag OAM Connectivity-Fault Management	679
Chapter 44	Configuring ITU-T Y.1731 Ethernet Service OAM	711
Chapter 45	Configuring IEEE 802.1x Port-Based Network Access Control	741
Chapter 46	Configuring IEEE 802.3ah OAM Link-Fault Management	745
Chapter 47	Configuring VRRP and VRRP for IPv6	753
Chapter 48	Configuring Gigabit Ethernet Accounting and Policing	755
Chapter 49	Configuring Gigabit Ethernet Autonegotiation	767
Chapter 50	Configuring Gigabit Ethernet OTN Options	773
Chapter 51	Configuring the Management Ethernet Interface	775
Chapter 52	Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength	779
Chapter 53	Configuring 10-Gigabit Ethernet Framing	781

Chapter 54	Configuring 10-Gigabit Ethernet Notification of Link Down Alarm	783
Chapter 55	Configuring Point-to-Point Protocol over Ethernet	785
Chapter 56	Configuring Ethernet Ring Protection Switching	799
Chapter 57	Example Ethernet Configurations	813
Part 11	Configuring ISDN Interfaces	
Chapter 58	Configuring ISDN Interfaces	819
Part 12	Configuring SONET Interfaces	
Chapter 59	Configuring SONET/SDH Interfaces	843
Part 13	Interface Configuration Statements	
Chapter 60	Summary of Interface Configuration Statements	889
Part 14	Index	
	Index	1341
	Index of Statements and Commands	1367

Table of Contents

About This Guide Ixi

JUNOS Documentation and Release Notes	Ixi
Objectives	Ixii
Audience	Ixii
Supported Routing Platforms	Ixii
Using the Indexes	Ixiii
Using the Examples in This Manual	Ixiii
Merging a Full Example	Ixiii
Merging a Snippet	Ixiv
Documentation Conventions	Ixiv
Documentation Feedback	Ixvi
Requesting Technical Support	Ixvii

Part 1

Network Interfaces Configuration Statements Overview

Chapter 1

Network Interfaces Configuration Statements and Hierarchy 3

[edit chassis] Hierarchy Level	3
[edit interfaces] Hierarchy Level	4
[edit logical-systems] Hierarchy Level	19
[edit protocols connections] Hierarchy Level	24
[edit protocols dot1x] Hierarchy Level	25
[edit protocols lacp] Hierarchy Level	25
[edit protocols oam] Hierarchy Level	25
[edit protocols ppp] Hierarchy Level	27
[edit protocols protection-group] Hierarchy Level	27
[edit protocols vrrp] Hierarchy Level	27
[edit system processes] Hierarchy Level	27

Part 2**Router Interfaces Configuration Concepts**

Chapter 2**Router Interfaces Overview****31**

Types of Interfaces	32
Permanent Interfaces	32
Management Ethernet Interfaces	32
Internal Ethernet Interfaces	33
Transient Interfaces	35
Services Interfaces	36
Container Interfaces	37
Traditional APS Concept	37
Container Interfaces Concept	37
APS Support for Container-Based Interfaces	38
Autocopy of APS Parameters	38
Interface Encapsulations	39
Interface Descriptors	50
Interface Naming	51
Physical Part of an Interface Name	52
Logical Part of an Interface Name	56
Separators in an Interface Name	56
Channel Part of an Interface Name	56
Interface Naming for a Routing Matrix Based on a TX Matrix Router	57
Interface Naming for a Routing Matrix Based on a TX Matrix Plus Router	59
Chassis Interface Naming	61
Examples: Interface Naming	62
Displaying Interface Configurations	64
Interface and Router Clock Sources	64
Configuring an External Synchronization Interface	65

Chapter 3**Configuring Physical Interface Properties****67**

Physical Interface Configuration Statements	68
Physical Interfaces Properties Statements List	77
Specifying an Aggregated Interface	92
Specifying a USB Modem Interface on J Series Routers	93
Specifying OC768-over-OC192 Mode	95
Adding an Interface Description to the Configuration	96
Example: Adding an Interface Description to the Configuration	96
Configuring the Link Characteristics	97
Configuring the Media MTU	98
Configuring Interface Encapsulation on Physical Interfaces	106
Configuring the Encapsulation on a Physical Interface	106
Encapsulation Capabilities	110
Example: Configuring the Encapsulation on a Physical Interface	111
Configuring the PPP Challenge Handshake Authentication Protocol	112
Assigning an Access Profile to an Interface	113
Configuring a Default CHAP Secret	113

Configuring the Local Name	113
Configuring Passive Mode	114
Example: Configuring the PPP Challenge Handshake Authentication Protocol	114
Configuring the PPP Password Authentication Protocol	114
Configuring the Local Name	116
Configuring the Local Password	116
Configuring Passive Mode	117
Example: Configuring PAP Authentication Protocol	117
Monitoring a PPP Session	118
Tracing Operations of the pppd Process	119
Configuring PPP Address and Control Field Compression	120
Configuring the PPP Protocol Field Compression	121
Configuring the Interface Speed	122
Management Ethernet Interface on M Series and T Series routers	122
Gigabit Ethernet Interfaces on J Series Routers	123
Fast Ethernet Interface	123
Tri-Rate Ethernet Copper Interface	124
SONET/SDH Interface	124
Configuring Keepalives	126
Configuring the Clock Source	128
Configuring the Router as a DCE	128
Configuring Receive and Transmit Leaky Bucket Properties	129
Configuring Accounting for the Physical Interface	130
Applying an Accounting Profile to the Physical Interface	130
Example: Applying an Accounting Profile to the Physical Interface	131
Interface Diagnostics	131
Configuring Loopback Testing	131
Interface Diagnostics	134
Starting and Stopping a BERT Test	136
Example: Configuring Bit Error Rate Testing	137
Tracing Operations of an Individual Router Interface	137
Damping Interface Transitions	138
Configuring Multiservice Physical Interface Properties	138
Enabling or Disabling SNMP Notifications on Physical Interfaces	139
Enabling Unidirectional Traffic Flow on Physical Interfaces	139
Disabling a Physical Interface	140
Example: Disabling a Physical Interface	141

Chapter 4

Configuring Logical Interface Properties

143

Logical Interfaces Configuration Statements	144
Logical Interfaces Statements List	147
Specifying the Logical Interface Number	155
Configuring Logical System Interface Properties	155
Example: Configuring Logical System Interface Properties	156
Adding a Logical Unit Description to the Configuration	156
Configuring a Point-to-Point Connection	157
Configuring a Multipoint Connection	157

Configuring Accounting for the Logical Interface	157
Applying an Accounting Profile to the Logical Interface	158
Example: Applying an Accounting Profile to the Logical Interface	158
Configuring the Interface Bandwidth	159
Enabling or Disabling SNMP Notifications on Logical Interfaces	159
Configuring Interface Encapsulation on Logical Interfaces	160
Configuring the Encapsulation on a Logical Interface	160
Configuring the LCP Configure-Request Maximum Sent	161
Configuring the NCP Configure-Request Maximum Sent	161
Configuring the PPP Restart Timers	162
Configuring the PPP Clear Loop Detected Timer	162
Configuring Dynamic Profiles for PPP	163
Configuring PPP CHAP Authentication	163
Configuring PPP PAP Authentication	164
Configuring a Default PAP Password	165
Configuring the Local Name	165
Configuring the Local Password	165
Configuring Passive Mode	166
Configuring Dynamic Call Admission Control	166
Example: Configuring Dynamic CAC	167
Disabling a Logical Interface	167

Chapter 5

Configuring Protocol Family and Interface Address Properties 169

Protocol Family Configuration and Interface Address Statements	169
Configuring the Protocol Family	172
IPv6 Overview	174
IPv4-to-IPv6 Transition	174
VRRP Properties	174
Configuring the Interface Address	174
Configuring an Interface IPv4 Address	176
Configuring the Interface IPv6 Address	176
Configuring ICCP for MC-LAG	177
Configuring IPCP Options	177
Configuring an IP Address for an Interface	178
Negotiating an IP Address Assignment from the Remote End	178
Configuring an Interface to Be Unnumbered	179
Assigning a Destination Profile to the Remote End	179
Configuring LLC2 Options	180
Configuring LLC2 Properties	180
Configuring DLSw Ethernet Redundancy Using LLC2 Properties	181
Example: Configuring LLC Options on an Interface	183
Example: Configuring DLSw Ethernet Redundancy	184

Configuring an Unnumbered Interface	185
Configuring an Unnumbered Point-to-Point Interface	185
Example: Configuring an Unnumbered Point-to-Point Interface	186
Configuring an Unnumbered Ethernet or Demux Interface	186
Configuring a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces	187
Configuring Static Routes on Unnumbered Ethernet Interfaces	188
Restrictions for Configuring Unnumbered Ethernet Interfaces	189
Example: Configuring an Unnumbered Ethernet Interface	189
Example: Configuring the Preferred Source Address for an Unnumbered Ethernet Interface	190
Example: Configuring an Unnumbered Ethernet Interface as the Next Hop for a Static Route	190
Setting the Protocol MTU	191
Disabling the Removal of Address and Control Bytes	192
Disabling the Transmission of Redirect Messages on an Interface	192
Configuring Default, Primary, and Preferred Addresses and Interfaces	192
Configuring the Primary Interface for the router	193
Configuring the Primary Address for an Interface	193
Configuring the Preferred Address for an Interface	194
Applying Policers	194
Applying Aggregate Policers	195
Example: Applying Aggregate Policers	196
Applying Hierarchical Policers on Enhanced Intelligent Queuing PICs	197
Hierarchical Policer Overview	198
Hierarchical Policing Characteristics	199
Configuring Hierarchical Policers	200
Configuring a Single-Rate Two-Color Policer	200
Configuring a Single-Rate Tricolor Policer	201
Configuring a Two-Rate Tricolor Marker Policer	202
Applying a Filter to an Interface	203
Defining Interface Groups in Firewall Filters	205
Filter-Based Forwarding on the Output Interface	206
Example: Applying a Filter to an Interface	206
Configuring Unicast RPF	208
Configuring Unicast RPF Strict Mode	209
Configuring Unicast RPF Loose Mode	210
Unicast RPF and Default Routes	210
Unicast RPF Behavior with a Default Route	211
Unicast RPF Behavior Without a Default Route	211
Unicast RPF with Routing Asymmetry	212
Configuring Unicast RPF on a VPN	212
Example: Configuring Unicast RPF on a VPN	213
Example: Configuring Unicast RPF	213
Enabling Source Class and Destination Class Usage	214
Examples: Enabling Source Class and Destination Class Usage	217

Chapter 6	Configuring Circuit and Translational Cross-Connects	223
	Circuit and Translational Cross-Connects Overview	223
	Defining the Encapsulation for Switching Cross-Connects	225
	Configuring PPP or Cisco HDLC Circuits	225
	Configuring ATM Circuits	225
	Configuring Frame Relay Circuits	226
	Configuring Ethernet CCC Circuits	227
	Configuring Ethernet VLAN Circuits	228
	Defining the Connection for Switching Cross-Connects	228
	Configuring MPLS for Switching Cross-Connects	229
	Configuring IS-IS or MPLS Traffic for TCC Interfaces	229
	Configuring ATM-to-Ethernet Interworking	229
	Enabling ATM-to-Ethernet Interworking	230
	Configuring the ATM-to-Ethernet Interworking Ethernet Interface	230
	Configuring the ATM-to-Ethernet Interworking Ethernet Encapsulation	231
	Configuring the ATM-to-Ethernet Interworking Outer VLAN Identifier	231
	Configuring the ATM-to-Ethernet Interworking Inner VLAN Identifier Range	231
	Configuring the ATM-to-Ethernet Interworking Physical Interface VPI	232
	Configuring the ATM-to-Ethernet Interworking ATM Logical Interface	232
	Configuring the ATM-to-Ethernet Interworking Protocol Family	232
	Configuring the ATM-to-Ethernet Interworking Logical Interface VPI	233
	Configuring the ATM-to-Ethernet Interworking Logical Interface VCI	233
	Examples: Configuring Switching Cross-Connects	233
	Example: Configuring a CCC over Frame Relay Encapsulated Interface	233
	Example: Configuring a TCC	234
	Example: Configuring CCC over Aggregated Ethernet	236
	Example: Configuring a Remote LSP CCC over Aggregated Ethernet	237
	Example: Configuring ATM-to-Ethernet Interworking	239
Chapter 7	Tracing Interface Operations	241
	Tracing Operations of an Individual Router Interface	241
	Tracing Operations of the Interface Process	241

Part 3	Configuring Special Router Interfaces	
Chapter 8	Displaying the Internal Ethernet Interface	245
	Displaying the Internal Ethernet Interface for M Series, MX Series, and Most T Series Routers	245
	Displaying Internal Ethernet Interfaces for a Routing Matrix with a TX Matrix Plus Router	246
Chapter 9	Configuring Discard Interfaces	249
Chapter 10	Configuring IP Demultiplexing Interfaces	251
	Configuring an IP Demultiplexing Interface	251
	Configuring an IP Demux Underlying Interface	252
	Specifying the Demux Underlying Interface	253
	Configuring IP Demux Prefixes	253
	Configuring MAC Address Validation on Static Demux Interfaces	254
	Example: Configuring a Demux Interface	255
Chapter 11	Configuring the Loopback Interface	257
	Configuring the Loopback Interface	257
	Example: Configuring the Loopback Interface	258
Part 4	Configuring Serial Interfaces	
Chapter 12	Configuring Serial Interfaces	263
	Serial Interfaces Overview	263
	Physical Interface Configuration Statements for Serial Interfaces	264
	Configuring the Serial Line Protocol	265
	Serial Interface Default Settings	265
	EIA-530 Interface Default Settings	266
	V.35 Interface Default Settings	266
	X.21 Interface Default Settings	267
	Invalid Serial Interface Statements	267
	Invalid EIA-530 Interface Statements	267
	Invalid V.35 interface Statements	268
	Invalid X.21 Interface Statements	268
	Configuring the Serial Clocking Mode	269
	Inverting the Serial Interface Transmit Clock	270
	Configuring the DTE Clock Rate	270
	Configuring the Serial Idle Cycle Flag	271
	Configuring the Serial Signal Handling	271

Configuring the Serial DTR Circuit	274
Configuring Serial Signal Polarities	274
Configuring Serial Loopback Capability	275
Example: Configuring Serial Loopback Capability	276
Configuring Serial Line Encoding	277

Part 5

Configuring ATM Interfaces

Chapter 13

Configuring ATM Interfaces 281

ATM Interfaces Overview	282
ATM1 Physical and Logical Configuration Statement Hierarchies	283
ATM2 IQ Physical and Logical Configuration Statement Hierarchies	285
Supported Features on ATM1 and ATM2 IQ Interfaces	287
Configuring Communication with Directly Attached ATM Switches and Routers	291
Example: Configuring Communication with Directly Attached ATM Switches and Routers	292
Enabling ILMI for Cell Relay	292
Example: Enabling ILMI for Cell Relay	293
Enabling Passive Monitoring on ATM Interfaces	293
Removing MPLS Labels from Incoming Packets	294
Configuring the ATM PIC Type	295
Example: Configuring the ATM PIC Type	296
Configuring ATM Cell-Relay Promiscuous Mode	296
Examples: Configuring ATM Cell-Relay Promiscuous Mode	297
Configuring the Maximum Number of ATM1 VCs on a VP	300
Configuring Layer 2 Circuit Transport Mode	300
Examples: Configuring IQ Layer 2 Circuit Transport Mode	303
Configuring Layer 2 Circuit Cell-Relay Promiscuous Mode	308
Example: Configuring Layer 2 Circuit Cell-Relay Promiscuous Mode	308
Configuring Layer 2 Circuit Trunk Mode Scheduling	309
Example: Configuring Layer 2 Circuit Trunk Mode Scheduling	310
Configuring CoS Queues in Layer 2 Circuit Trunk Mode	311
Example: Configuring CoS Queues in Layer 2 Circuit Trunk Mode	313
Configuring the Layer 2 Circuit Cell-Relay Cell Maximum	313
Class-Based Cell Bundling	314
Configuring the OAM F4 Cell Flows	315
Defining Virtual Path Tunnels	316
Configuring a Point-to-Point ATM1 or ATM2 IQ Connection	316
Configuring a Point-to-Multipoint ATM1 or ATM2 IQ Connection	317
Configuring a Multicast-Capable ATM1 or ATM2 IQ Connection	318
Configuring Inverse ATM1 or ATM2 ARP	318
Defining the ATM Traffic-Shaping Profile	319
Configuring ATM CBR	320
Configuring ATM2 IQ Real-Time VBR	321
Configuring ATM VBR	321

Specifying ATM1 Shaping Values	322
Example: Specifying ATM1 Shaping Values	324
Specifying ATM2 IQ Shaping Values	325
Configuring the ATM1 Queue Length	325
Configuring the ATM2 IQ EPD Threshold	326
Example: Configuring the ATM2 IQ EPD Threshold	328
Configuring Two EPD Thresholds per Queue	328
Configuring the ATM2 IQ Transmission Weight	329
Defining the ATM OAM F5 Loopback Cell Period	329
Configuring the ATM OAM F5 Loopback Cell Threshold	329
Configuring ATM Interface Encapsulation	330
Configuring an ATM1 Cell-Relay Circuit	332
Example: Configuring an ATM1 Cell-Relay Circuit	333
Configuring PPP over ATM2 Encapsulation	334
Example: Configuring PPP over ATM2 IQ Encapsulation	335
Configuring E3 and T3 Parameters on ATM Interfaces	337
Configuring SONET/SDH Parameters on ATM Interfaces	338
Configuring ATM2 IQ VC Tunnel CoS Components	339
Configuring Linear RED Profiles	340
Configuring an ATM Scheduler Map	341
Enabling Eight Queues on ATM2 IQ Interfaces	342
Example: Enabling Eight Queues on T Series, M120, and M320 Platforms	343
Configuring VC CoS Mode	348
Enabling the PLP Setting to Be Copied to the CLP Bit	348
Configuring ATM CoS on the Logical Interface	349
Example: Configuring ATM2 IQ VC Tunnel CoS Components	349
Example: Configuring ATM1 Interfaces	350
Example: Configuring ATM2 IQ Interfaces	352

Chapter 14

Configuring ATM-over-ADSL Interfaces 355

ATM-over-ADSL Overview	355
Configuring Physical ATM Interfaces and Logical Interface Properties for ADSL	356
Configuring the ATM-over-ADSL Virtual Path Identifier	356
Configuring the ATM-over-ADSL Physical Interface Operating Mode	357
Configuring the ATM-over-ADSL Physical Interface Encapsulation Type	358
Configuring the ATM-over-ADSL Logical Interface Encapsulation Type	358
Configuring the ATM-over-ADSL Protocol Family	359
Configuring the ATM-over-ADSL Virtual Channel Identifier	360

Chapter 15

Configuring ATM-over-SHDSL Interfaces 361

ATM-over-SHDSL Overview	361
Configuring ATM Mode for SHDSL Overview	362
Configuring ATM Mode on the PIM	363
Configuring SHDSL Operating Mode on an ATM Physical Interface	364
Configuring Encapsulation on the ATM Physical Interface	364
Configuring Logical Interface Properties	365

Example: Configuring an ATM-over-SHDSL Interface	366
Verifying an ATM-over-SHDSL Interface Configuration	367

Part 6 Configuring Frame Relay

Chapter 16	Configuring Frame Relay	371
	Frame Relay Overview	371
	Configuring Frame Relay Interface Encapsulation	372
	Configuring the Frame Relay Encapsulation on a Physical Interface	372
	Example: Configuring the Encapsulation on a Physical Interface	374
	Configuring the Frame Relay Encapsulation on a Logical Interface	375
	Configuring Frame Relay Control Bit Translation	375
	Configuring the Media MTU on Frame Relay Interfaces	377
	Setting the Protocol MTU with Frame Relay Encapsulation	377
	Configuring Frame Relay Keepalives	378
	Configuring Tunable Keepalives for Frame Relay LMI	378
	Configuring Inverse Frame Relay ARP	379
	Configuring the Router as a DCE with Frame Relay Encapsulation	380
	Configuring Frame Relay DLCIs	380
	Configuring a Point-to-Point Frame Relay Connection	380
	Configuring a Point-to-Multipoint Frame Relay Connection	381
	Configuring a Multicast-Capable Frame Relay Connection	381

Part 7 Configuring Channelized Interfaces

Chapter 17	Channelized Interfaces	385
	Channelized Interfaces Overview	385
	Channelized Interface Capabilities	386
	Data-Link Connection Identifiers on Channelized Interfaces	388
	Clock Sources on Channelized Interfaces	390
	Channelized E1 and T1 PIM Properties	393
	Channelized IQ and IQE Interfaces Properties	393
	Structure of Channelized IQ and Channelized IQE PICs	396
Chapter 18	Configuring Channelized OC48/STM16 IQE Interfaces	405
	Channelized OC48/STM16 IQE Interfaces Overview	405
	Configuring Channelized OC48/STM16 IQE Interfaces in SONET Mode	407
	Configuring OC12 Interfaces	407
	Example: Configuring OC12 Interfaces	408
	Configuring OC3 Interfaces	408
	Example: Configuring OC3 Interfaces	409

Configuring T3 Interfaces	409
Example: Configuring T3 Interfaces	410
Configuring T1 Interfaces	410
Example: Configuring T1 Interfaces	411
Configuring Fractional T1 Interfaces	412
Example: Configuring Fractional T1 Interfaces	412
Configuring NxDS0 Interfaces	413
Example: Configuring NxDS0 Interfaces	414
Configuring Channelized OC48/STM16 IQE Interfaces (SDH Mode)	415
Configuring a Channelized OC48/STM16 IQE PIC for SDH Mode	415
Configuring Clear Channel STM1 and STM4 Interfaces	416
Configuring Channelized AU-4 Interfaces	416
Example: Configuring Channelized AU-4 Interfaces	416
Configuring E3 Interfaces	417
Example: Configuring E3 Interfaces	417
Configuring E1 or Channelized E1 Interfaces	418
Example: Configuring E1 and Channelized E1 Interfaces	418
Configuring NxDS0 IQE Interfaces	418
Example: Configuring NxDS0 IQE Interfaces	419
Configuring Link PIC Failover on Channelized OC48/STM16 IQE Interfaces	419
Example: Configuring Channelized OC48 Interfaces with Partitioned Channels	419

Chapter 19

Configuring Channelized OC12/STM4 Interfaces **423**

Channelized OC12/STM4 IQ and IQE Interfaces Overview	423
Channelization of OC12/STM4 IQ and Channelized OC12/STM4 IQE PICs (SONET Mode)	424
Channelization of OC12/STM4 IQE PIC (SDH Mode)	425
Channelization of OC12/STM4 IQ PIC (SDH Mode)	425
Channelization of OC12 PIC (SONET Mode)	426
Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode)	427
Configuring an OC12/STM4 Interface	427
Configuring T3 Interfaces	427
Example: Configuring T3 Interfaces	428
Configuring OC3 Interfaces	429
Example: Configuring OC3 Interfaces	429
Configuring T1 Interfaces	429
Example: Configuring T1 Interfaces	431

Configuring NxDS0 Interfaces	431
Example: Configuring NxDS0 Interfaces	433
Configuring Fractional T1 Interfaces	433
Example: Configuring Fractional T1 Interfaces	434
Configuring Channelized OC12/STM4 IQE Interfaces (SDH Mode)	434
Configuring Channelized OC12/STM4 IQE PICs for SDH Mode	434
Configuring an Unpartitioned SDH (VC-4-4C) Interface on a Channelized OC12/STM4 IQE PIC	435
Example: Configuring an Unpartitioned SDH (VC-4-4C) Interface	435
Configuring SDH (VC-4) Interfaces on Channelized OC12/STM4 IQE PICs	436
Example: Configuring SDH (VC-4) Interfaces	436
Configuring Channelized AU-4 Interfaces	436
Example: Configuring Channelized AU-4 Interfaces	437
Configuring E3 Interfaces	437
Example: Configuring E3 Interfaces	438
Configuring E1 or Channelized E1 Interfaces	438
Example: Configuring E1 or Channelized CE1 Interfaces	439
Configuring NxDS0 Interfaces on Channelized OC12/STM4 IQE PICs	439
Example: Configuring NxDS0 Interfaces	439
Configuring Channelized OC12/STM4 IQ Interfaces (SDH Mode)	440
Configuring Channelized OC12/STM4 IQ PICs for SDH Mode	440
Configuring an Unpartitioned SDH (VC-4-4C) Interface on a Channelized OC12/STM4 IQ PIC	441
Example: Configuring an Unpartitioned SDH (VC-4-4C) Interface	441
Configuring SDH (VC-4) Interfaces on Channelized OC12/STM4 IQ PICs	441
Example: Configuring SDH (VC-4) Interfaces	442
Configuring Channelized AU-4 Interfaces	442
Example: Configuring Channelized AU-4 Interfaces	442
Configuring T3 or Channelized T3 Interfaces	443
Example: Configuring T3 or Channelized T3 Interfaces	443
Configuring T1 or Channelized T1 Interfaces	443
Example: Configuring T1 or Channelized T1 Interfaces	444
Configuring NxDS0 Interfaces on Channelized OC12/STM4 IQ PICs	444
Example: Configuring NxDS0 Interfaces	444
Configuring Channelized OC12 Interfaces	445
Example: Configuring Channelized OC12 Interfaces	446
Configuring Link PIC Failover on Channelized OC12/STM4 IQ and IQE Interfaces	447
Example: Configuring a Channelized OC12 IQ Interface as an Unpartitioned, Clear Channel	448
Example: Configuring Channelized OC12 Interfaces with Partitioned Channels	451

Chapter 20	Configuring Channelized OC3 IQ and IQE Interfaces	455
	Channelized OC3 IQ and IQE Overview	455
	Partitions, OC Slices, Interface Types, and Time Slots	456
	Configuring a Clear Channel on Channelized OC3 IQ and IQE PICs	457
	Configuring T3 IQ Interfaces	457
	Example: Configuring T3 Interfaces	458
	Configuring T1 and NxDS0 Interfaces	458
	Example: Configuring T1 and NxDS0 Interfaces	460
	Example: Setting Remote Loopback and Running BERT Tests on NxDS0 Interfaces	461
	Configuring Fractional T1 IQ Interfaces	462
	Example: Configuring Fractional T1 IQ Interfaces	462
	Configuring Link PIC Failover on Channelized OC3 IQ and IQE Interfaces	462
Chapter 21	Configuring Channelized STM1 Interfaces	465
	Configuring Channelized STM1 IQ and IQE Interfaces	465
	Configuring an STM1 IQ or STM1 IQE Interface	465
	Configuring E1 IQ and IQE Interfaces	466
	Example: Configuring E1 IQ and IQE Interfaces	466
	Configuring Fractional E1 IQ and IQE Interfaces	467
	Example: Configuring Fractional E1 Interfaces	468
	Configuring an NxDS0 IQ Interface	468
	Example: Configuring an NxDS0 IQ Interface	469
	Example: Configuring Channelized STM1 IQ and IQE Interfaces	469
	Configuring Channelized STM1 Interfaces	471
	Configuring Channelized STM1 Interface Properties	471
	Configuring Virtual Tributary Mapping of Channelized STM1 Interfaces	472
	Configuring Link PIC Failover on Channelized STM1 Interfaces	475
	Example: Configuring Channelized STM1 Interfaces	475
Chapter 22	Configuring Channelized T3 Interfaces	479
	Configuring Channelized T3 IQ Interfaces	479
	Configuring T3 IQ Interfaces	479
	Configuring T1 IQ Interfaces	480
	Example: Configuring T1 IQ and IQE Interfaces	480

	Configuring Fractional T1 IQ and IQE Interfaces	480
	Example: Configuring Fractional T1 IQ Interfaces	481
	Configuring an NxDS0 IQ Interface	481
	Example: Configuring an NxDS0 IQ Interface	482
	Configuring Channelized DS3-to-DS0 Interfaces	482
	Configuring Channelized DS3-to-DS1 Interfaces	485
	Example: Configuring Channelized T3 IQ Interfaces	486
	Examples: Configuring Channelized DS3-to-DS0 Interfaces	487
	Examples: Configuring Channelized DS3-to-DS1 Interfaces	490
Chapter 23	Configuring Channelized T1 Interfaces	495
	Configuring Channelized T1 IQ and IQE Interfaces	495
	Configuring T1 IQ and IQE Interfaces	495
	Configuring Fractional T1 IQ and IQE Interfaces	496
	Example: Configuring Fractional T1 IQ and IQE Interfaces	496
	Configuring NxDS0 IQ and IQE Interfaces	497
	Example: Configuring an NxDS0 IQ or IQE Interface	497
	Configuring Payload Loopback	497
	Configuring Channelized T1 Interface Properties	499
	Example: Configuring Channelized T1 IQ and IQE Interfaces	499
Chapter 24	Configuring Channelized E1 Interfaces	501
	Configuring Channelized E1 IQ and IQE Interfaces	501
	Configuring E1 IQ and IQE Interfaces	501
	Configuring Fractional E1 IQ and IQE Interfaces	502
	Example: Configuring Fractional E1 IQ and IQE Interfaces	502
	Configuring NxDS0 IQ and IQE Interfaces	502
	Example: Configuring an NxDS0 IQ or IQE Interface	503
	Configuring Channelized E1 Interfaces	503
	Configuring Channelized E1 Interface Properties	505
	Example: Configuring Channelized E1 IQ or IQE Interfaces	505
	Example: Configuring Channelized E1 Interfaces	506
Chapter 25	Configuring Channelized E1 PRI and T1 PRI Interfaces	509
	Channelized E1 PRI and T1 PRI Overview	509
	Configuring a Clear Channel on a Dual-Port Channelized T1-E1 PIM	510
	Configuring a Channelized T1/E1 Interface to Drop and Insert Time Slots	510
	Configuring Primary Rate Interfaces	512
	Allocating B-Channels for Dialout	513
	Configuring PRI Interfaces	513
	Example: Configuring a Channelized T1 Interface as Primary Rate Interface	514

Part 8**Configuring Circuit Emulation PICs****Chapter 26****Circuit Emulation PICs Overview 519**

Mobile Backhaul and Circuit Emulation Overview	519
Mobile Backhaul Application Overview	519
Circuit Emulation PIC Types	520
Four-Port Channelized OC3/STM1 Circuit Emulation PIC	520
Twelve-Port T1/E1 Circuit Emulation PIC	520
Circuit Emulation PICs Clocking Features	520
T1 and E1 Options Exceptions on Circuit Emulation PICs	521
T1 and E1 Options Exceptions on 12-Port T1/E1 Circuit Emulation PICs	521
T1 and E1 Options Exceptions on 4-Port Channelized OC3/STM1 Circuit Emulation PICs	522
Displaying Information About Circuit Emulation PICs	522

Chapter 27**Configuring SAToP Support on Circuit Emulation PICs 523**

Configuring SAToP on 4-port Channelized OC3/STM1 CE PICs	523
Configuring SONET/SDH Framing Mode at the PIC Level	523
Configuring SONET/SDH Framing Mode at the Port Level	524
Configuring COC3 Ports Down to T1 Channels	524
Configuring CSTM1 Ports Down to E1 Channels	524
Configuring SAToP Emulation on T1/E1 Interfaces on CE PICs	525
Setting the Emulation Mode	525
Configuring SAToP Emulation on T1/E1 Interfaces	525
Setting the Encapsulation Mode	526
T1/E1 Loopback Support	526
T1 FDL Support	526
Setting the SAToP Options	526
Pseudowire Interface Configuration	527

Chapter 28**Configuring ATM Support on Circuit Emulation PICs 529**

Overview of ATM Support on Circuit Emulation PICs	529
Configuring the 12-Port Channelized T1/E1 CE PIC Operating Mode	530
T1/E1 Mode Selection	530
12-Port Channelized T1/E1 CE PIC Configuration Statements	531
Configuring the 4-Port Channelized COC3/STM1 CE PIC Operating Mode	532
T1/E1 Mode Selection	532
Configuring a Port for SONET or SDH Mode on a 4-Port Channelized COC3/STM1 CE PIC	533
Configuring an ATM Interface on a COC1	534
Configuring ATM Pseudowires	534
Cell Relay Mode (atm-l2circuit-mode cell)	535
Configuring VP or Port Promiscuous Mode	535
Configuring AAL5 SDU Mode (atm-l2circuit-mode aal5)	536

ATM OAM	536
VP Pseudowires (CCC Encapsulation)	536
Port Pseudowires (CCC Encapsulation)	537
VC Pseudowires (CCC Encapsulation)	537
Scaling	537
Congestion Control	537
QoS/Shaping	537
Configuring the PIC Type	537
Configuring Layer 2 Circuit and Layer 2 VPN Pseudowires	537
Supported Interface Configurations	538
ATM Limitations	539

Part 9

Configuring E1, E3, T1, and T3 Interfaces

Chapter 29

Configuring E1 Interfaces 543

E1 Interfaces Overview	543
Configuring E1 Physical Interface Properties	544
Configuring E1 BERT Properties	544
Configuring the E1 Frame Checksum	545
Configuring E1 Framing	546
Configuring the E1 Idle Cycle Flag	546
Configuring E1 Data Inversion	546
Configuring E1 Loopback Capability	547
Example: Configuring E1 Loopback Capability	548
Configuring E1 Start and End Flags	548
Configuring Fractional E1 Time Slots	548
Example: Configuring Fractional E1 Time Slots	549

Chapter 30

Configuring E3 Interfaces 551

E3 Interfaces Overview	551
Configuring E3 Physical Interface Properties	552
Configuring E3 BERT Properties	552
Configuring the E3 CSU Compatibility Mode	553
Configuring the E3 Frame Checksum	554
Configuring the E3 Idle Cycle Flag	555
Configuring E3 Data Inversion	555
Configuring E3 Loopback Capability	555
Example: Configuring E3 Loopback Capability	556
Configuring E3 HDLC Payload Scrambling	557
Configuring the E3 Start and End Flags	557
Configuring E3 IQ and IQE Unframed Mode	558

Chapter 31 Configuring T1 Interfaces 559

T1 Interfaces Overview	559
Configuring T1 Physical Interface Properties	560
Configuring T1 BERT Properties	560
Configuring the T1 Buildout	561
Configuring T1 Byte Encoding	561
Configuring T1 CRC Error Major Alarm Thresholds	562
Configuring T1 CRC Error Minor Alarm Thresholds	562
Configuring T1 Data Inversion	563
Configuring the T1 Frame Checksum	563
Configuring the T1 Remote Loopback Response	564
Configuring T1 Framing	564
Configuring T1 Line Encoding	564
Configuring T1 Loopback Capability	565
Configuring the T1 Idle Cycle Flag	566
Configuring T1 Start and End Flags	567
Configuring Fractional T1 Time Slots	567
Example: Configuring Fractional T1 Time Slots	567

Chapter 32 Configuring T3 Interfaces 569

T3 Interfaces Overview	569
Configuring T3 Physical Interface Properties	570
Configuring T3 BERT Properties	570
Disabling T3 C-Bit Parity Mode	571
Configuring the T3 CSU Compatibility Mode	572
Configuring the T3 Frame Checksum	574
Configuring the T3 FEAC Response	575
Configuring the T3 Idle Cycle Flag	575
Configuring the T3 Line Buildout	575
Configuring the Channelized T3 Loop Timing	576
Configuring T3 Loopback Capability	576
Configuring T3 HDLC Payload Scrambling	578
Configuring T3 Start and End Flags	579
Examples: Configuring T3 Interfaces	579

Part 10 Configuring Ethernet Interfaces

Chapter 33 Configuring Ethernet Interfaces 585

Ethernet Interfaces Overview	585
Configuring Ethernet Physical Interface Properties	586
Configuring J Series Services Router Switching Interfaces	589
Example: Configuring J Series Services Router Switching Interfaces	590
MX Series Router Interface Identifiers	591

Enabling Ethernet MAC Address Filtering	591
Filtering Specific MAC Addresses	592
Configuring Ethernet Loopback Capability	593
Configuring Flow Control	594
Ignoring Layer 3 Incomplete Errors	594
Configuring the Link Characteristics on Ethernet Interfaces	595
Configuring Gratuitous ARP	596
Adjusting the ARP Aging Timer	596
Configuring the Interface Speed on Ethernet Interfaces	597
Configuring the Ingress Rate Limit	597
Configuring Weighted Random Early Detection	598

Chapter 34

Configuring 802.1Q VLANs 599

802.1Q VLANs Overview	599
Configuring Dynamic 802.1Q VLANs	600
802.1Q VLAN IDs and Ethernet Interface Types	600
Enabling VLAN Tagging	601
Configuring Single-Tag Framing	602
Configuring Dual Tagging	602
Configuring Mixed Tagging	602
Configuring Mixed Tagging Support for Untagged Packets	603
Example: Configuring Mixed Tagging	603
Example: Configuring Mixed Tagging to Support Untagged Packets	604
Binding VLAN IDs to Logical Interfaces	604
Binding VLAN IDs to Logical Interfaces Overview	604
Binding a VLAN ID to a Logical Interface	605
Binding a VLAN ID to a Single-Tag Logical Interface	605
Binding a VLAN ID to a Dual-Tag Logical Interface	605
Binding a Range of VLAN IDs to a Logical Interface	606
Binding a Range of VLAN IDs to a Single-Tag Logical Interface	606
Binding a Range of VLAN IDs to a Dual-Tag Logical Interface	606
Example: Binding Ranges VLAN IDs to Logical Interfaces	607
Binding a List of VLAN IDs to a Logical Interface	607
Binding a List of VLAN IDs to a Single-Tag Logical Interface	607
Binding a List of VLAN IDs to a Dual-Tag Logical Interface	608
Example: Binding Lists of VLAN IDs to Logical Interfaces	608
Configuring VLAN Encapsulation	609
Example: Configuring VLAN Encapsulation on a Gigabit Ethernet Interface	610
Example: Configuring VLAN Encapsulation on an Aggregated Ethernet Interface	610

Configuring Extended VLAN Encapsulation	610
Example: Configuring Extended VLAN Encapsulation on a Gigabit Ethernet Interface	611
Example: Configuring Extended VLAN Encapsulation on an Aggregated Ethernet Interface	611
Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs	612
Guidelines for Configuring Physical Link-Layer Encapsulation to Support CCCs	612
Guidelines for Configuring Logical Link-Layer Encapsulation to Support CCCs	612
Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface	614
Configuring a VLAN-Bundled Logical Interface	614
Specifying the Interface Over Which VPN Traffic Travels to the CE Router	614
Specifying the Interface to Handle Traffic for a CCC	615
Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface	615
Configuring a VLAN-Bundled Logical Interface	616
Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit	616
Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface	617
Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface	618
Configuring a Logical Interface for Access Mode	619
Example: Configuring a Logical Interface for Access Mode	620
Configuring a Logical Interface for Trunk Mode	620
Configuring the VLAN ID List for a Trunk Interface	620
Configuring a Trunk Interface on a Bridge Network	621

Chapter 35

Configuring Aggregated Ethernet Interfaces

623

Aggregated Ethernet Interfaces Overview	623
Configuring Aggregated Ethernet Interfaces	624
Configuring Ethernet Link Aggregation	625
Configuring Aggregated Ethernet Link Protection	626
Setting the Number of Aggregated Ethernet Interfaces on the Chassis	627
Configuring Aggregated Ethernet LACP	627
Configuring the LACP Interval	628
Configuring LACP Link Protection	629
Enabling LACP Link Protection	629
Configuring LACP System Priority	630
Configuring LACP Port Priority	630
Tracing LACP Operations	631
Example: Configuring Aggregated Ethernet LACP	631
Configuring Tagged Aggregated Ethernet Interfaces	632
Configuring Untagged Aggregated Ethernet Interfaces	633
Example: Configuring Untagged Aggregated Ethernet Interfaces	634
Configuring Aggregated Ethernet Link Speed	634

Configuring Aggregated Ethernet Minimum Links	635
Configuring Hierarchical Scheduler on Aggregated Ethernet Interfaces	635
Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers	636
Overview of Symmetrical Load Balancing on an 802.3ad LAG on MX Series Routers	636
Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers	637
Example Configurations	640
Example Configurations of Chassis Wide Settings	640
Example Configurations of Per PFE Settings	640

Chapter 36 Stacking and Rewriting Gigabit Ethernet VLAN Tags 641

Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview	641
Stacking and Rewriting Gigabit Ethernet VLAN Tags	642
Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames	644
Configuring Stacked VLAN Tagging	645
Configuring Dual VLAN Tags	645
Configuring Inner and Outer TPIDs and VLAN IDs	645
Stacking a VLAN Tag	648
Removing a VLAN Tag	649
Removing the Outer and Inner VLAN Tags	649
Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag	650
Stacking Two VLAN Tags	651
Rewriting the VLAN Tag on Tagged Frames	651
Rewriting a VLAN Tag on Untagged Frames	652
Rewriting a VLAN Tag and Adding a New Tag	655
Rewriting the Inner and Outer VLAN Tags	655
Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags	656

Chapter 37 Configuring Layer 2 Bridging Interfaces 663

Layer 2 Bridging Interfaces Overview	663
Configuring Layer 2 Bridging Interfaces	663
Example: Configuring Layer 2 Bridging Interfaces	664

Chapter 38 Configuring TCC and Layer 2.5 Switching 665

TCC and Layer 2.5 Switching Overview	665
Configuring VLAN TCC Encapsulation	665
Configuring Ethernet TCC	666
Example: Configuring an Ethernet TCC or Extended VLAN TCC	667

Chapter 39	Configuring Static ARP Table Entries	669
	Static ARP Table Entries Overview	669
	Configuring Static ARP Table Entries	669
	Example: Configuring Static ARP Table Entries	670
Chapter 40	Configuring Unrestricted Proxy ARP	671
	Unrestricted Proxy ARP Overview	671
	Configuring Unrestricted Proxy ARP	672
Chapter 41	Configuring MAC Address Validation on Static Ethernet Interfaces	675
	MAC Address Validation on Static Ethernet Interfaces Overview	675
	Configuring MAC Address Validation on Static Ethernet Interfaces	675
	Example of Strict MAC Validation on a Static Ethernet Interface	676
Chapter 42	Enabling Passive Monitoring on Ethernet Interfaces	677
	Passive Monitoring on Ethernet Interfaces Overview	677
	Enabling Passive Monitoring on Ethernet Interfaces	677
Chapter 43	Configuring IEEE 802.1ag OAM Connectivity-Fault Management	679
	IEEE 802.1ag OAM Connectivity Fault Management Overview	680
	Connectivity Fault Management Key Elements	680
	Continuity Check Protocol	681
	Linktrace Protocol	682
	Creating the Maintenance Domain	682
	Configuring the Maintenance Domain Name Format	682
	Configuring the Maintenance Domain Level	682
	Configuring MIP for Bridge Domains of a Virtual Switch	683
	Configuring Maintenance Intermediate Points	683
	Configuring the Maintenance Domain Routing Instances Bridge Domain	684
	Configuring the Maintenance Domain Routing Instance	684
	Configuring the Maintenance Domain Instance	684
	Configuring the Maintenance Domain MIP Half Function	684
	Creating the Maintenance Association	684
	Configuring the Maintenance Association Short Name Format	685
	Configuring the Continuity Check	685
	Configuring the Continuity Check Hold Interval	685
	Configuring the Continuity Check Interval	686
	Configuring the Continuity Check Loss Threshold	686
	Configuring a Maintenance End Point	686

Enabling Maintenance End Point Automatic Discovery	686
Configuring the Maintenance End Point Direction	686
Configuring the Maintenance End Point Interface	687
Configuring the Maintenance End Point Priority	687
Configuring a Remote Maintenance End Point	688
Configuring a Remote Maintenance End Point Action Profile	688
Configuring a Connectivity-Fault Management Action Profile	688
Configuring a CFM Action Profile Action	688
Configuring a CFM Interface Down Action Profile Action	688
Configuring the Linktrace Path Age Timer	689
Configuring the Linktrace Database Size	690
Configuring Ethernet Local Management Interface	690
Ethernet Local Management Interface Overview	690
Configuring the Ethernet Local Management Interface	692
Configuring an OAM Protocol (CFM)	692
Assigning the OAM Protocol to an EVC	692
Enabling E-LMI on an Interface and Mapping CE VLAN IDs to an EVC	692
Example E-LMI Configuration	693
Configuring PE1	694
Configuring PE2	695
Configuring Two UNIs Sharing the Same EVC	697
Configuring Port Status TLV and Interface Status TLV	697
Overview of TLVs	698
Various TLVs for CFM PDUs	698
Support for Additional Optional TLVs	700
Port Status TLV	700
Interface Status TLV	703
MAC Status Defects	705
Configuring Remote MEP Action Profile Support	707
Monitoring a Remote MEP Action Profile	708
Configuring M120 and MX Series Routers for CCC Encapsulated Packets	708
Overview of IEEE 802.1ag CFM OAM Support for CCC Encapsulated Packets	709
CFM Features Supported on Layer 2 VPN Circuits	709
Configuring CFM for CCC Encapsulated Packets	709

Chapter 44

Configuring ITU-T Y.1731 Ethernet Service OAM

711

Ethernet Frame Delay Measurements Overview	711
ITU-T Y.1731 Frame Delay Measurement Feature	711
Ethernet CFM	712
Ethernet Frame Delay Measurement	713
One-Way Ethernet Frame Delay Measurement	713
Two-Way Ethernet Frame Delay Measurement	714
Choosing Between One-Way and Two-Way ETH-DM	715
Restrictions for Ethernet Frame Delay Measurement	716
Guidelines for Configuring Routers to Support an ETH-DM Session	717
Configuration Requirements for ETH-DM	717
Configuration Options for ETH-DM	717

Guidelines for Starting an ETH-DM Session	718
ETH-DM Session Prerequisites	718
ETH-DM Session Parameters	718
Restrictions for an ETH-DM Session	719
Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts	720
ETH-DM Statistics	720
ETH-DM Statistics Retrieval	722
ETH-DM Frame Counts	722
ETH-DM Frame Count Retrieval	723
Configuring Routers to Support an ETH-DM Session	724
Configuring MEP Interfaces	724
Ensuring that Distributed ppmdd Is Not Disabled	725
Enabling the Hardware-Assisted Timestamping Option	726
Starting an ETH-DM Session	727
Using the monitor ethernet delay-measurement Command	727
Starting a One-Way ETH-DM Session	728
Starting a Two-Way ETH-DM Session	728
Managing ETH-DM Statistics and ETH-DM Frame Counts	729
Displaying ETH-DM Statistics Only	729
Displaying ETH-DM Statistics and Frame Counts	730
Displaying ETH-DM Frame Counts for MEPs by Enclosing CFM Entity	730
Displaying ETH-DM Frame Counts for MEPs by Interface or Domain Level	731
Clearing ETH-DM Statistics and Frame Counts	732
Example: One-Way Ethernet Frame Delay Measurement	732
Description of the Example One-Way Frame Delay Measurement	733
Steps for the Example One-Way Frame Delay Measurement	734

Chapter 45**Configuring IEEE 802.1x Port-Based Network Access Control 741**

IEEE 802.1x Port-Based Network Access Control Overview	741
Administrative State of the Authenticator Port	742
Administrative Mode of the Authenticator Port	742
Configuring the Authenticator	742
Viewing the dot1x Configuration	743

Chapter 46**Configuring IEEE 802.3ah OAM Link-Fault Management 745**

IEEE 802.3ah OAM Link-Fault Management Overview	745
Configuring IEEE 802.3ah OAM Link-Fault Management	746
Enabling IEEE 802.3ah OAM Support	746
Configuring Link Discovery	746
Configuring the OAM PDU Interval	747
Configuring the OAM PDU Threshold	747
Configuring Threshold Values for Local Fault Events on an Interface	747
Disabling the Sending of Link Event TLVs	748
Detecting Remote Faults	748
Configuring an OAM Action Profile	748
Specifying the Actions to Be Taken for Link-Fault Management Events	749

	Monitoring the Loss of Link Adjacency	750
	Monitoring Protocol Status	750
	Configuring Threshold Values for Fault Events in an Action Profile	750
	Applying an Action Profile	751
	Setting a Remote Interface into Loopback Mode	751
	Enabling Remote Loopback Support on the Local Interface	751
	Example: Configuring IEEE 802.3ah OAM Support on an Interface	752
Chapter 47	Configuring VRRP and VRRP for IPv6	753
	VRRP and VRRP for IPv6 Overview	753
	Configuring VRRP and VRRP for IPv6	753
Chapter 48	Configuring Gigabit Ethernet Accounting and Policing	755
	Gigabit Ethernet Accounting and Policing Overview	755
	Configuring Gigabit Ethernet Policers	757
	Configuring a Policer	757
	Specifying an Input Priority Map	758
	Specifying an Output Priority Map	759
	Applying a Policer	759
	Configuring MAC Address Filtering	761
	Example: Configuring Gigabit Ethernet Policers	762
	Configuring Gigabit Ethernet Two-Color and Tricolor Policers	763
	Configuring a Policer	764
	Applying a Policer	765
	Example: Configuring and Applying a Policer	765
	Configuring MAC Address Accounting	766
Chapter 49	Configuring Gigabit Ethernet Autonegotiation	767
	Gigabit Ethernet Autonegotiation Overview	767
	Configuring Gigabit Ethernet Autonegotiation	767
	Configuring Gigabit Ethernet Autonegotiation with Remote Fault	767
	Configuring Flow Control	768
	Configuring Autonegotiation Speed on MX Series Routers	768
	Displaying Autonegotiation Status	768
Chapter 50	Configuring Gigabit Ethernet OTN Options	773
	Gigabit Ethernet OTN Options Configuration Overview	773
	Gigabit Ethernet OTN Options	773

Chapter 51	Configuring the Management Ethernet Interface	775
	Management Ethernet Interface Overview	775
	Configuring a Consistent Management IP Address	775
	Configuring the MAC Address on the Management Ethernet Interface	777
Chapter 52	Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength	779
	10-Gigabit Ethernet DWDM Interface Wavelength Overview	779
	Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength	779
Chapter 53	Configuring 10-Gigabit Ethernet Framing	781
	10-Gigabit Ethernet Framing Overview	781
	Configuring 10-Gigabit Ethernet Framing	781
Chapter 54	Configuring 10-Gigabit Ethernet Notification of Link Down Alarm	783
	10-Gigabit Ethernet Notification of Link Down Alarm Overview	783
	Configuring 10-Gigabit Ethernet Notification of Link Down Alarm	783
Chapter 55	Configuring Point-to-Point Protocol over Ethernet	785
	PPPoE Overview	785
	PPPoE Interfaces	786
	Ethernet Interface	786
	PPPoE Stages	786
	PPPoE Discovery Stage	787
	PPPoE Session Stage	787
	Optional CHAP Authentication	788
	Configuring PPPoE	788
	Setting the Appropriate Encapsulation on the PPPoE Interface	789
	Configuring PPPoE Encapsulation on an Ethernet Interface	790
	Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface	790
	Configuring a PPPoE Interface	790
	Configuring the PPPoE Underlying Interface	791
	Identifying the Access Concentrator	791
	Configuring the PPPoE Automatic Reconnect Wait Timer	792
	Configuring the PPPoE Service Name	792
	Configuring the PPPoE Server Mode	793
	Configuring the PPPoE Client Mode	793
	Configuring the PPPoE Source and Destination Addresses	793
	Deriving the PPPoE Source Address From a Specified Interface	794
	Configuring the PPPoE IP Address by Negotiation	794
	Configuring the Protocol MTU PPPoE	794

Example: Configuring a PPPoE Client Interface on a J Series Services Router	795
Example: Configuring a PPPoE Server Interface on an M120 or M320 Router	796
Disabling the Sending of PPPoE Keepalive Messages	796
Verifying a PPPoE Configuration	796

Chapter 56 Configuring Ethernet Ring Protection Switching 799

Ethernet Ring Protection Switching Overview	799
Ethernet Ring Protection Switching Functionality	800
Acronyms	800
Ring Nodes	800
Ring Node States	801
Failure Detection	801
Logical Ring	801
FDB Flush	801
Traffic Blocking and Forwarding	801
RAPS Message Blocking and Forwarding	802
Dedicated Signaling Control Channel	803
RAPS Message Termination	803
Manual Switch	803
Non-Revertive Switch	803
Multiple Rings	803
Node ID	804
Bridge Domains with the Ring Port	804
Configuring Ethernet Ring Protection Switching	804
Ethernet Ring Protection Switching Configuration Example	805
Examples: Ethernet RPS Output	809
Normal Situation	809
Failure Situation	811

Chapter 57 Example Ethernet Configurations 813

Example: Configuring Fast Ethernet Interfaces	813
Example: Configuring Gigabit Ethernet Interfaces	813
Example: Configuring Aggregated Ethernet Interfaces	814
Example: Configuring Aggregated Ethernet Link Protection	815

Part 11 Configuring ISDN Interfaces

Chapter 58 Configuring ISDN Interfaces 819

ISDN Interfaces Overview	819
Configuring ISDN Services Physical and Logical Interface Properties	820
Configuring ISDN Physical Interface Properties	821
Configuring an ISDN Interface to Screen Incoming Calls	823

Configuring ISDN Logical Interface Properties	823
Configuring an ISDN Dialer Interface as a Backup Interface	826
Example: Configuring an ISDN Interface as the Backup Interface	827
Applying the Dial-on-Demand Dialer Filter to the Dialer Interfaces	828
Example: Applying the Dialer Filter	828
Configuring Bandwidth on Demand	829
Configuring the Dialer Interface	830
Configuring the ISDN Interface	831
Example: Configuring Bandwidth on Demand	831
Configuring Dial-In and Callback	832
Configuring Dial-In	833
Disabling Dial-In	833
Configuring Callback	834
Example: Configuring Dial-In and Callback	834
Configuring Dialer Watch	835
Configuring the Dialer Interface	835
Configuring the Physical Interface	836
Example: Configuring Dialer Watch	836
Example: Complete ISDN Called-Calling Router Configuration	837
Disabling ISDN Processes	840

Part 12

Configuring SONET Interfaces

Chapter 59

Configuring SONET/SDH Interfaces 843

SONET/SDH Interfaces Overview	843
Configuring SONET/SDH Physical Interface Properties	844
Configuring SONET/SDH Framing	846
Configuring SONET/SDH Interface Speed	847
Configuring SONET/SDH Header Byte Values	849
Configuring an Incrementing STM ID	850
Configuring the SONET/SDH Frame Checksum	851
Configuring Channelized IQ and IQE SONET/SDH Loop Timing	852
Configuring SONET/SDH Loopback Capability	852
Example: Configuring SONET/SDH Loopback Capability	853
Configuring the SONET/SDH Path Trace Identifier	853
Configuring SONET/SDH HDLC Payload Scrambling	854
Configuring SONET/SDH RFC 2615 Support	855
Configuring SONET/SDH Defect Triggers to Be Ignored	855
Configuring SONET/SDH Defect Hold Times	856
Example: Configuring SONET/SDH Defects to Be Ignored	858
Configuring Virtual Tributary Mapping	858
Configuring APS and MSP	859
Configuring Basic APS Support	861
Configuring Container Interfaces	863
Configuring Switching Between the Working and Protect Circuits	866

Configuring Revertive Mode	867
Configuring Unidirectional Switching Mode Support	867
Configuring APS Timers	868
Configuring Link PIC Redundancy	869
Example: Configuring Link PIC Redundancy	870
Configuring APS Load Sharing Between Circuit Pairs	870
Example: Configuring APS Load Sharing Between Circuit Pairs	872
Configuring SONET Options for 10-Gigabit Ethernet Interfaces	872
Configuring the Media MTU on SONET/SDH Interfaces	873
Enabling Passive Monitoring on SONET/SDH Interfaces	874
Removing MPLS Labels from Incoming Packets	874
Configuring the Clock Source on SONET/SDH Interfaces	875
Configuring Receive and Transmit Leaky Bucket Properties on SONET/SDH Interfaces	876
Damping Interface Transitions on SONET/SDH Interfaces	877
Configuring Interface Encapsulation on SONET/SDH Interfaces	878
Configuring the Encapsulation on a Physical SONET/SDH Interface	878
Example: Configuring the Encapsulation on a Physical SONET/SDH Interface	880
Configuring the Encapsulation on a Logical SONET/SDH Interface	880
Example: Configuring SONET/SDH Interfaces	881
Configuring Aggregated SONET/SDH Interfaces	881
Configuring SONET/SDH Link Aggregation	882
Configuring Aggregated SONET/SDH Link Speed	883
Configuring Aggregated SONET/SDH Minimum Links	883
Configuring Filters or Sampling on Aggregated SONET/SDH Links	884
Examples: Configuring Filters or Sampling on Aggregated SONET/SDH Links	884
Example: Configuring Aggregated SONET/SDH Interfaces	885

Part 13

Interface Configuration Statements

Chapter 60

Summary of Interface Configuration Statements	889
802.3ad	889
accept	890
accept-source-mac	891
access-concentrator	892
access-profile	893
accounting	894
accounting-profile	894
acfc	895
ack-delay-time	895
ack-max	896
acknowledge-retries	896

acknowledge-timer	897
action	898
action (OAM)	898
action (Policer)	898
action-profile	899
action-profile (Applying to CFM)	899
action-profile (Defining for CFM)	899
action-profile (Defining for LFM)	900
action-red-differential-delay	901
activation-delay	901
activation-priority	902
address	903
advertise-interval	905
advertise-interval (APS)	905
advertise-interval (DLSw)	906
age	906
aggregate	907
aggregate (Gigabit Ethernet CoS Policer)	907
aggregate (Hierarchical Policer)	908
aggregate (SONET/SDH)	908
aggregate-ports	909
aggregated-ether-options	910
aggregated-sonet-options	911
allow-any-vci	911
allow-fragmentation	912
allow-remote-loopback	912
annex	913
apply-action-profile	913
aps	914
arp	915
asynchronous-notification	916
atm-encapsulation	916
atm-options	917
atm-scheduler-map	918
authentication-key	918
authentication-profile-name	919
authenticator	920
auto-configure	921
auto-discovery	921
auto-negotiation	922
auto-negotiation (Gigabit Ethernet)	922
auto-negotiation (J Series uPIM)	923
auto-reconnect	923
backup-destination	924
backup-interface	924
backup-options	925
bandwidth	925
bandwidth-limit	926
bandwidth-limit (Hierarchical Policer)	926
bandwidth-limit (Policer for Gigabit Ethernet Interfaces)	927
bchannel-allocation	927

bearer-bandwidth-limit	928
bert-algorithm	929
bert-error-rate	931
bert-period	932
bridge-domain	933
broadcast	933
buildout	934
buildout (E3 or T3 over ATM Interfaces)	934
buildout (T1 Interfaces)	935
bundle	936
burst-size-limit	937
burst-size-limit (Hierarchical Policer)	937
burst-size-limit (Policer for Gigabit Ethernet Interfaces)	937
byte-encoding	938
bytes	939
callback	940
callback-wait-period	941
caller	942
calling-number	943
cbit-parity	943
cbr	944
cell-bundle-size	945
chap	946
chap-secret	947
cisco-interoperability	947
classifier	948
clear-dont-fragment-bit	948
client	949
clocking	950
clocking-mode	951
clock-rate	952
compatibility-mode	953
compression	954
compression (PPP Properties)	954
compression (Voice Services)	955
compression-device	955
connections	956
connectivity-fault-management	957
container-devices	958
container-list	958
container-options	959
container-type	959
continuity-check	960
control-channel	961
control-polarity	961
control-signal	962
copy-tos-to-outer-ip-header	962
core-dump	963
crc-major-alarm-threshold	963
crc-minor-alarm-threshold	964
cts	964

cts-polarity	965
current	965
data-input	966
dcd	967
dcd-polarity	967
dce	968
dce-options	968
deactivation-delay	969
default-action	969
default-chap-secret	970
default-pap-password	970
demux0	971
demux-destination	972
demux-destination (Underlying Interface)	972
demux-destination (Demux Interface)	973
demux-options	973
demux-source	974
demux-source (Underlying Interface)	974
demux-source (Demux Interface)	975
description	976
destination	977
destination (DLSw)	977
destination (IPCP)	978
destination (Routing Instance)	978
destination (Tunnels)	979
destination-class-usage	979
destination-profile	980
dialer	980
dialer-options	981
dialin	982
dial-options	982
dial-string	983
direction	983
disable	984
disable (Interface)	984
disable (Link Protection)	984
disable-mlppp-inner-ppp-pfc	985
dlci	985
dls w	986
do-not-fragment	986
dot1 x	987
down-count	988
drop-timeout	989
ds0-options	989
dsl-options	990
dsr	990
dsr-polarity	991
dte-options	991
dtr	992
dtr-circuit	993
dtr-polarity	993

dump-on-flow-control	994
dynamic-call-admission-control	994
dynamic-profile	995
dynamic-profile (Stacked VLAN)	995
dynamic-profile (VLAN)	996
e1-options	997
e3-options	998
east-interface	999
encapsulation	1000
encapsulation (Container Interface)	1000
encapsulation (Logical Interface)	1001
encapsulation (Physical Interface)	1004
encoding	1007
epd-threshold	1008
epd-threshold (Logical Interface)	1008
epd-threshold (Physical Interface)	1009
es-options	1009
ethernet	1010
ethernet-policer-profile	1012
ethernet-ring	1013
ethernet-switch-profile	1014
evcs	1015
event	1016
event-thresholds	1016
eui-64	1017
facility-override	1017
failover-delay	1018
family	1019
fastether-options	1023
fcs	1024
feac-loop-respond	1025
filter	1026
flexible-vlan-tagging	1027
flow-control	1028
f-max-period	1028
force	1029
forwarding-class	1030
forwarding-class (ATM2 IQ Scheduler Maps)	1030
forwarding-class (Gigabit Ethernet IQ Classifier)	1031
fragment-threshold	1031
frame-error	1032
frame-period	1033
frame-period-summary	1034
framing	1035
framing (E1, E3, and T1 Interfaces)	1036
framing (10-Gigabit Ethernet Interfaces)	1037
framing (SONET and SDH Interfaces)	1037
gether-options	1038
gratuitous-arp-reply	1039
guard-interval	1040
hardware-assisted-timestamping	1040

hello-timer	1041
high-plp-max-threshold	1041
high-plp-threshold	1042
hierarchical-policer	1043
hold-interval	1044
hold-interval (OAM)	1044
hold-interval (Protection Group)	1044
hold-time	1045
hold-time (APS)	1045
hold-time (DLSw)	1046
hold-time (Physical Interface)	1047
hold-time (SONET/SDH Defect Triggers)	1048
host	1049
idle-cycle-flag	1050
idle-time	1051
idle-timeout	1051
ieee802.1p	1052
if-exceeding	1052
ignore	1053
ignore-all	1053
ignore-l3-incompletes	1054
ilmi	1054
inactivity-timeout	1055
incoming-called-number	1055
incoming-map	1056
indication	1057
indication-polarity	1057
ingress-rate-limit	1058
init-command-string	1059
initial-route-check	1060
inner-tag-protocol-id	1060
inner-vlan-id	1061
inner-vlan-id-range	1062
input	1062
input-list	1063
input-policer	1063
input-priority-map	1064
input-three-color	1064
input-vlan-map	1065
input-vlan-map (Gigabit Ethernet IQ)	1065
input-vlan-map (Aggregated Ethernet)	1066
instance	1066
interface	1067
interface (DLSw Ethernet Redundancy)	1067
interface (Hierarchical CoS Schedulers)	1068
interface (IEEE 802.1x)	1069
interface (IEEE 802.1ag OAM Connectivity-Fault Management)	1070
interface (OAM Link-Fault Management)	1071
interfaces	1072
interface-down	1072
interface-mode	1073

interface-set	1074
interface-set (Ethernet Interfaces)	1074
interface-set (IP Demux Interfaces)	1074
interface-status-tlv	1075
interface-switch	1075
interface-type	1076
interleave-fragments	1077
interval	1078
inverse-arp	1079
invert-data	1080
ipsec-sa	1080
isdn-options	1081
keep-address-and-control	1082
keepalives	1083
key	1084
l2tp-interface-id	1084
lACP	1085
lACP (802.3ad)	1085
lACP (Aggregated Ethernet)	1086
layer2-policer	1087
lcp-max-conf-req	1088
lcp-restart-timer	1088
level	1089
linear-red-profile	1089
linear-red-profiles	1090
line-encoding	1091
line-protocol	1091
line-rate	1092
link-adjacency-loss	1092
link-down	1093
link-discovery	1093
link-event-rate	1094
link-fault-management	1095
link-layer-overhead	1096
link-mode	1097
link-protection	1098
link-speed	1099
link-speed (Aggregated Ethernet)	1099
link-speed (Aggregated SONET/SDH)	1100
linktrace	1100
llc2	1101
lmi	1102
lmi (Frame Relay)	1103
lmi (Ethernet OAM)	1104
lmi-type	1105
load-interval	1105
load-threshold	1106
local-mac	1106
local-name	1107
local-password	1108
local-window	1108

lockout	1109
logical-interface-policer	1109
logical-systems	1110
log-prefix	1110
long-buildout	1111
loopback	1112
loopback (ADSL, DS0, E1/E3, SONET/SDH, SHDSL, and T1/T3)	1113
loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet) ...	1114
loopback (Serial)	1115
loopback-clear-timer	1115
loop-timing	1116
loss-priority	1116
loss-threshold	1117
low-plp-max-threshold	1117
low-plp-threshold	1118
lsq-failure-options	1118
mac	1119
mac-address	1119
mac-learn-enable	1120
mac-validate	1121
maintenance-association	1122
maintenance-domain	1123
map	1124
master-only	1125
maximum-contexts	1125
maximum-vcs	1126
maximum-requests	1126
max-retry	1127
member-interface-speed	1127
member-interface-type	1128
mep	1129
minimum-links	1130
mip-half-function	1131
mlfr-uni-nni-bundle-options	1132
mode	1133
modem-options	1133
monitor-session	1134
mpls	1134
mrru	1135
mtu	1136
multicast-dlci	1137
multicast-vci	1138
multicast-only	1138
multilink-max-classes	1139
multipoint	1139
multipoint-destination	1140
multiservice-options	1141
n391	1141
n392	1142
n393	1142
name-format	1143

native-vlan-id	1144
nbp-max-conf-req	1145
nbp-restart-timer	1145
negotiate-address	1146
negotiation-options	1146
neighbor	1147
no-allow-link-events	1147
no-asynchronous-notification	1147
no-auto-mdix	1148
no-auto-negotiation	1148
no-cbit-parity	1148
no-core-dump	1148
no-feac-loop-respond	1148
no-flow-control	1148
no-gratuitous-arp-reply	1148
no-gratuitous-arp-request	1149
no-keepalives	1149
no-long-buildout	1150
no-loopback	1150
no-mac-learn-enable	1150
node-id	1150
non-revertive	1150
no-partition	1151
no-partition (Channelized E1 IQ Interfaces)	1151
no-partition (Channelized OC1 IQ Interfaces)	1152
no-partition (Channelized OC12 IQ Interfaces)	1152
no-partition (Channelized STM1 IQ Interfaces)	1153
no-partition (Channelized T3 IQ Interfaces)	1153
no-payload-scrambler	1153
no-preempt	1154
no-redirects	1154
no-source-filtering	1154
no-syslog	1154
no-termination-request	1155
no-translate-discard-eligible	1155
no-translate-fecn-and-becn	1155
no-unframed	1155
no-z0-increment	1155
oam	1156
oam-liveness	1158
oam-period	1159
oc-slice	1160
open-timeout	1160
operating-mode	1161
optics-options	1162
otn-options	1163
output	1165
output-list	1165
output-policer	1166
output-priority-map	1166
output-three-color	1167

output-vlan-map	1168
output-vlan-map (Gigabit Ethernet IQ)	1168
output-vlan-map (Aggregated Ethernet)	1169
overflow	1170
overflow (Receive Bucket)	1170
overflow (Transmit Bucket)	1170
paired-group	1171
pap	1172
pap-password	1173
partition	1174
passive	1175
passive (CHAP)	1175
passive (PAP)	1176
passive-monitor-mode	1176
path-database-size	1177
path-trace	1178
payload-scrambler	1179
payload-size	1180
p-bit-timeout	1180
pdu-interval	1181
pdu-threshold	1181
peer	1182
peer-unit	1182
performance-monitoring	1183
periodic	1183
per-unit-scheduler	1184
pfc	1185
pic-type	1185
plp1	1186
plp-to-clp	1187
point-to-point	1187
policer	1188
policer (CoS)	1188
policer (Interface)	1189
policer (MAC)	1190
pool	1191
pop	1192
pop-all-labels	1193
pop-pop	1194
pop-swap	1194
port	1195
port-priority	1195
port-status-tlv	1196
post-service-filter	1196
pppoe-options	1197
ppp-options	1198
preempt	1199
preferred	1200
preferred-source-address	1201

premium	1202
premium (Hierarchical Policer)	1202
premium (Output Priority Map)	1203
premium (Policer)	1203
preserve-interface	1204
primary	1205
primary (Address on Interface)	1205
primary (AS PIC or MultiServices PIC Interfaces)	1205
priority	1206
priority (DLSw)	1206
priority (OAM Connectivity-Fault Management)	1207
priority (Schedulers)	1207
priority-cost	1208
promiscuous-mode	1208
protect-circuit	1209
protection-group	1210
protocol-down	1210
protocols	1211
proxy	1211
proxy-arp	1212
push	1212
push-push	1213
queue-depth	1213
queue-length	1214
queues	1214
quiet-period	1215
ranges	1216
ranges (Dynamic Stacked VLAN)	1216
ranges (Dynamic VLAN)	1216
rate	1217
reassemble-packets	1217
reauthentication	1218
receive-bucket	1218
receive-options-packets	1219
receive-ttl-exceeded	1219
red-differential-delay	1220
redial-delay	1220
redundancy-group	1221
redundancy-options	1222
remote	1223
remote-loopback	1223
remote-loopback-respond	1224
remote-mep	1225
request	1225
required-depth	1226
restore-interval	1226
retries	1227
revertive	1227
revert-time	1228
rfc-2615	1228
ring-protection-link-end	1229

ring-protection-link-owner	1229
routing-instance	1230
rpf-check	1231
rtp	1232
rts	1232
rts-polarity	1233
rtvbr	1234
sampling	1235
satop-options	1236
scheduler-maps	1237
schedulers	1237
secondary	1238
send-critical-event	1238
serial-options	1239
server	1240
server-timeout	1240
service	1241
service-domain	1241
service-filter	1242
service-name	1242
service-set	1243
services	1243
services-options	1244
shaping	1245
shdsl-options	1246
short-name-format	1246
short-sequence	1247
snext	1247
snr-margin	1248
sonet-options	1249
source	1251
source-address-filter	1252
source-class-usage	1253
source-filtering	1254
speed	1255
speed (Ethernet)	1255
speed (MX Series DPC)	1256
speed (SONET/SDH)	1257
spid1	1257
spid2	1258
stacked-vlan-ranges	1258
stacked-vlan-tagging	1259
start-end-flag	1260
static-tei-val	1261
supplicant	1261
supplicant-timeout	1262
swap	1262
swap-push	1263
swap-swap	1263
switching-mode	1264
switch-options	1264

switch-port	1265
switch-type	1266
symbol-period	1267
syslog	1268
syslog (Interfaces)	1268
syslog (Monitoring)	1269
syslog (OAM Action)	1269
system-priority	1270
t1-options	1271
t1-time	1272
t2-time	1272
t310	1273
t391	1273
t392	1274
t3-options	1275
tag-protocol-id	1276
tag-protocol-id (TPIDs Expected to Be Sent or Received)	1276
tag-protocol-id (TPID to Rewrite)	1277
tei-option	1277
then	1278
threshold	1278
timeslots	1279
tm	1280
tm-polarity	1280
traceoptions	1281
traceoptions (Individual Interfaces)	1282
traceoptions (Interface Process)	1283
traceoptions (LACP)	1285
traceoptions (PPP Process)	1287
track	1290
translate-discard-eligible	1291
translate-fecn-and-becn	1291
transmit-bucket	1292
transmit-clock	1292
transmit-period	1293
transmit-weight	1294
transmit-weight (ATM2 IQ CoS Forwarding Class)	1294
transmit-weight (ATM2 IQ Virtual Circuit)	1295
traps	1295
trej-time	1296
trigger	1297
trigger-link-failure	1298
trunk-bandwidth	1298
trunk-id	1299
ttl	1299
tunnel	1300
underlying-interface	1301
unframed	1302
unidirectional	1302
unit	1303

unnumbered-address	1309
unnumbered-address (Demux)	1309
unnumbered-address (Ethernet)	1310
unnumbered-address (PPP)	1310
up-count	1311
vbr	1312
vc-cos-mode	1313
vci	1314
vci-range	1315
virtual-switch	1315
vlan-id	1316
vlan-id (VLAN ID to Be Bound to a Logical Interface)	1316
vlan-id (Logical Port in Bridge Domain)	1317
vlan-id (VLAN ID to Rewrite)	1317
vlan-id (Outer VLAN ID)	1318
vlan-id-list	1319
vlan-id-list (Interface in Bridge Domain)	1319
vlan-id-list (Ethernet VLAN Circuit)	1320
vlan-id-range	1322
vlan-ranges	1323
vlan-rewrite	1324
vlan-tagging	1324
vlan-tags	1325
vlan-tags (Dual-Tagged Logical Interface)	1326
vlan-tags (Stacked VLAN Tags)	1328
vlan-tags-outer	1329
vlan-vci-tagging	1330
vpi	1331
vpi (ATM CCC Cell-Relay Promiscuous Mode)	1331
vpi (Define Virtual Path)	1332
vpi (Logical Interface and Interworking)	1333
vtmapping	1333
watch-list	1334
wavelength	1335
west-interface	1337
working-circuit	1337
yellow-differential-delay	1338
z0-increment	1338

Part 14

Index

Index	1341
Index of Statements and Commands	1367

List of Figures

Part 2

Router Interfaces Configuration Concepts

Chapter 2	Router Interfaces Overview	31
	Figure 1: APS Interface	37
	Figure 2: Container Interface	38
	Figure 3: Routing Matrix	58
	Figure 4: Routing Matrix Based on a TX Matrix Plus Router	60
	Figure 5: Interface Slot, PIC, and Port Locations	63
	Figure 6: Clock Sources	65
Chapter 5	Configuring Protocol Family and Interface Address Properties	169
	Figure 7: DLSw Ethernet Redundancy Topology	184
	Figure 8: Hierarchical Policer	198
	Figure 9: Unicast RPF with Routing Asymmetry	212
	Figure 10: Prefix Accounting with Source and Destination Classes	215
Chapter 6	Configuring Circuit and Translational Cross-Connects	223
	Figure 11: Layer 2 Switching Circuit Cross-Connect	224
	Figure 12: Example Topology of a Switching Circuit Cross-Connect with Frame Relay CCC Encapsulation	234
	Figure 13: Layer 2.5 Switching Translational Cross-Connect	235
	Figure 14: Interface-to-Interface Circuit Cross-Connect over Aggregated Ethernet Interfaces	236
	Figure 15: Remote Interface-LSP-Interface Circuit Cross-Connect over Aggregated Ethernet Interfaces	238
	Figure 16: ATM-to-Ethernet Interworking	240

Part 4

Configuring Serial Interfaces

Chapter 12	Configuring Serial Interfaces	263
	Figure 17: Serial Interface Clocking Mode	269
	Figure 18: Serial Interface LIU Loopback	275
	Figure 19: Serial Interface Local Loopback	276

Part 5

Configuring ATM Interfaces

Chapter 13	Configuring ATM Interfaces	281
	Figure 20: Layer 2 Circuit Trunk Topology	304
	Figure 21: Example Topology for Router with Eight Queues	344

Part 7

Configuring Channelized Interfaces

Chapter 17	Channelized Interfaces	385
-------------------	-------------------------------	------------

	Figure 22: Channelized OC48/STM16 IQE PIC (in SONET Mode)	396
	Figure 23: Channelized OC48/STM16 IQE PIC (in SDH Mode)	396
	Figure 24: Channelized OC12 IQ PIC and Channelized OC12/STM4 IQE PIC (in SONET Mode)	397
	Figure 25: Channelized OC12/STM4 IQE PIC (in SDH Mode)	397
	Figure 26: Channelized OC12/STM4 IQ PIC (in SDH Mode)	398
	Figure 27: Channelized OC3 Ports (in SONET Mode) on Channelized OC3 IQ and Channelized OC3/STM1 IQE PICs	398
	Figure 28: Channelized CSTM1 Ports (in SDH Mode) on Channelized OC3/STM1 IQE PIC	399
	Figure 29: Channelized STM1 IQ PIC	399
	Figure 30: Channelized CDS3/E3 IQE PIC (in DS3 Mode)	399
	Figure 31: Channelized CDS3/E3 IQE PIC (in E3 Mode)	400
	Figure 32: Channelized DS3 IQ PIC	400
	Figure 33: Channelized T1 IQ and IQE PIC	400
	Figure 34: Channelized E1 IQ and IQE PIC	400
Chapter 18	Configuring Channelized OC48/STM16 IQE Interfaces	405
	Figure 35: Sample Channelization of OC48/STM16 IQE PIC (SONET Mode)	405
	Figure 36: Sample Channelization of OC48/STM16 IQE PIC (SDH Mode)	406
	Figure 37: Sample Channelization of OC48/STM16 IQE PIC to E3 Channels	407
	Figure 38: T1 Interfaces on a Channelized OC48 PIC	411
	Figure 39: Sample Channelization of OC48 IQE PIC	414
Chapter 19	Configuring Channelized OC12/STM4 Interfaces	423
	Figure 40: Sample Channelization of OC12/STM4 IQ or IQE PIC (SONET Mode)	424
	Figure 41: Sample Channelization of OC12/STM4 IQE PIC (SDH Mode)	425
	Figure 42: Sample Channelization of OC12/STM4 IQ PIC (SDH Mode)	426
	Figure 43: Sample Channelization of OC12 PIC (non IQ and IQE)	426
	Figure 44: T1 Interfaces on a Channelized OC12 PIC	430
	Figure 45: Sample Channelization of OC12 IQE PIC	432
Chapter 20	Configuring Channelized OC3 IQ and IQE Interfaces	455
	Figure 46: Channelized OC3 IQ Interface Example for Show Interfaces Controller	456
	Figure 47: T1 Interfaces on a Channelized OC3 PIC	460
	Figure 48: Sample Channelization of OC3 IQ or IQE PIC	460
Chapter 22	Configuring Channelized T3 Interfaces	479
	Figure 49: Sample Channelization of DS3 IQ or IQE PIC	487
 Part 8	 Configuring Circuit Emulation PICs	
Chapter 26	Circuit Emulation PICs Overview	519
	Figure 50: Mobile Backhaul Application	520
Chapter 28	Configuring ATM Support on Circuit Emulation PICs	529
	Figure 51: 12-Port T1/E1 CE PIC Possible Interfaces (T1 Size)	531
	Figure 52: 12-Port T1/E1 CE PIC Possible Interfaces (E1 Size)	531
	Figure 53: 4-Port Channelized COC3/STM1 CE PIC Possible Interfaces (T1 Size)	533

Figure 54: 4-Port Channelized COC3/STM1 CE PIC Possible Interfaces (E1 Size)	533
--	-----

Part 9

Configuring E1, E3, T1, and T3 Interfaces

Chapter 29	Configuring E1 Interfaces	543
	Figure 55: Remote and Local E1 Loopback	547
Chapter 30	Configuring E3 Interfaces	551
	Figure 56: Remote and Local E3 Loopback	556
Chapter 31	Configuring T1 Interfaces	559
	Figure 57: Remote and Local T1 Loopback	565
Chapter 32	Configuring T3 Interfaces	569
	Figure 58: Remote and Local T3 Loopback	577

Part 10

Configuring Ethernet Interfaces

Chapter 35	Configuring Aggregated Ethernet Interfaces	623
	Figure 59: Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers	637
Chapter 38	Configuring TCC and Layer 2.5 Switching	665
	Figure 60: Topology of Layer 2.5 Translational Cross-Connect	667
Chapter 40	Configuring Unrestricted Proxy ARP	671
	Figure 61: Edge Device Case for Unrestricted Proxy ARP	672
	Figure 62: Core Device Case for Unrestricted Proxy ARP	672
Chapter 43	Configuring IEEE 802.1ag OAM Connectivity-Fault Management	679
	Figure 63: Relationship Among MEPs, MIPs, and Maintenance Domain Levels	681
	Figure 64: Relationship Among Bridges, Maintenance Domains, Maintenance Associations, and MEPs	681
	Figure 65: Scope of the E-LMI Protocol	690
	Figure 66: E-LMI Configuration for a Point-to-Point EVC (SVLAN) Monitored by CFM	694
	Figure 67: Layer 2 VPN Topology	709
Chapter 44	Configuring ITU-T Y.1731 Ethernet Service OAM	711
	Figure 68: Relationship of MEPs, MIPs, and Maintenance Domain Levels	712
Chapter 55	Configuring Point-to-Point Protocol over Ethernet	785
	Figure 69: PPPoE Session on an Ethernet Loop	786
Chapter 56	Configuring Ethernet Ring Protection Switching	799
	Figure 70: Protocol Packets from the Network to the Router	802
	Figure 71: Protocol Packets from the Router to the Network	802
	Figure 72: Example of a Three Node Ring	805

Part 11

Configuring ISDN Interfaces

Chapter 58	Configuring ISDN Interfaces	819
	Figure 73: ISDN Backup Topology	827
	Figure 74: Dialer Filter Topology	829

Figure 75: Bandwidth-on-Demand Topology831

Figure 76: Dialer Watch Topology836

Part 12

Configuring SONET Interfaces

Chapter 59	Configuring SONET/SDH Interfaces	843
	Figure 77: APS/MSP Configuration Topologies	860
	Figure 78: APS Load Sharing Between Circuit Pairs	871

List of Tables

About This Guide	lxi
Table 1: Notice Icons	lxv
Table 2: Text and Syntax Conventions	lxv

Part 2

Router Interfaces Configuration Concepts

Chapter 2	Router Interfaces Overview	31
	Table 3: Encapsulation Support by Interface Type	39
	Table 4: FPC Numbering for T640 Routers in a Routing Matrix	59
	Table 5: One-to-One FPC Numbering for T640 Routers in a Routing Matrix	59
	Table 6: FPC Numbering for T1600 Routers in a Routing Matrix	61
	Table 7: One-to-One FPC Numbering for T1600 Routers in a Routing Matrix	61
Chapter 3	Configuring Physical Interface Properties	67
	Table 8: Statements for Physical Interface Properties	77
	Table 9: Media MTU Sizes by Interface Type for M5, M7i with CFEB, M10, M10i with CFEB, M20, and M40 Routers	98
	Table 10: Media MTU Sizes by Interface Type for M40e Routers	99
	Table 11: Media MTU Sizes by Interface Type for M160 Routers	100
	Table 12: Media MTU Sizes by Interface Type for M7i with CFEB-E, M10i with CFEB-E, M320 and M120 Routers	100
	Table 13: Media MTU Sizes by Interface Type for T320 Routers	101
	Table 14: Media MTU Sizes by Interface Type for T640 Platforms	101
	Table 15: Media MTU Sizes by Interface Type for J2300 Platforms	102
	Table 16: Media MTU Sizes by Interface Type for J4300 and J6300 Platforms	102
	Table 17: Media MTU Sizes by Interface Type for J4350 and J6350 Platforms	102
	Table 18: Encapsulation Overhead by Encapsulation Type	104
	Table 19: Type 1 PIC Mode Combinations	125
	Table 20: Type 2 PIC Mode Combinations	125
	Table 21: Loopback Modes by Interface Type	132
	Table 22: BERT Capabilities by Interface Type	136
Chapter 4	Configuring Logical Interface Properties	143
	Table 23: Statements for Logical Interface Properties	147

Part 4

Configuring Serial Interfaces

Chapter 12	Configuring Serial Interfaces	263
	Table 24: Signal Handling by Serial Interface Type	272

Part 5	Configuring ATM Interfaces	
Chapter 13	Configuring ATM Interfaces	281
	Table 25: ATM1 and ATM2 IQ Supported Features	287
	Table 26: ILMI Support by Encapsulation Type	292
	Table 27: Shaping Rate Range by Interface Type	320
	Table 28: ATM1 Traffic-Shaping Rates	323
	Table 29: EPD Threshold Range by Interface Type	327
	Table 30: ATM Logical Interface Encapsulation Types	331
Chapter 14	Configuring ATM-over-ADSL Interfaces	355
	Table 31: ATM-over-ADSL Operational Modes	357
	Table 32: ATM-over-ADSL Encapsulation Types	358
Part 6	Configuring Frame Relay	
Chapter 16	Configuring Frame Relay	371
	Table 33: PIC Support for Enhanced Frame Relay Encapsulation Types	374
Part 7	Configuring Channelized Interfaces	
Chapter 17	Channelized Interfaces	385
	Table 34: Frame Relay DLCI Limitations for Channelized Interfaces	388
	Table 35: Per Unit Scheduler DLCI Limitations for Channelized Interfaces	389
	Table 36: Protocol Family Combinations	389
	Table 37: Clocking Capabilities by Channelized PIC Type	391
	Table 38: Structural Differences: Channelized IQE PICs	401
	Table 39: Structural Differences: Channelized IQ PICs	402
	Table 40: Structural Differences: Channelized PICs	402
Chapter 19	Configuring Channelized OC12/STM4 Interfaces	423
	Table 41: OC12-to-DS3 Numbering Scheme	445
Chapter 21	Configuring Channelized STM1 Interfaces	465
	Table 42: Channelized STM1-to-E1 Channel Mapping	473
Chapter 22	Configuring Channelized T3 Interfaces	479
	Table 43: Ranges for Channelized DS3-to-DS0 Configuration	483
Chapter 23	Configuring Channelized T1 Interfaces	495
	Table 44: Ranges for Channelized T1 IQ Configuration	498
Chapter 24	Configuring Channelized E1 Interfaces	501
	Table 45: Ranges for Channelized E1 Configuration	504
Part 9	Configuring E1, E3, T1, and T3 Interfaces	
Chapter 30	Configuring E3 Interfaces	551
	Table 46: Subrate Values for E3 Digital Link Compatibility Mode	554
Chapter 32	Configuring T3 Interfaces	569
	Table 47: Subrate Values for T3 Digital Link Compatibility Mode	573

Part 10**Configuring Ethernet Interfaces**

Chapter 34	Configuring 802.1Q VLANs	599
	Table 48: VLAN ID Range by Interface Type	600
	Table 49: Configuration Statements Used to Bind VLAN IDs to Logical Interfaces	605
	Table 50: Encapsulation Inside Circuits CCC-Connected by VLAN-Bundled Logical Interfaces	613
Chapter 35	Configuring Aggregated Ethernet Interfaces	623
	Table 51: Untagged Aggregated Ethernet and LACP Support by PIC and Platform	633
Chapter 36	Stacking and Rewriting Gigabit Ethernet VLAN Tags	641
	Table 52: Rewrite Operations on Not Tagged, Single-Tagged, and Dual-Tagged Frames	643
	Table 53: Applying Rewrite Operations to VLAN Maps	644
	Table 54: Rewrite Operations and Statement Usage for Input VLAN Maps	647
	Table 55: Rewrite Operations and Statement Usage for Output VLAN Maps	647
	Table 56: Input VLAN map statements allowed for ethernet-ccc and ethernet-vpls encapsulations	653
	Table 57: Output VLAN map statements allowed for ethernet-ccc and ethernet-vpls encapsulations	653
	Table 58: Rules for applying rewrite operations to VLAN maps	653
Chapter 43	Configuring IEEE 802.1ag OAM Connectivity-Fault Management	679
	Table 59: Format of TLVs	698
	Table 60: Type Field Values for Various TLVs for CFM PDUs	698
	Table 61: Port Status TLV Format	700
	Table 62: Port Status TLV Values	701
	Table 63: Interface Status TLV Format	703
	Table 64: Interface Status TLV Values	703
Chapter 44	Configuring ITU-T Y.1731 Ethernet Service OAM	711
	Table 65: ETH-DM Statistics	721
	Table 66: ETH-DM Frame Counts	722
Chapter 48	Configuring Gigabit Ethernet Accounting and Policing	755
	Table 67: Capabilities of Gigabit Ethernet IQ and Gigabit Ethernet with SFPs	755
	Table 68: Default Forwarding Classes	759
Chapter 49	Configuring Gigabit Ethernet Autonegotiation	767
	Table 69: Mode and Autonegotiation Status (Local)	768
	Table 70: Mode and Autonegotiation Status (Remote)	770
Chapter 52	Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength	779
	Table 71: Wavelength-to-Frequency Conversion Matrix	780

Part 12**Configuring SONET Interfaces**

Chapter 59	Configuring SONET/SDH Interfaces	843
	Table 72: Type 1 PIC Mode Combinations	847
	Table 73: Type 2 PIC Mode Combinations	848

Table 74: SONET/SDH Framing Bytes for Specific Speeds	849
Table 75: SONET/SDH Default Settings	855
Table 76: SONET/SDH and ATM Active Alarms and Defects	855

About This Guide

This preface provides the following guidelines for using the *JUNOS® Software Network Interfaces Configuration Guide*:

- JUNOS Documentation and Release Notes on page lxi
- Objectives on page lxii
- Audience on page lxii
- Supported Routing Platforms on page lxii
- Using the Indexes on page lxiii
- Using the Examples in This Manual on page lxiii
- Documentation Conventions on page lxiv
- Documentation Feedback on page lxvi
- Requesting Technical Support on page lxvii

JUNOS Documentation and Release Notes

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

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

Juniper Networks supports a technical book program to publish books by Juniper Networks engineers and subject matter experts with book publishers around the world. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration using JUNOS Software and Juniper Networks devices. In addition, the Juniper Networks Technical Library, published in conjunction with O'Reilly Media, explores improving network security, reliability, and availability using JUNOS configuration techniques. All the books are for sale at technical bookstores and book outlets around the world. The current list can be viewed at <http://www.juniper.net/books>.

Objectives

This guide provides an overview of the network interfaces features of the JUNOS Software and describes how to configure these properties on the routing platform.



NOTE: For additional information about JUNOS Software—either corrections to or information that might have been omitted from this guide—see the software release notes at <http://www.juniper.net/>.

Audience

This guide is designed for network administrators who are configuring and monitoring a Juniper Networks M Series, MX Series, T Series, EX Series, or J Series router or switch.

To use this guide, you need a broad understanding of networks in general, the Internet in particular, networking principles, and network configuration. You must also be familiar with one or more of the following Internet routing protocols:

- Border Gateway Protocol (BGP)
- Distance Vector Multicast Routing Protocol (DVMRP)
- Intermediate System-to-Intermediate System (IS-IS)
- Internet Control Message Protocol (ICMP) router discovery
- Internet Group Management Protocol (IGMP)
- Multiprotocol Label Switching (MPLS)
- Open Shortest Path First (OSPF)
- Protocol-Independent Multicast (PIM)
- Resource Reservation Protocol (RSVP)
- Routing Information Protocol (RIP)
- Simple Network Management Protocol (SNMP)

Personnel operating the equipment must be trained and competent; must not conduct themselves in a careless, willfully negligent, or hostile manner; and must abide by the instructions provided by the documentation.

Supported Routing Platforms

For the features described in this manual, the JUNOS Software currently supports the following routing platforms:

- J Series
- M Series

- MX Series
- T Series

Using the Indexes

This reference contains two indexes: a complete index that includes topic entries, and an index of statements and commands only.

In the index of statements and commands, an entry refers to a statement summary section only. In the complete index, the entry for a configuration statement or command contains at least two parts:

- The primary entry refers to the statement summary section.
- The secondary entry, *usage guidelines*, refers to the section in a configuration guidelines chapter that describes how to use the statement or command.

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

```
system {
  scripts {
    commit {
      file ex-script.xsl;
    }
  }
}
interfaces {
  fxp0 {
    disable;
    unit 0 {
```

```

        family inet {
            address 10.0.0.1/24;
        }
    }
}

```

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

```

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

```

Merging a Snippet

To merge a snippet, follow these steps:

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

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

```

commit {
    file ex-script-snippet.xsl; }

```

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

```

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

```

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

```

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

```

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

Documentation Conventions

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

Table 1: Notice Icons





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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the <code>configure</code> command: user@host> configure
Fixed-width text like this	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic text like this</i>	<ul style="list-style-type: none"> Introduces important new terms. Identifies book names. Identifies RFC and Internet draft titles. 	<ul style="list-style-type: none"> A policy <i>term</i> is a named structure that defines match conditions and actions. <i>JUNOS System Basics Configuration Guide</i> RFC 1997, <i>BGP Communities Attribute</i>
<i>Italic text like this</i>	Represents variables (options for which you substitute a value) in commands or configuration statements.	Configure the machine's domain name: [edit] root@# set system domain-name <i>domain-name</i>
Plain text like this	Represents names of configuration statements, commands, files, and directories; IP addresses; configuration hierarchy levels; or labels on routing platform components.	<ul style="list-style-type: none"> To configure a stub area, include the stub statement at the [edit protocols ospf area area-id] hierarchy level. The console port is labeled CONSOLE.
< > (angle brackets)	Enclose optional keywords or variables.	stub <default-metric <i>metric</i> >;

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

Convention	Description	Examples
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (string1 string2 string3)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [community-ids]
Indentation and braces ({ })	Identify a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop address; retain; } } }
;(semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
J-Web GUI Conventions		
Bold text like this	Represents J-Web graphical user interface (GUI) items you click or select.	<ul style="list-style-type: none">■ In the Logical Interfaces box, select All Interfaces.■ To cancel the configuration, click Cancel.
> (bold right angle bracket)	Separates levels in a hierarchy of J-Web selections.	In the configuration editor hierarchy, select Protocols > Ospf .

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can 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 name
- Document part number
- Page number
- Software release version (not required for Network Operations Guides [NOGs])

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 postsales 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/710059.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 located at <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, visit us at <http://www.juniper.net/support/requesting-support.html>.

Part 1

Network Interfaces Configuration Statements Overview

- Network Interfaces Configuration Statements and Hierarchy on page 3

Chapter 1

Network Interfaces Configuration Statements and Hierarchy

This chapter shows the complete configuration statement hierarchy, listing all possible configuration statements and showing their level in the configuration hierarchy. When you are configuring the JUNOS Software, your current hierarchy level is shown in the banner on the line preceding the `user@host#` prompt.

This section contains the following topics:

- [edit chassis] Hierarchy Level on page 3
- [edit interfaces] Hierarchy Level on page 4
- [edit logical-systems] Hierarchy Level on page 19
- [edit protocols connections] Hierarchy Level on page 24
- [edit protocols dot1x] Hierarchy Level on page 25
- [edit protocols lacp] Hierarchy Level on page 25
- [edit protocols oam] Hierarchy Level on page 25
- [edit protocols ppp] Hierarchy Level on page 27
- [edit protocols protection-group] Hierarchy Level on page 27
- [edit protocols vrrp] Hierarchy Level on page 27
- [edit system processes] Hierarchy Level on page 27

[edit chassis] Hierarchy Level

```
chassis {
  aggregated-devices {
    ethernet {
      device-count number;
    }
    sonet {
      device-count number;
    }
  }
  channel-group number {
    ethernet {
      device-count number;
    }
    fpc slot-number{
```

```

pic pic-number {
  adaptive-services{
    service-package (layer-2 | layer-3);
  }
  aggregate-ports;
  atm-cell-relay-accumulation;
  atm-l2circuit-mode (aal5 | cell | trunk trunk);
  ce1 {
    e1 link-number {
      channel-group group-number;
      timeslots time-slot-range;
    }
  }
  ct1 {
    t1 link-number {
      channel-group group-number;
      timeslots time-slot-range;
    }
  }
  ct3 {
    port port-number {
      t1 link-number {
        channel-group group-number;
        timeslots time-slot-range;
      }
    }
    framing sdh;
  }
  max-queues-per-interface number;
  mlfr-uni-nni-bundles num-intf;
  no-concatenate;
  shdsl {
    pic-mode (1-port-atm | 2-port-atm);
  }
  vtmapping (klm | itu-t);
}
}
}

```

[edit interfaces] Hierarchy Level

The statements at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level can also be configured at the [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*] hierarchy level.



NOTE: The accounting-profile statement is an exception to this rule. The accounting-profile statement can be configured at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level, but it cannot be configured at the [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*] hierarchy level.

```

interfaces {
  traceoptions {
    file filename <files number> <match regular-expression> <size size>
      <world-readable | no-world-readable> ;
    flag flag <disable>;
  }
  interface-name {
    accounting-profile name;
    aggregated-ether-options {
      (flow-control | no-flow-control);
      lacp {
        (active | passive);
        link-protection{
          disable;
          (revertive | non-revertive);
          periodic interval;
          system-priority priority;
        }
        link-protection;
        link-speed speed;
        (loopback | no-loopback);
        minimum-links number;
        source-address-filter {
          mac-address;
        }
        (source-filtering | no-source-filtering);
      }
    }
    aggregated-sonet-options {
      link-speed speed | mixed;
      minimum-links number;
    }
    atm-options {
      cell-bundle-size cells;
      ilmi;
      linear-red-profiles profile-name {
        high-plp-max-threshold percent;
        low-plp-max-threshold percent;
        queue-depth cells high-plp-threshold percent low-plp-threshold percent;
      }
      mpls {
        pop-all-labels {
          required-depth number;
        }
      }
      pic-type (atm1 | atm2);
      plp-to-clp;
      promiscuous-mode {
        vpi vpi-identifier;
      }
      scheduler-maps map-name {
        forwarding-class class-name {
          epd-threshold cells plp1 cells;
          linear-red-profile profile-name;
          priority (high | low);
          transmit-weight (cells number | percent number);
        }
      }
    }
  }
}

```

```

        vc-cos-mode (alternate | strict);
    }
    vpi vpi-identifier {
        maximum-vcs maximum-vcs;
        oam-liveness {
            down-count cells;
            up-count cells;
        }
        oam-period (seconds | disable);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate
             sustained rate burst length);
            queue-length number;
        }
    }
}
clocking clock-source;
data-input (system | interface interface-name);
dce;
serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dsr-polarity (negative | positive);
    dte-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dtr-circuit (balanced | unbalanced);
    dtr-polarity (negative | positive);
    encoding (nrz | nrzi);
    indication-polarity (negative | positive);
    line-protocol protocol;
    loopback mode;
    rts-polarity (negative | positive);
    tm-polarity (negative | positive);

```

```

        transmit-clock invert;
    }
    description text;
    dialer-options {
        pool pool-name <priority priority>;
    }
    disable;
    ds0-options {
        bert-algorithm algorithm;
        bert-error-rate rate;
        bert-period seconds;
        byte-encoding (nx56 | nx64);
        fcs (16 | 32);
        idle-cycle-flag (flags | ones);
        invert-data;
        loopback payload;
        start-end-flag (filler | shared);
    }
    e1-options {
        bert-error-rate rate;
        bert-period seconds;
        fcs (16 | 32);
        framing (g704 | g704-no-crc4 | unframed);
        idle-cycle-flag (flags | ones);
        invert-data;
        loopback (local | remote);
        start-end-flag (filler | shared);
        timeslots time-slot-range;
    }
    e3-options {
        atm-encapsulation (direct | plcp);
        bert-algorithm algorithm;
        bert-error-rate rate;
        bert-period seconds;
        buildout feet;
        compatibility-mode (digital-link | kentrox | larscom) <subrate value>;
        fcs (16 | 32);
        framing (g.751 | g.832);
        idle-cycle-flag (filler | shared);
        invert-data;
        loopback (local | remote);
        (payload-scrambler | no-payload-scrambler);
        start-end-flag (filler | shared);
        (unframed | no-unframed);
    }
    encapsulation type;
    es-options {
        backup-interface es-fpc/pic/port;
    }
    fastether-options {
        802.3ad aex;
        (flow-control | no-flow-control);
        ignore-l3-incompletes;
        ingress-rate-limit rate;
        (loopback | no-loopback);
        mpls {

```

```

        pop-all-labels {
            required-depth number;
        }
    }
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
}
flexible-vlan-tagging;
gigether-options {
    802.3ad aex;
    (asynchronous-notification | no-asynchronous-notification);
    (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
        local-interface-offline>;
    auto-reconnect seconds;
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    (loopback | no-loopback);
    mpls {
        pop-all-labels {
            required-depth number;
        }
    }
    no-auto-mdix;
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
    ethernet-switch-profile {
        (mac-learn-enable | no-mac-learn-enable);
        tag-protocol-id [ tpids ];
        ethernet-policer-profile {
            input-priority-map {
                ieee802.1p premium [ values ];
            }
            output-priority-map {
                classifier {
                    premium {
                        forwarding-class class-name {
                            loss-priority (high | low);
                        }
                    }
                }
            }
        }
    }
    policer cos-policer-name {
        aggregate {
            bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
            burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
        }
        premium {
            bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
            burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
        }
    }
}

```



```

    }
}
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
interface-set interface-set-name {
    interface ethernet-interface-name {
        (unit unit-number | vlan-tags-outer vlan-tag);
    }
    interface interface-name {
        (unit unit-number);
    }
}
}
isdn-options {
    bchannel-allocation (ascending | descending);
    calling-number number;
    pool pool-name <priority priority>;
    spid1 spid-string;
    spid2 spid-string;
    static-tei-val value;
    switch-type (att5e | etsi | ni1 | ntdms100 | ntt);
    t310 seconds;
    tei-option (first-call | power-up);
}
keepalives <down-count number> <interval seconds> <up-count number>;
link-mode mode;
lmi (Frame Relay) {
    lmi-type (ansi | itu);
    n391dte number;
    n392dce number;
    n392dte number;
    n393dce number;
    n393dte number;
    t391dte seconds;
    t392dce seconds;
}
lsq-failure-options {
    no-termination-request;
    [ trigger-link-failure interface-name ];
}
}
mac mac-address;
mlfr-uni-nni-bundle-options {
    acknowledge-retries number;
    acknowledge-timer milliseconds;
    action-red-differential-delay (disable-tx | remove-link);
    drop-timeout milliseconds;
    fragment-threshold bytes;
    cisco-interoperability send-lip-remove-link-for-link-reject;
    hello-timer milliseconds;
    link-layer-overhead percent;
    lmi-type (ansi | itu);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;

```

```

        red-differential-delay milliseconds;
        t391 seconds;
        t392 seconds;
        yellow-differential-delay milliseconds;
        encapsulation type;
    }
    modem-options {
        dialin (console | routable);
        init-command-string initialization-command-string;
    }
    mtu bytes;
    multiservice-options {
        (core-dump | no-core-dump);
        (syslog | no-syslog);
    }
    native-vlan-id number;
    no-gratuitous-arp-request;
    no-keepalives;
    no-partition {
        interface-type type;
    }
    optics-options {
        wavelength nm;
    }
    partition partition-number oc-slice oc-slice-range interface-type type;
    timeslots time-slot-range;
    passive-monitor-mode;
    per-unit-scheduler;
    ppp-options {
        chap {
            access-profile name;
            default-chap-secret name;
            local-name name;
            passive;
        }
        compression {
            acfc;
            pfc;
        }
        dynamic-profile profile-name;
        no-termination-request;
        pap {
            access-profile name;
            local-name name;
            local-password password;
            passive;
        }
    }
    receive-bucket {
        overflow (discard | tag);
        rate percentage;
        threshold bytes;
    }
    redundancy-options {
        primary sp-fpc/pic/port;
        secondary sp-fpc/pic/port;
    }

```

```

}
satop-options {
    payload-size n;
}
schedulers number;
serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dsr-polarity (negative | positive);
    dte-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    }
    dtr-circuit (balanced | unbalanced);
    dtr-polarity (negative | positive);
    encoding (nrz | nrzi);
    indication-polarity (negative | positive);
    line-protocol protocol;
    loopback mode;
    rts-polarity (negative | positive);
    tm-polarity (negative | positive);
    transmit-clock invert;
}
services-options {
    inactivity-timeout seconds;
    open-timeout seconds;
    syslog {
        host hostname {
            facility-override facility-name;
            log-prefix prefix-number;
            services priority-level;
        }
    }
}
shdsl-options {

```

```

    annex (annex-a | annex-b);
    line-rate line-rate;
    loopback (local | remote);
    snr-margin {
        current margin;
        snext margin;
    }
}
sonet-options {
    aggregate asx;
    aps {
        advertise-interval milliseconds;
        annex-b
        authentication-key key;
        force;
        hold-time milliseconds;
        lockout;
        neighbor address;
        paired-group group-name;
        preserve-interface;
        protect-circuit group-name;
        request;
        revert-time seconds;
        switching-mode (bidirectional | unidirectional);
        working-circuit group-name;
    }
    bytes {
        c2 value;
        e1-quiet value;
        f1 value;
        f2 value;
        s1 value;
        z3 value;
        z4 value;
    }
    fcs (16 | 32);
    loopback (local | remote);
    mpls {
        pop-all-labels {
            required-depth number;
        }
    }
    path-trace trace-string;
    (payload-scrambler | no-payload-scrambler);
    rfc-2615;
    trigger {
        defect ignore;
        hold-time up milliseconds down milliseconds;
    }
    vtmapping (itu-t | klm);
    (z0-increment | no-z0-increment);
}
speed (10m | 100m | 1g | oc3 | oc12 | oc48);
stacked-vlan-tagging;
switch-options {
    switch-port port-number {

```

```

        (auto-negotiation | no-auto-negotiation);
        speed (10m | 100m | 1g);
        link-mode (full-duplex | half-duplex);
    }
}
t1-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout value;
    byte-encoding (nx56 | nx64);
    crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);
    crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);
    fcs (16 | 32);
    framing (esf | sf);
    idle-cycle-flag (flags | ones);
    invert-data;
    line-encoding (ami | b8zs);
    loopback (local | payload | remote);
    remote-loopback-respond;
    start-end-flag (filler | shared);
    timeslots time-slot-range;
}
t3-options {
    atm-encapsulation (direct | plcp);
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout feet;
    (cbit-parity | no-cbit-parity);
    compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate
        value>;
    fcs (16 | 32);
    (feac-loop-respond | no-feac-loop-respond);
    idle-cycle-flag value;
    (long-buildout | no-long-buildout);
    (loop-timing | no-loop-timing);
    loopback (local | payload | remote);
    (mac | no-mac);
    (payload-scrambler | no-payload-scrambler);
    start-end-flag (filler | shared);
}
traceoptions {
    flag flag <flag-modifier> <disable>;
}
transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
}
}
(traps | no-traps);
unidirectional;
vlan-tagging;
vlan-vci-tagging;
unit logical-unit-number {
    accept-source-mac {

```

```

    mac-address mac-address {
        policer {
            input cos-policer-name;
            output cos-policer-name;
        }
    }
}
accounting-profile name;
allow-any-vci;
atm-scheduler-map (map-name | default);
backup-options {
    interface interface-name;
}
bandwidth rate;
cell-bundle-size cells;
clear-dont-fragment-bit;
compression {
    rtp {
        f-max-period number;
        maximum-contexts number <force>;
        queues [ queue-numbers ];
        port {
            minimum port-number;
            maximum port-number;
        }
    }
}
compression-device interface-name;
copy-tos-to-outer-ip-header;
demux-destination family;
demux-source family;
demux-options {
    underlying-interface interface-name;
}
description text;
dial-options {
    l2tp-interface-id name;
    (dedicated | shared);
}
dialer-options {
    activation-delay seconds;
    callback;
    callback-wait-period time;
    deactivation-delay seconds;
    dial-string [ dial-string-numbers ];
    idle-timeout seconds;
    incoming-map {
        caller caller-id | accept-all;
        initial-route-check seconds;
        load-interval seconds;
        load-threshold percent;
        pool pool-name;
        redial-delay time;
        watch-list {
            [ routes ];
        }
    }
}

```

```

    }
}
disable;
disable-mlppp-inner-ppp-pfc;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
    activation-priority priority;
    bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
fragment-threshold bytes;
inner-vlan-id-range start start-id end end-id;
input-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
    down-count cells;
    up-count cells;
}
oam-period (seconds | disable);
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
    }
}

```

```

        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
        pap;
        default-pap-password password;
        local-name name;
        local-password password;
        passive;
    }
    dynamic-profile profile-name;
    lcp-max-conf-req number
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-max-conf-req number
    ncp-restart-timer milliseconds;
}
pppoe-options {
    access-concentrator name;
    auto-reconnect seconds;
    (client | server);
    service-name name;
    underlying-interface interface-name;
}
proxy-arp;
service-domain (inside | outside);
shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
        rate burst length);
    queue-length number;
}
short-sequence;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
    backup-destination address;
    destination address;
    key number;
    routing-instance {
        destination routing-instance-name;
    }
    source source-address;
    ttl number;
}
vci vpi-identifier.vci-identifier;
vci-range start start-vci end end-vci;
vpi vpi-identifier;
vlan-id number;
vlan-id-list [vlan-id vlan-id-vlan-id]
vlan-id-range number-number;
vlan-tags (Stacked VLAN Tags) inner tpid.vlan-id outer tpid.vlan-id;

```



```

vlan-tags outer <tpid.>vlan-id inner-list [vlan-id vlan-id-vlan-id]
family family {
    accounting {
        destination-class-usage;
        source-class-usage {
            direction;
        }
    }
    address address {
        destination address;
    }
    bundle ml-fpc-pic/port | ls-fpc/pic/port);
    filter {
        group filter-group-number;
        input filter-name;
        input-list {
            [ filter-names ];
            output filter-name;
        }
        output-list {
            [ filter-names ];
        }
    }
    ipsec-sa sa-name;
    keep-address-and-control;
    mtu bytes;
    multicast-only;
    negotiate-address;
    no-redirects;
    policer {
        arp policer-template-name;
        input policer-template-name;
        output policer-template-name;
    }
    primary;
    proxy inet-address address;
    receive-options-packets;
    receive-ttl-exceeded;
    remote (inet-address address | mac-address address);
    rpf-check <fail-filter filter-name> {
        <mode loose>;
    }
    sampling {
        direction;
    }
    service {
        input {
            service-set service-set-name <service-filter filter-name>;
            post-service-filter filter-name;
        }
        output {
            service-set service-set-names <service-filter filter-name>;
        }
    }
    (translate-discard-eligible | no-translate-discard-eligible);
    (translate-fecn-and-becn | no-translate-fecn-and-becn);

```

```
unnumbered-address interface-name <destination destination-profile  
    profile-name | preferred-source-address address>;  
address address {  
    arp ip-address (mac | multicast-mac) mac-address [<publish>;  
    broadcast address;  
    destination address;  
    destination-profile name;  
    eui-64;  
    multipoint-destination address (dlci dlci-identifier | vci vci-identifier);  
    multipoint-destination address {  
        epd-threshold cells plp1 cells;  
        inverse-arp;  
        oam-liveness {  
            up-count cells;  
            down-count cells;  
        }  
        oam-period (seconds | disable);  
        shaping {  
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate  
                sustained rate burst length);  
            queue-length number;  
        }  
        vci vpi-identifier.vci-identifier;  
    }  
    preferred;  
    primary;  
    (vrmp-group | vrmp-inet6-group) group-number {  
        (accept-data | no-accept-data);  
        advertise-interval seconds;  
        authentication-type authentication;  
        authentication-key key;  
        fast-interval milliseconds;  
        (preempt | no-preempt) {  
            hold-time seconds;  
        }  
        priority-number number;  
        track {  
            priority-cost seconds;  
            priority-hold-time interface-name {  
                bandwidth-threshold bits-per-second {  
                    priority;  
                }  
            }  
            interface priority;  
        }  
        route ip-address/mask routing-instance instance-name priority-cost  
            cost;  
    }  
    virtual-address [ addresses ];  
}  
  
}  
  
}
```

[edit logical-systems] Hierarchy Level

The following lists the statements that can be configured at the [edit logical-systems] hierarchy level that are also documented in this manual. For more information about logical systems, see the *JUNOS Routing Protocols Configuration Guide*.

```
logical-systems logical-system-name {
  interfaces interface-name {
    unit logical-unit-number {
      accept-source-mac {
        mac-address mac-address {
          policer {
            input cos-policer-name;
            output cos-policer-name;
          }
        }
      }
    }
  }
  allow-any-vci;
  atm-scheduler-map (map-name | default);
  bandwidth rate;
  backup-options {
    interface interface-name;
  }
  cell-bundle-size cells;
  clear-dont-fragment-bit;
  compression {
    rtp {
      f-max-period number;
      port {
        minimum port-number;
        maximum port-number;
      }
    }
    queues [ queue-numbers ];
  }
}
compression-device interface-name;
description text;
dial-options {
  l2tp-interface-id name;
  (dedicated | shared);
}
dialer-options {
  activation-delay seconds;
  deactivation-delay seconds;
  dial-string [ dial-string-numbers ];
  idle-timeout seconds;
  initial-route-check seconds;
  load-threshold number;
  pool pool;
  remote-name remote-callers;
  watch-list {
    [ routes ];
  }
}
```

```

}
disable;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
    activation-priority priority;
    bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
fragment-threshold bytes;
input-vlan-map {
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
    up-count cells;
    down-count cells;
}
oam-period (seconds | disable);
output-vlan-map {
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
}

```

```

compression {
    acfc;
    pfc;
}
dynamic-profile profile-name;
pap {
    default-pap-password password;
    local-name name;
    local-password password;
    passive;
}
proxy-arp;
service-domain (inside | outside);
shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
    rate burst length);
    queue-length number;
}
short-sequence;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
    backup-destination address;
    destination address;
    key number;
    routing-instance {
        destination routing-instance-name;
    }
    source source-address;
    ttl number;
}
vci vpi-identifier.vci-identifier;
vlan-id number;
vlan-id-list [vlan-id vlan-id-vlan-id]
vlan-tags (Stacked VLAN Tags) inner tpid.vlan-id outer tpid.vlan-id;
vlan-tags outer <tpid.>vlan-id inner-list [vlan-id vlan-id-vlan-id]
vpi vpi-identifier;
family family {
    accounting {
        destination-class-usage;
        source-class-usage {
            direction;
        }
    }
}
bundle interface-name;
filter {
    group filter-group-number;
    input filter-name;
    input-list {
        [filter-names ];
    }
    output filter-name;
}

```

```

    output-list {
        [ filter-names ];
    }
}
ipsec-sa sa-name;
keep-address-and-control;
llc2 {
    ack-delay-time time;
    ack-max count;
    idle-time time;
    local-window count;
    max-retry count;
    p-bit-timeout time;
    redundancy-group group-number {
        advertise-interval seconds;
        map {
            local-mac mac-address request mac-address;
        }
        preempt hold-time seconds;
        no-preempt;
        priority priority;
        track {
            dls {
                peer ip-address priority-cost priority;
                destination mac-address priority-cost priority;
            }
            interface interface-name priority-cost priority;
        }
    }
    t1-time time;
    t2-time time;
    trej-time time;
}
mtu bytes;
multicast-only;
no-redirects;
policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
}
primary;
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check <fail-filter filter-name> {
    <mode loose>;
}
sampling {
    direction;
}
service {
    input {
        service-set service-set-name <service-filter filter-name>;
        post-service-filter filter-name;
    }
}

```

```

    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
unnumbered-address interface-name destination address destination-profile
    profile-name;
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    multipoint-destination address (dlci dlci-identifier | vci vci-identifier);
    multipoint-destination address {
        epd-threshold cells plp1 cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (seconds | disable);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate
                sustained rate burst length);
            queue-length number;
        }
        vci vpi-identifier.vci-identifier;
    }
    preferred;
    primary;
    (vrrp-group | vrrp-inet6-group) group-number {
        (accept-data | no-accept-data);
        advertise-interval seconds;
        authentication-type authentication;
        authentication-key key;
        fast-interval milliseconds;
        (preempt | no-preempt) {
            hold-time seconds;
        }
        priority-number number;
        track {
            priority-cost seconds;
            priority-hold-time interface-name {
                interface priority;
                bandwidth-threshold bits-per-second {
                    priority;
                }
            }
            route ip-address/mask routing-instance instance-name priority-cost
                cost;
        }
    }
}
virtual-address [ addresses ];

```

```

    }
  }
}

```

[edit protocols connections] Hierarchy Level

The following statements can also be configured at the [edit logical-systems *logical-system-name* protocols connections] hierarchy level.

```

interface-switch connection-name {
  interface interface-name.unit-number;
  interface interface-name.unit-number;
}

```


[edit protocols dot1x] Hierarchy Level

```

dot1x {
  authenticator
    authentication-profile-name access-profile-name;
    interface interface-ids {
      maximum-requests integer;
      retries integer;
      quiet-period seconds;
      transmit-period seconds;
      reauthentication (disable | interval seconds);
      server-timeout seconds;
      supplicant (single);
      supplicant-timeout seconds;
    }
  }
}

```

[edit protocols lacp] Hierarchy Level

```

traceoptions {
  file filename <files number> <size size> <world-readable | no-world-readable>;
  flag flag <disable>;
}

```

[edit protocols oam] Hierarchy Level

```

ethernet {
  connectivity-fault-management {
    action-profile profile-name {
      default-action {
        interface-down;
      }
    }
  }
  linktrace {
    age (30m | 10m | 1m | 30s | 10s);
    path-database-size path-database-size;
  }
  maintenance-domain domain-name {
    bridge-domain name;
    routing-instance r1 {
      bridge-domain name;
      instance vpls-instance;
      interface (<ge> | <xe>) fpc/pic/port.domain;
      level number;
      maintenance-association name{
        mep identifier {
          direction (up | down)
          interface (ge | xe) fpc/pic/port.domain;
          auto-discovery;
          priority number;
        }
      }
    }
  }
}

```

```

    }
  }
  mip-half-function (none | default | explicit);
  name-format (character-string | none | dns | mac+2oct);
  short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
  continuity-check {
    hold-interval minutes;
    interval (10m | 10s | 1m | 1s | 100ms);
    loss-threshold number;
  }
  maintenance-association ma-name {
    mip-half-function (none | default | explicit);
    mep mep-id {
      auto-discovery;
      direction (up | down);
      interface interface-name;
      priority number;
      remote-mep mep-id {
        action-profile profile-name;
      }
    }
  }
}
}
performance-monitoring {
  hardware-assisted-timestamping;
}
}
link-fault-management {
  action-profile profile-name {
    action {
      syslog;
      link-down;
      send-critical-event;
    }
    event {
      link-adjacency-loss;
      link-event-rate {
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
      }
      protocol-down;
    }
  }
}
interface interface-name {
  apply-action-profile profile-name;
  event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
  }
  link-discovery (active | passive);
  negotiation-options {
    allow-remote-loopback;
  }
}

```

```

        no-allow-link-events;
    }
    pdu-interval interval;
    pdu-threshold threshold-value;
    remote-loopback;
}
}
}

```

[edit protocols ppp] Hierarchy Level

```

monitor-session (interface-name | all);
traceoptions {
    file filename <files number> <match regular-expression> <size size> <world-readable |
        no-world-readable> ;
    flag flag <disable>;
}

```

[edit protocols protection-group] Hierarchy Level

```

ethernet-ring ring-name {
    east-interface {
        control-channel channel-name {
            vlan number;
        }
    }
    guard-interval number;
    node-id mac-address;
    restore-interval number;
    ring-protection-link-owner;
    west-interface {
        control-channel channel-name {
            vlan number;
        }
    }
}

```

[edit protocols vrrp] Hierarchy Level

```

traceoptions {
    file <filename> <files number <match regular-expression <microsecond-stamp>
        <size size> <world-readable | no-world-readable>;
    flag flag;
}

```

[edit system processes] Hierarchy Level

```

dialer-services {
    disable;
}
isdn-signaling {
    disable;
}

```

```
    reject-incoming;  
}
```

Part 2

Router Interfaces Configuration Concepts

This part of the *Network Interfaces Configuration Guide* describes the various interface types and the processes used to configure them for typical usage.

- Router Interfaces Overview on page 31
- Configuring Physical Interface Properties on page 67
- Configuring Logical Interface Properties on page 143
- Configuring Protocol Family and Interface Address Properties on page 169
- Configuring Circuit and Translational Cross-Connects on page 223
- Tracing Interface Operations on page 241

Chapter 2

Router Interfaces Overview

Routers typically contain several different types of interfaces suited to various functions. For the interfaces on a router to function, you must configure them, specifying properties such as the interface location (that is, the slot in which the Flexible PIC Concentrator [FPC] or Dense Port Concentrator [DPC] is installed, and the location where the Physical Interface Card [PIC] is installed), the interface type (such as SONET/SDH, Asynchronous Transfer Mode [ATM], or Ethernet), encapsulation, and interface-specific properties. You can configure the interfaces that are currently present in the router, and you can also configure interfaces that are not currently present but that you might add in the future. When a configured interface appears, the JUNOS Software detects its presence and applies the appropriate configuration to it.

To determine which interfaces are currently installed in the router, issue the **show interfaces terse** operational mode command. If an interface is listed in the output, it is installed in the router. If an interface is not listed in the output, it is not installed in the router.

For information about which PICs are supported on your router, see your router's PIC guide.

You can configure JUNOS class-of-service (CoS) properties to provide a variety of classes of service for different applications, including multiple forwarding classes for managing packet transmission, congestion management, and CoS-based forwarding. For more information about configuring CoS properties, see the *JUNOS Class of Service Configuration Guide*.

This chapter discusses the following topics:

- Types of Interfaces on page 32
- Interface Encapsulations on page 39
- Interface Descriptors on page 50
- Interface Naming on page 51
- Displaying Interface Configurations on page 64
- Interface and Router Clock Sources on page 64

Types of Interfaces

Interfaces can be permanent or transient, and are used for networking or services:

- Permanent interfaces—Interfaces that are always present in the router.
- Transient interfaces—Interfaces that can be inserted into or removed from the router depending on your network configuration needs.
- Networking interfaces—Interfaces, such as Ethernet or SONET/SDH interfaces, that primarily provide traffic connectivity.
- Services interfaces—Interfaces that provide specific capabilities for manipulating traffic before it is delivered to its destination.
- Container interfaces—Interfaces that support APS on physical SONET links using a virtual container infrastructure.

The JUNOS Software internally generates nonconfigurable interfaces which are described in *Interfaces Command Reference* and *Services Interfaces*.

Permanent Interfaces

Permanent interfaces in the router consist of management Ethernet interfaces and internal Ethernet interfaces, as described in the following sections:



NOTE: The Routing Engines in the TX Matrix Plus router and in the T1600 routers configured in a routing matrix do not support the management Ethernet interface `fxp0` or the internal Ethernet interfaces `fxp1` or `fxp2`.

Management Ethernet Interfaces

The management Ethernet interface on the router provides an out-of-band method for connecting to the router. You can connect to the management interface over the network using utilities such as ssh and telnet. The Simple Network Management Protocol (SNMP) can use the management interface to gather statistics from the router.

- M Series, MX Series, and most T Series routers—For M Series and MX Series routers, and for T Series routers other than TX Matrix Plus routers or T1600 routers configured in a routing matrix, the JUNOS Software automatically creates the router's management Ethernet interface, `fxp0`. To use `fxp0` as a management port, you must configure its logical port, `fxp0.0`, with a valid IP address.
- TX Matrix Plus routers and T1600 routers in a routing matrix—For Juniper Networks TX Matrix Plus Routers and for T1600 Core Routers configured in a routing matrix, the JUNOS Software automatically creates the router's management Ethernet interface, `em0`. To use `em0` as a management port, you must configure its logical port, `em0.0`, with a valid IP address.



NOTE: Automated scripts that have been developed for standalone T1600 routers (T1600 routers not configured in a routing matrix) might contain references to the `fxp0`, `fxp1`, or `fxp2` interfaces. Before reusing the scripts on T1600 routers in a routing matrix, edit any command lines that reference the T1600 router management Ethernet interface `fxp0` by replacing “`fxp0`” with “`em0`.”

- J Series routers—For the Juniper Networks J Series Services Routers, you can use any of the built-in Ethernet ports as a management interface. To use a built-in interface as a management Ethernet interface, configure it with a valid IP address. The factory configuration for the J4350 and J6350 Services Routers automatically enables the J-Web user interface on the `ge-0/0/0`, `ge-0/0/1`, `ge-0/0/2`, and `ge-0/0/3` interfaces. To manually configure J-Web access, include the `interface interface-name` statement at the `[edit system services web-management http]` hierarchy level.

For information about establishing basic connectivity and configuring a management port, see the *Getting Started* guide for your router.

Internal Ethernet Interfaces

Internal Ethernet interfaces on the router provide communication between the Routing Engine and the Packet Forwarding Engine. The JUNOS Software boots the packet-forwarding component hardware. When these components are running, the Control Board uses the internal Ethernet interface to transmit hardware status information to the Routing Engine (the portion of the router running the JUNOS Software). Information transmitted includes the internal router temperature, the condition of the fans, whether an FPC has been removed or inserted, and information from the craft interface on the LCD panel. The internal Ethernet interface is configured automatically when the JUNOS Software boots.

- J Series, M Series, and MX Series routers and most T Series routers—For J Series, M Series, and MX Series routers, and for T Series routers other than TX Matrix Plus routers or T1600 routers configured in a routing matrix, the JUNOS Software creates the internal Ethernet interface `fxp1`. The internal Ethernet interface connects the Routing Engine `re0` (the portion of the router running the JUNOS Software) to the Packet Forwarding Engine. If the router has redundant Routing Engines, another internal Ethernet interface, `fxp2`, is created on each Routing Engine (`re0` and `re1`) in order to support fault tolerance. Two physical links between `re0` and `re1` connect the independent control planes. If one of the links fails, both Routing Engines can use the other link for IP communication.
- TX Matrix Plus routers—On a TX Matrix Plus router, the Routing Engine (RE-TXP-SFC) and Control Board (TXP-CB) function as a unit, or host subsystem. For each host subsystem in the router, the JUNOS Software automatically creates two internal Ethernet interfaces, `ixgbe0` and `ixgbe1`, for the two 10-Gigabit Ethernet ports on the Routing Engine.

The 10-Gigabit Ethernet port at the **ixgbe0** interface connects the TX Matrix Plus Routing Engine to the Routing Engines of every T1600 router configured in the routing matrix.

- The port connects the Routing Engine to a 10-Gigabit Ethernet switch on the local Control Board.
- The 10-Gigabit Ethernet switch connects the Control Board to a Gigabit Ethernet switch on the same local Control Board.
- The Gigabit Ethernet switch connects the Control Board to the remote Routing Engines of every T1600 router configured in the routing matrix.

If a TX Matrix Plus router contains redundant host subsystems, the independent control planes are connected by two physical links between the two 10-Gigabit Ethernet ports on their respective Routing Engines.

- The primary link to the remote Routing Engine is at the **ixgbe0** interface; the 10-Gigabit Ethernet switch on the local Control Board also connects the Routing Engine to the 10-Gigabit Ethernet port accessed by the **ixgbe1** interface on the remote Routing Engine.
- The alternate link to the remote Routing Engine is the 10-Gigabit Ethernet port at the **ixgbe1** interface. This second port connects the Routing Engine to the 10-Gigabit Ethernet switch on the remote Control Board, which connects to the 10-Gigabit Ethernet port at the **ixgbe0** interface on the remote Routing Engine.

If one of the two links between the host subsystems fails, both Routing Engines can use the other link for IP communication.

- T1600 routers in a routing matrix—On a T1600 router configured in a routing matrix, the Routing Engine (RE-TXP-LCC) and Control Board (LCC-CB) function as a unit, or host subsystem. For each host subsystem in the router, the JUNOS Software automatically creates two internal Ethernet interfaces, **bcm0** and **em1**, for the two Gigabit Ethernet ports on the Routing Engine.

The Gigabit Ethernet port at the **bcm0** interface connects the LCC Routing Engine to the Routing Engines of every other T1600 router configured in the routing matrix.

- The port connects the Routing Engine to a Gigabit Ethernet switch on the local Control Board.
- The switch connects the Control Board to the remote Routing Engines of every other T1600 router configured in the routing matrix.

If a T1600 router in a routing matrix contains redundant host subsystems, the independent control planes are connected by two physical links between the Gigabit Ethernet ports on their respective Routing Engines.

- The primary link to the remote Routing Engine is at the **bcm0** interface; the Gigabit Ethernet switch on the local Control Board also connects the Routing Engine to the Gigabit Ethernet port accessed by the **em1** interface on the remote Routing Engine.

- The alternate link to the remote Routing Engine is at the `em1` interface. This second port connects the Routing Engine to the Gigabit Ethernet switch on the remote Control Board, which connects to the Gigabit Ethernet port at the `bcm0` interface on the remote Routing Engine.

If one of the two links between the host subsystems fails, both Routing Engines can use the other link for IP communication.

Each router also has two serial ports, labeled *console* and *auxiliary*, for connecting tty type terminals to the router using standard PC-type tty cables. Although these ports are not network interfaces, they do provide access to the router.

Transient Interfaces

The M Series, MX Series, and T Series routers contain slots for installing FPCs. PICs can be installed in FPCs. The number of PICs that can be installed varies by router and type of FPC. The PICs provide the actual physical interfaces to the network. The MX Series routers contain slots for installing either DPC boards that provide the physical interfaces to the network or for installing FPCs in which PICs can be installed. These physical interfaces are transient interfaces of the router. They are referred to as transient because you can hot-swap a DPC or FPC and its PICs at any time.

You can insert any DPC or FPC into any slot that supports them in the appropriate router. Typically, you can place any combination of PICs, compatible with your router, in any location on an FPC. (You are limited by the total FPC bandwidth, and by the fact that some PICs physically require two or four of the PIC locations on the FPC. In some cases, power limitations or microcode limitations may also apply.) To determine DPC and PIC compatibility, see the *Hardware Guide*, *DPC Guide*, and *PIC Guide* for your router.

You must configure each transient interface based on the slot in which the FPC is installed, the location in which the PIC is installed, and for multiple port PICs, the port to which you are connecting.

You can configure the interfaces on PICs that are already installed in the router as well as interfaces on PICs that you plan to install later. The JUNOS Software detects which interfaces are actually present, so when the software activates its configuration, it activates only the present interfaces and retains the configuration information for the interfaces that are not present. When the JUNOS Software detects that an FPC containing PICs has been inserted into the router, the software activates the configuration for those interfaces.

Services Interfaces

Services interfaces enable you to incrementally add services to your network. The JUNOS Software supports the following services PICs:

- Adaptive Services (AS) PICs—Allow you to provide multiple services on a single PIC by configuring a set of services and applications. The AS PICs offer a special range of services you configure in one or more service sets.
- ES PIC—Provides a security suite for the IP version 4 (IPv4) and IP version 6 (IPv6) network layers. The suite provides functionality such as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. It also defines mechanisms for key generation and exchange, management of security associations, and support for digital certificates.
- Monitoring Services PICs—Enable you to monitor traffic flow and export the monitored traffic. Monitoring traffic allows you to gather and export detailed information about IPv4 traffic flows between source and destination nodes in your network; sample all incoming IPv4 traffic on the monitoring interface and present the data in cflowd record format; perform discard accounting on an incoming traffic flow; encrypt or tunnel outgoing cflowd records, intercepted IPv4 traffic, or both; and direct filtered traffic to different packet analyzers and present the data in its original format. On a Monitoring Services II PIC, you can configure either monitoring interfaces or collector interfaces. A collector interface allows you to combine multiple cflowd records into a compressed ASCII data file and export the file to an FTP server.
- Multilink Services, MultiServices, Link Services, and Voice Services PICs—Enable you to split, recombine, and sequence datagrams across multiple logical data links. The goal of multilink operation is to coordinate multiple independent links between a fixed pair of systems, providing a virtual link with greater bandwidth than any of the members.
- Tunnel Services PIC—By encapsulating arbitrary packets inside a transport protocol, tunneling provides a private, secure path through an otherwise public network. Tunnels connect discontinuous subnetworks and enable encryption interfaces, virtual private networks (VPNs), and Multiprotocol Label Switching (MPLS).
- On M Series and T Series routers, logical tunnel interfaces allow you to connect logical systems, virtual routers, or VPN instances. For more information about VPNs, see the *JUNOS VPNs Configuration Guide*. For more information about configuring tunnels, see the *JUNOS Services Interfaces Configuration Guide*.
- Services (J Series)—On J Series Services Routers, the It interface is an internal interface only and is not associated with a physical medium or PIM. You can configure the logical tunnel interface to provide class-of-service (CoS) support for data link switching (DLSw) traffic and real-time performance monitoring (RPM) probe packets. For more information, see the *J-series Services Router Basic LAN and WAN Access Configuration Guide*.



NOTE: The It interface on the J Series router does not support logical systems.

Container Interfaces

Container interfaces provide the following features:

- APS on SONET links are supported using container infrastructure.
- Container physical interfaces and logical interfaces remain up on switchover.
- APS parameters are auto-copied from the container interface to the member links.

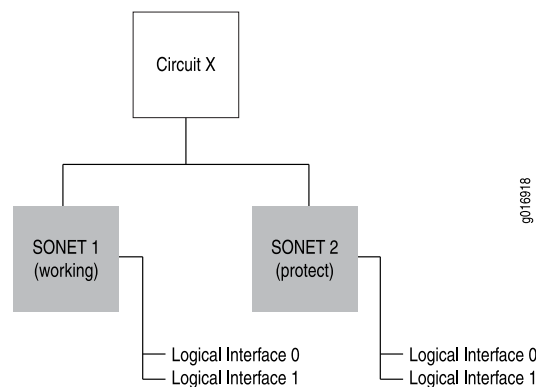


NOTE: Paired groups and true unidirectional APS are not currently supported.

Traditional APS Concept

Traditional APS is configured on two independent physical SONET interfaces: one configured as the working circuit and the other as the protect circuit (see Figure 1 on page 37). The circuit, named Circuit X in the figure, is the link between the two SONET interfaces.

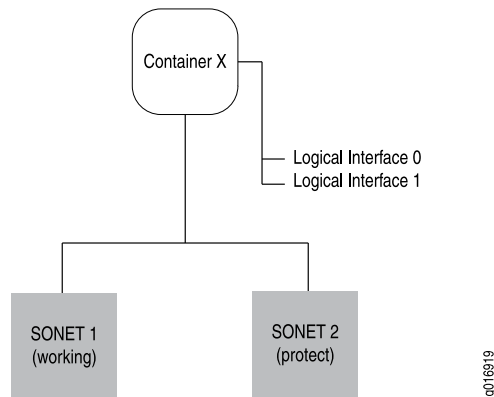
Figure 1: APS Interface



Traditional APS uses routing protocols that run on each individual SONET interface (since circuit is an abstract construct, instead of being an actual interface). When the working link goes down, the APS infrastructure brings up the protect link and its underlying logical interfaces, and brings down the working link and its underlying logical interfaces, causing the routing protocols to reconverge. This consumes time and leads to traffic loss even though the APS infrastructure has performed the switch quickly.

Container Interfaces Concept

To solve this problem, the JUNOS Software provides a soft interface construct called a container interface (see Figure 2 on page 38).

Figure 2: Container Interface

The container interface allows routing protocols to run on the logical interfaces associated with a virtual *container interface* instead of on the physical SONET interfaces. When APS switches the underlying physical link based on a fault condition, the container interface remains up, and the logical interface on the container interface does not flap. The routing protocols remain unaware of the APS switching.

APS Support for Container-Based Interfaces

With the container interface, APS is configured on the container interface itself. Individual member SONET links are either marked as primary (corresponding to the working circuit) or standby (corresponding to the protect circuit) in the configuration. No circuit or group name is specified in the container interface model; physical SONET links are put in an APS group by linking them to a single container interface. APS parameters are specified at the container interface level, and are propagated to the individual SONET links by the APS daemon.

Autocopy of APS Parameters

Typical applications require copying APS parameters from the working circuit to the protect circuit, since most of the parameters must be the same for both circuits. This is automatically done in the container interface. APS parameters are specified only once under the container physical interface configuration, and are internally copied over to the individual physical SONET links.

For more information, see “Configuring Container Interfaces” on page 863.

Interface Encapsulations

Table 3 on page 39 lists encapsulation support by interface type.

Table 3: Encapsulation Support by Interface Type

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
ae—Aggregated Ethernet interface	ethernet-ccc—Ethernet cross-connect	dix—Ethernet DIXv2 (RFC 894)
	extended-vlan-ccc—Nonstandard TPID tagging for a cross-connect	vlan-ccc—802.1Q tagging for a cross-connect
	extended-vlan-vpls—Extended VLAN virtual private LAN service	
	vlan-ccc—802.1Q tagging for a cross-connect	
	ethernet-vpls—Ethernet virtual private LAN service	
	vlan-vpls—VLAN virtual private LAN service	
as—Aggregated SONET/SDH interface	cisco-hdlc—Cisco-compatible HDLC framing	NA
	ppp—Serial PPP device	
at—ATM1 interface	atm-ccc-cell-relay—ATM cell relay encapsulation for a cross-connect	atm-ccc-cell-relay—ATM cell relay for CCC
	atm-pvc—ATM permanent virtual circuits	atm-ccc-vc-mux—ATM VC for CCC
	ethernet-over-atm—Ethernet over ATM encapsulation	atm-cisco-nlpid—Cisco-compatible ATM NLPID encapsulation
		atm-nlpid—ATM NLPID encapsulation
		atm-snap—ATM LLC/SNAP encapsulation
		atm-tcc-snap—ATM LLC/SNAP for a translational cross-connect
		atm-tcc-vc-mux—ATM VC for a translational cross-connect
		atm-vc-mux—ATM VC multiplexing
		ether-over-atm-llc—Ethernet over ATM (LLC/SNAP) encapsulation

Table 3: Encapsulation Support by Interface Type *(continued)*

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
at—ATM2 intelligent queuing (IQ) interface	atm-ccc-cell-relay—ATM cell relay encapsulation for a cross-connect	atm-ccc-cell-relay—ATM cell relay for CCC
	atm-pvc—ATM permanent virtual circuits	atm-ccc-vc-mux—ATM VC for CCC
	ethernet-over-atm—Ethernet over ATM encapsulation	atm-cisco-nlpid—Cisco-compatible ATM NLPID encapsulation
		atm-mlppp-llc—ATM MLPPP over AAL5/LLC
		atm-nlpid—ATM NLPID encapsulation
		atm-ppp-llc—ATM PPP over AAL5/LLC
		atm-ppp-vc-mux—ATM PPP over raw AAL5
		atm-snap—ATM LLC/SNAP encapsulation
		atm-tcc-snap—ATM LLC/SNAP for a translational cross-connect
		atm-tcc-vc-mux—ATM VC for a translational cross-connect
bcm—Gigabit Ethernet internal interfaces	NA	NA
	NA	NA
	NA	NA
	NA	NA
br—Integrated Services Digital Network (ISDN) interface	NA	NA
ci—Container interface	cisco-hdlc—Cisco-compatible HDLC framing	aps—SONET interface required for APS configuration.
	ppp—Serial PPP device	

Table 3: Encapsulation Support by Interface Type *(continued)*

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
ds—DS0 interface	cisco-hdlc—Cisco-compatible HDLC framing	frame-relay-ccc—Frame Relay DLCI for CCC
	cisco-hdlc-ccc—Cisco-compatible HDLC framing for a cross-connect	frame-relay-ppp—PPP over Frame Relay
	cisco-hdlc-tcc—Cisco-compatible HDLC framing for a translational cross-connect	frame-relay-tcc—Frame Relay DLCI for a translational cross-connect
	extended-frame-relay-ccc—Any Frame Relay DLCI for a cross-connect	
	extended-frame-relay-tcc—Any Frame Relay DLCI for a translational cross-connect	
	flexible-frame-relay—Multiple Frame Relay encapsulations	
	frame-relay—Frame Relay encapsulation	
	frame-relay-ccc—Frame Relay for a cross-connect	
	frame-relay-port-ccc—Frame Relay port encapsulation for a cross-connect	
	frame-relay-tcc—Frame Relay for a translational cross-connect	
	multilink-frame-relay-uni-nni—Multilink Frame Relay UNI NNI (FRF.16) encapsulation	
	ppp—Serial PPP device	
	ppp-ccc—Serial PPP device for a cross-connect	
	ppp-tcc—Serial PPP device for a translational cross-connect	
dsc—Discard interface	NA	NA

Table 3: Encapsulation Support by Interface Type *(continued)*

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
e1—E1 interface (including channelized STM1-to-E1 interfaces)	cisco-hdlc—Cisco-compatible HDLC framing	frame-relay-ccc—Frame Relay DLCI for CCC
	cisco-hdlc-ccc—Cisco-compatible HDLC framing for a cross-connect	frame-relay-ppp—PPP over Frame Relay
	cisco-hdlc-tcc—Cisco-compatible HDLC framing for a translational cross-connect	frame-relay-tcc—Frame Relay DLCI for a translational cross-connect
	extended-frame-relay-ccc—Any Frame Relay DLCI for a cross-connect	
	extended-frame-relay-tcc—Any Frame Relay DLCI for a translational cross-connect	
	flexible-frame-relay—Multiple Frame Relay encapsulations	
	frame-relay—Frame Relay encapsulation	
	frame-relay-ccc—Frame Relay for a cross-connect	
	frame-relay-port-ccc—Frame Relay port encapsulation for a cross-connect	
	frame-relay-tcc—Frame Relay for a translational cross-connect	
	multilink-frame-relay-uni-nni—Multilink Frame Relay UNI NNI (FRF.16) encapsulation	
	ppp—Serial PPP device	
	ppp-ccc—Serial PPP device for a cross-connect	
	ppp-tcc—Serial PPP device for a translational cross-connect	

Table 3: Encapsulation Support by Interface Type *(continued)*

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
e3—E3 interface (including E3 IQ and IQE interfaces)	cisco-hdlc—Cisco-compatible HDLC framing	frame-relay-ccc—Frame Relay DLCI for CCC
	cisco-hdlc-ccc—Cisco-compatible HDLC framing for a cross-connect	frame-relay-ppp—PPP over Frame Relay
	cisco-hdlc-tcc—Cisco-compatible HDLC framing for a translational cross-connect	frame-relay-tcc—Frame Relay DLCI for a translational cross-connect
	extended-frame-relay-ccc—Any Frame Relay DLCI for a cross-connect	
	extended-frame-relay-tcc—Any Frame Relay DLCI for a translational cross-connect	
	flexible-frame-relay—Multiple Frame Relay encapsulations	
	frame-relay—Frame Relay encapsulation	
	frame-relay-ccc—Frame Relay for a cross-connect	
	frame-relay-port-ccc—Frame Relay port encapsulation for a cross-connect	
	frame-relay-tcc—Frame Relay for a translational cross-connect	
	ppp—Serial PPP device	
	ppp-ccc—Serial PPP device for a cross-connect	
	ppp-tcc—Serial PPP device for a translational cross-connect	
em—Management and internal Ethernet interfaces	NA	NA

Table 3: Encapsulation Support by Interface Type *(continued)*

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
fe—Fast Ethernet interface	ethernet-ccc—Ethernet cross-connect	dix—Ethernet DIXv2 (RFC 894)
	ethernet-tcc—Ethernet translational cross-connect	vlan-ccc—802.1Q tagging for a cross-connect
	ethernet-vpls—Ethernet virtual private LAN service	vlan-vpls—VLAN virtual private LAN service
	extended-vlan-ccc—Nonstandard TPID tagging for a cross-connect	
	extended-vlan-tcc—802.1Q tagging for a translational cross-connect	
	extended-vlan-vpls—Extended VLAN virtual private LAN service	
	vlan-ccc—802.1Q tagging for a cross-connect	
	vlan-vpls—VLAN virtual private LAN service	
fxp—Management and internal Ethernet interfaces	NA	NA
ge—Gigabit Ethernet interface (including Gigabit Ethernet IQ interfaces)	ethernet-ccc—Ethernet cross-connect	dix—Ethernet DIXv2 (RFC 894)
	ethernet-tcc—Ethernet translational cross-connect	vlan-ccc—802.1Q tagging for a cross-connect
	ethernet-vpls—Ethernet virtual private LAN service	vlan-tcc—802.1Q tagging for a translational cross-connect
	extended-vlan-ccc—Nonstandard TPID tagging for a cross-connect	vlan-vpls—VLAN virtual private LAN service
	extended-vlan-tcc—802.1Q tagging for a translational cross-connect	
	extended-vlan-vpls—Extended VLAN virtual private LAN service	
	flexible-ethernet-services—Allows per-unit Ethernet encapsulation configuration	
	vlan-ccc—802.1Q tagging for a cross-connect	
	vlan-vpls—VLAN virtual private LAN service	
ixgbe—10-Gigabit Ethernet internal interfaces	NA	NA

Table 3: Encapsulation Support by Interface Type *(continued)*

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
lo—Loopback interface; the JUNOS Software automatically configures one loopback interface (lo0)	NA	NA
ls—Link services interface	multilink-frame-relay-uni-nni—Multilink Frame Relay UNI NNI (FRF.16) encapsulation	multilink-frame-relay-end-to-end—Multilink Frame Relay end-to-end (FRF.15) multilink-ppp—Multilink PPP
lsq—Link services IQ interface	multilink-frame-relay-uni-nni—Multilink Frame Relay UNI NNI (FRF.16) encapsulation	multilink-frame-relay-end-to-end—Multilink Frame Relay end-to-end (FRF.15) multilink-ppp—Multilink PPP
lt—Logical tunnel interface	NA	ethernet—Ethernet service ethernet-vpls—Ethernet virtual private LAN service ethernet-ccc—Ethernet cross-connect frame-relay—Frame Relay encapsulation frame-relay-ccc—Frame Relay for a cross-connect vlan—VLAN service vlan-ccc—802.1Q tagging for a cross-connect vlan-vpls—VLAN virtual private LAN service
ml—Multilink interface (including Multilink Frame Relay and MLPPP)	NA	multilink-frame-relay-end-to-end—Multilink Frame Relay end-to-end (FRF.15) multilink-ppp—Multilink PPP

Table 3: Encapsulation Support by Interface Type (*continued*)

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
se—Serial interface (including EIA-530, V.35, and X.21 interfaces)	cisco-hdlc—Cisco-compatible HDLC framing	frame-relay-ccc—Frame Relay DLCI for CCC
	cisco-hdlc-ccc—Cisco-compatible HDLC framing for a cross-connect	frame-relay-ppp—PPP over Frame Relay
	cisco-hdlc-tcc—Cisco-compatible HDLC framing for a translational cross-connect	frame-relay-tcc—Frame Relay DLCI for a translational cross-connect
	frame-relay—Frame Relay encapsulation	
	frame-relay-ccc—Frame Relay for a cross-connect	
	frame-relay-port-ccc—Frame Relay port encapsulation for a cross-connect	
	frame-relay-tcc—Frame Relay for a translational cross-connect	
	ppp—Serial PPP device	
	ppp-ccc—Serial PPP device for a cross-connect	
	ppp-tcc—Serial PPP device for a translational cross-connect	

Table 3: Encapsulation Support by Interface Type *(continued)*

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
so—SONET/SDH interface	cisco-hdlc—Cisco-compatible HDLC framing	frame-relay-ccc—Frame Relay DLCI for CCC
	cisco-hdlc-ccc—Cisco-compatible HDLC framing for a cross-connect	frame-relay-ppp—PPP over Frame Relay
	cisco-hdlc-tcc—Cisco-compatible HDLC framing for a translational cross-connect	frame-relay-tcc—Frame Relay DLCI for a translational cross-connect
	extended-frame-relay-ccc—Any Frame Relay DLCI for a cross-connect	multilink-frame-relay-end-to-end—IQE SONET PICs support Multilink Frame Relay end-to-end (FRF.15)
	extended-frame-relay-tcc—Any Frame Relay DLCI for a translational cross-connect	multilink-ppp—IQE SONET PICs support Multilink PPP
	flexible-frame-relay—Multiple Frame Relay encapsulations	
	frame-relay—Frame Relay encapsulation	
	frame-relay-ccc—Frame Relay for a cross-connect	
	frame-relay-port-ccc—Frame Relay port encapsulation for a cross-connect	
	frame-relay-tcc—Frame Relay for a translational cross-connect	
	ppp—Serial PPP device	
	ppp-ccc—Serial PPP device for a cross-connect	
	ppp-tcc—Serial PPP device for a translational cross-connect	

Table 3: Encapsulation Support by Interface Type *(continued)*

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
t1—T1 interface (including channelized DS3-to-DS1 interfaces)	cisco-hdlc—Cisco-compatible HDLC framing	frame-relay-ccc—Frame Relay DLCI for CCC
	cisco-hdlc-ccc—Cisco-compatible HDLC framing for a cross-connect	frame-relay-ppp—PPP over Frame Relay
	cisco-hdlc-tcc—Cisco-compatible HDLC framing for a translational cross-connect	frame-relay-tcc—Frame Relay DLCI for a translational cross-connect
	extended-frame-relay-ccc—Any Frame Relay DLCI for a cross-connect	
	extended-frame-relay-tcc—Any Frame Relay DLCI for a translational cross-connect	
	flexible-frame-relay—Multiple Frame Relay encapsulations	
	frame-relay—Frame Relay encapsulation	
	frame-relay-ccc—Frame Relay for a cross-connect	
	frame-relay-port-ccc—Frame Relay port encapsulation for a cross-connect	
	frame-relay-tcc—Frame Relay for a translational cross-connect	
	multilink-frame-relay-uni-nni—Multilink Frame Relay UNI NNI (FRF.16) encapsulation	
	ppp—Serial PPP device	
	ppp-ccc—Serial PPP device for a cross-connect	
	ppp-tcc—Serial PPP device for a translational cross-connect	

Table 3: Encapsulation Support by Interface Type (*continued*)

Interface Type	Physical Interface Encapsulation	Logical Interface Encapsulation
t3—T3 interface (including channelized OC12-to-DS3 interfaces)	cisco-hdlc—Cisco-compatible HDLC framing	frame-relay-ccc—Frame Relay DLCI for CCC
	cisco-hdlc-ccc—Cisco-compatible HDLC framing for a cross-connect	frame-relay-ppp—PPP over Frame Relay
	cisco-hdlc-tcc—Cisco-compatible HDLC framing for a translational cross-connect	frame-relay-tcc—Frame Relay DLCI for a translational cross-connect
	extended-frame-relay-ccc—Any Frame Relay DLCI for a cross-connect	
	extended-frame-relay-tcc—Any Frame Relay DLCI for a translational cross-connect	
	flexible-frame-relay—Multiple Frame Relay encapsulations	
	frame-relay—Frame Relay encapsulation	
	frame-relay-ccc—Frame Relay for a cross-connect	
	frame-relay-port-ccc—Frame Relay port encapsulation for a cross-connect	
	frame-relay-tcc—Frame Relay for a translational cross-connect	
Controller-level channelized IQ interfaces (cau4, coc1, coc3, coc12, cstm1, ct1, ct3, ce1)	ppp—Serial PPP device	
	ppp-ccc—Serial PPP device for a cross-connect	
	ppp-tcc—Serial PPP device for a translational cross-connect	
Services interfaces (cp, gr, ip, mo, vt, es, mo, rsp, sp)	NA	NA
Unconfigurable, internally generated interfaces (gre, ipip, learning-chip (lc), lsi, tap, mt, mtun, pd, pe, pimd, pime)	NA	NA

Interface Descriptors

When you configure an interface, you are effectively specifying the properties for a physical interface descriptor. In most cases, the physical interface descriptor corresponds to a single physical device and consists of the following parts:

- The interface name, which defines the media type
- The slot in which the FPC or DPC is located
- The location on the FPC in which the PIC is installed
- The PIC or DPC port
- The interface's channel and logical unit numbers (optional)

Each physical interface descriptor can contain one or more logical interface descriptors. These allow you to map one or more logical (or virtual) interfaces to a single physical device. Creating multiple logical interfaces is useful for ATM, Frame Relay, and Gigabit Ethernet networks, in which you can associate multiple virtual circuits, data-link connections, or virtual LANs (VLANs) with a single interface device.

Each logical interface descriptor can have one or more family descriptors to define the protocol family that is associated with and allowed to run over the logical interface.

The following protocol families are supported:

- Internet Protocol version 4 (IPv4) suite (inet)
- Internet Protocol version 6 (IPv6) suite (inet6)
- Circuit cross-connect (CCC)
- Translational cross-connect (TCC)
- International Organization for Standardization (ISO)
- Multilink Frame Relay end-to-end (MLFR end-to-end)
- Multilink Frame Relay user-to-network interface network-to-network interface (MLFR UNI NNI)
- Multilink Point-to-Point Protocol (MLPPP)
- Multiprotocol Label Switching (MPLS)
- Trivial Network Protocol (TNP)
- (M Series, T Series, and MX Series routers only) Virtual private LAN service (VPLS)

Finally, each family descriptor can have one or more address entries, which associate a network address with a logical interface and hence with the physical interface.

You configure the various interface descriptors as follows:

- You configure the physical interface descriptor by including the **interfaces** *interface-name* statement.
- You configure the logical interface descriptor by including the **unit** statement within the **interfaces** *interface-name* statement or by including the *.logical* descriptor at the end of the interface name, as in **t3-0/0/0.1**, where the logical unit number is 1, as shown in the following examples:

```
[edit]
user@host# set interfaces t3-0/0/0 unit 1
[edit]
user@host# edit interfaces t3-0/0/0.1
[edit interfaces t3-0/0/0]
user@host# set unit 1
```

- You configure the family descriptor by including the **family** statement within the **unit** statement.
- You configure address entries by including the **address** statement within the **family** statement.
- You configure tunnels by including the **tunnel** statement within the **unit** statement.

Interface Naming

Each interface has an interface name, which specifies the media type, the slot in which the FPC or DPC is located, the location on the FPC where the PIC is installed, and the PIC or DPC port. The interface name uniquely identifies an individual network connector in the system. You use the interface name when configuring interfaces and when enabling various functions and properties, such as routing protocols, on individual interfaces. The system uses the interface name when displaying information about the interface, for example, in the **show interfaces** command.

The interface name is represented by a physical part, a channel part, and a logical part in the following format:

physical<:channel>.logical

The channel part of the name is optional for all interfaces except channelized DS3, E1, OC12, and STM1 interfaces.

The following sections provide interface naming configuration guidelines:

- Physical Part of an Interface Name on page 52
- Logical Part of an Interface Name on page 56
- Separators in an Interface Name on page 56
- Channel Part of an Interface Name on page 56
- Interface Naming for a Routing Matrix Based on a TX Matrix Router on page 57
- Interface Naming for a Routing Matrix Based on a TX Matrix Plus Router on page 59

- Chassis Interface Naming on page 61
- Examples: Interface Naming on page 62

Physical Part of an Interface Name

The physical part of an interface name identifies the physical device, which corresponds to a single physical network connector. This part of the interface name has the following format:

type-fpc/pic/port

type is the media type, which identifies the network device where it can be one of the following:

- **ae**—Aggregated Ethernet interface. This is a virtual aggregated link and has a different naming format from most PICs; for more information, see “Configuring Aggregated Ethernet Interfaces” on page 623.
- **as**—Aggregated SONET/SDH interface. This is a virtual aggregated link and has a different naming format from most PICs; for more information, see “Configuring Aggregated SONET/SDH Interfaces” on page 881.
- **at**—ATM1 or ATM2 intelligent queuing (IQ) interface or a virtual ATM interface on a Circuit Emulation (CE) interface.
- **bcm**—Gigabit Ethernet internal interface
- **br**—Integrated Services Digital Network (ISDN) interface (configured on a 1-port or 4-port Basic Rate Interface (BRI) card). This interface has a different naming format from most PICs: **br-pim/0/port**. The second number is always 0. For more information, see “Configuring ISDN Physical Interface Properties” on page 821.
- **cau4**—Channelized AU-4 IQ interface (configured on the Channelized STM1 IQ or IQE PIC or Channelized OC12 IQ and IQE PICs).
- **ce1**—Channelized E1 IQ interface (configured on the Channelized E1 IQ PIC or Channelized STM1 IQ or IQE PIC).
- **ci**—Container interface.
- **coc1**—Channelized OC1 IQ interface (configured on the Channelized OC12 IQ and IQE or Channelized OC3 IQ and IQE PICs).
- **coc3**—Channelized OC3 IQ interface (configured on the Channelized OC3 IQ and IQE PICs).
- **coc12**—Channelized OC12 IQ interface (configured on the Channelized OC12 IQ and IQE PICs).
- **coc48**—Channelized OC48 interface (configured on the Channelized OC48 and Channelized OC48 IQE PICs).
- **cp**—Collector interface (configured on the Monitoring Services II PIC).
- **cstm1**—Channelized STM1 IQ interface (configured on the Channelized STM1 IQ or IQE PIC).
- **cstm4**—Channelized STM4 IQ interface (configured on the Channelized OC12 IQ and IQE PICs).

- **cstm16**—Channelized STM16 IQ interface (configured on the Channelized OC48/STM16 and Channelized OC48/STM16 IQE PICs).
- **ct1**—Channelized T1 IQ interface (configured on the Channelized DS3 IQ and IQE PICs, Channelized OC3 IQ and IQE PICs, Channelized OC12 IQ and IQE PICs, or Channelized T1 IQ PIC).
- **ct3**—Channelized T3 IQ interface (configured on the Channelized DS3 IQ and IQE PICs, Channelized OC3 IQ and IQE PICs, or Channelized OC12 IQ and IQE PICs).
- **demux**—Interface that supports logical IP interfaces that use the IP source or destination address to demultiplex received packets. Only one demux interface (**demux0**) exists per chassis. All demux logical interfaces must be associated with an underlying logical interface.
- **dfc**—Interface that supports dynamic flow capture processing on T Series or M320 routers containing one or more Monitoring Services III PICs. Dynamic flow capture enables you to capture packet flows on the basis of dynamic filtering criteria. Specifically, you can use this feature to forward passively monitored packet flows that match a particular filter list to one or more destinations using an on-demand control protocol.
- **ds**—DS0 interface (configured on the Multichannel DS3 PIC, Channelized E1 PIC, Channelized OC3 IQ and IQE PICs, Channelized OC12 IQ and IQE PICs, Channelized DS3 IQ and IQE PICs, Channelized E1 IQ PIC, Channelized STM1 IQ or IQE PIC, or Channelized T1 IQ).
- **dsc**—Discard interface.
- **e1**—E1 interface (including channelized STM1-to-E1 interfaces).
- **e3**—E3 interface (including E3 IQ interfaces).
- **em**—Management and internal Ethernet interfaces for TX Matrix Plus routers and T1600 routers in a TX Matrix Plus routing matrix. The JUNOS Software automatically configures the router's management interface, **em0**, which is an out-of-band management interface, and the internal Ethernet interface **em1**, which connects the Routing Engine with the router's packet-forwarding components. If the router has redundant Routing Engines, the JUNOS Software configures another internal Ethernet interface, **em2**, on each Routing Engine (**re0** and **re1**) in order to support fault tolerance. Two physical links between **re0** and **re1** connect the independent control planes. If one of the links fails, both Routing Engines can use the other link for IP communication.
- **es**—Encryption interface.
- **fe**—Fast Ethernet interface.
- **fxp**—Management and internal Ethernet interfaces for J Series, M Series, and MX Series routers, and for T Series routers other than the TX Matrix Plus router and T1600 routers configured in a routing matrix. The JUNOS Software automatically configures the router's management Ethernet interface, **fxp0**, which is an out-of-band management interface, and the internal Ethernet interface, **fxp1**, which connects the Routing Engine with the router's packet-forwarding components. If the router has redundant Routing Engines, another internal Ethernet interface, **fxp2**, is created on each Routing Engine (**re0** and **re1**) in order to support fault tolerance. Two physical links between **re0** and **re1** connect the

independent control planes. If one of the links fails, both Routing Engines can use the other link for IP communication.

- **ge**—Gigabit Ethernet interface. Some older 10-Gigabit Ethernet interfaces use the **ge** media type to identify the physical part of the network device, but newer 10-Gigabit Ethernet interfaces use the **xe** media type.
- **gr**—Generic routing encapsulation (GRE) tunnel interface.
- **gre**—Internally generated interface that is configurable only as the control channel for Generalized MPLS (GMPLS). For more information about GMPLS, see the *JUNOS MPLS Applications Configuration Guide* and the *JUNOS Feature Guide*.
- **ip**—IP-over-IP encapsulation tunnel interface.
- **ipip**—Internally generated interface that is not configurable.
- **ixgbe**—10-Gigabit Ethernet internal interface
- **lc**—Internally generated interface that is not configurable.
- **lo**—Loopback interface. The JUNOS Software automatically configures one loopback interface (**lo0**). The logical interface **lo0.16383** is a nonconfigurable interface for router control traffic.
- **ls**—Link services interface.
- **lsi**—Internally generated interface that is not configurable.
- **ml**—Multilink interface (including Multilink Frame Relay and MLPPP).
- **mo**—Monitoring services interface (including monitoring services and monitoring services II). The logical interface **mo-fpc/pic/port.16383** is an internally generated, nonconfigurable interface for router control traffic.
- **ms**—MultiServices interface.
- **mt**—Multicast tunnel interface (internal router interface for VPNs). If your router has a Tunnel PIC, the JUNOS Software automatically configures one multicast tunnel interface (**mt**) for each virtual private network (VPN) you configure. Although it is not necessary to configure multicast interfaces, you can use the **multicast-only** statement to configure the unit and family so that the tunnel can transmit and receive multicast traffic only. For more information, see **multicast-only**.
- **mtun**—Internally generated interface that is not configurable.
- **oc3**—OC3 IQ interface (configured on the Channelized OC12 IQ and IQE PICs or Channelized OC3 IQ and IQE PICs).
- **pd**—Interface on the rendezvous point (RP) that de-encapsulates packets.
- **pe**—Interface on the first-hop RP that encapsulates packets destined for the RP router.
- **pimd**—Internally generated interface that is not configurable.
- **pime**—Internally generated interface that is not configurable.
- **rlsq**—Container interface, numbered from 0 through 127, used to tie the primary and secondary LSQ PICs together in high availability configurations. Any failure of the primary PIC results in a switch to the secondary PIC and vice versa.

- **rms**—Redundant interface for two MultiServices interfaces.
- **rsp**—Redundant virtual interface for the adaptive services interface.
- **se**—Serial interface (including EIA-530, V.35, and X.21 interfaces).
- **so**—SONET/SDH interface.
- **sp**—Adaptive services interface. The logical interface **sp-fpc/pic/port.16383** is an internally generated, nonconfigurable interface for router control traffic.
- **stm1**—STM1 interface (configured on the OC3/STM1 interfaces).
- **stm4**—STM4 interface (configured on the OC12/STM4 interfaces).
- **stm16**—STM16 interface (configured on the OC48/STM16 interfaces).
- **t1**—T1 interface (including channelized DS3-to-DS1 interfaces).
- **t3**—T3 interface (including channelized OC12-to-DS3 interfaces).
- **tap**—Internally generated interface that is not configurable.
- **umd**—USB modem interface.
- **vsp**—Voice services interface.
- **vc4**—Virtually concatenated interface.
- **vt**—Virtual loopback tunnel interface.
- **xe**—10-Gigabit Ethernet interface. Some older 10-Gigabit Ethernet interfaces use the **ge** media type (rather than **xe**) to identify the physical part of the network device.
- **xt**—Logical interface for Protected System Domains to establish a Layer 2 tunnel connection.

fpc identifies the number of the FPC or DPC card on which the physical interface is located. Specifically, it is the number of the slot in which the card is installed.

M40, M40e, M160, M320, M120, T320, T640, and T1600 routers each have eight FPC slots that are numbered 0 through 7, from left to right as you are facing the front of the chassis. For information about compatible FPCs and PICs, see the *Hardware Guide* for your router.

The M20 router has four FPC slots that are numbered 0 through 3, from top to bottom as you are facing the front of the chassis. The slot number is printed adjacent to each slot.

MX Series routers support DPCs and FPCs. Each FPC occupies two DPC slots and combinations are permitted. The MX960 router has slots that support up to 12 DPCs or up to 6 FPCs (without redundant SCB) or 5 FPCs (with redundant SCB). The MX480 router has slots that support up to three FPCs. The MX240 router has slots that support up to three DPCs or one FPC. FPCs use the lower DPC slot number for FPC slot numbering. For information about compatible FPCs and PICs, see the *MX Series PIC Guide*. For information about DPCs, see the *MX Series DPC Guide*.

MX Series routers support Type 2 and Type 3 FPCs. On MX Series routers, the Type 2 and Type 3 FPCs support Type 2 and Type 3 SONET/SDH PICs, respectively. Type

1 PICs are not supported on MX Series routers. For a complete list of supported PICs, see the *MX Series Hardware Guide*.

For M5, M7i, M10, and M10i routers, the FPCs are built into the chassis; you install the PICs into the chassis.

The M5 and M7i routers have space for up to four PICs. The M7i router also comes with an integrated Tunnel PIC, or an optional integrated AS PIC, or an optional integrated MS PIC.

The M10 and M10i routers have space for up to eight PICs.

A routing matrix can have up to 32 FPCs (numbered 0 through 31).

For more information about interface naming for a routing matrix, see “Interface Naming for a Routing Matrix Based on a TX Matrix Router” on page 57.

pic identifies the number of the PIC on which the physical interface is located. Specifically, it is the number of the PIC location on the FPC. FPCs with four PIC slots are numbered 0 through 3. FPCs with three PIC slots are numbered 0 through 2. The PIC location is printed on the FPC carrier board. For PICs that occupy more than one PIC slot, the lower PIC slot number identifies the PIC location.

port identifies a specific port on a PIC or DPC. The number of ports varies depending on the PIC. The port numbers are printed on the PIC.

Logical Part of an Interface Name

The logical unit part of the interface name corresponds to the logical unit number, which can be a number from 0 through 16,384.

Separators in an Interface Name

There is a separator between each element of an interface name.

In the physical part of the name, a hyphen (-) separates the media type from the FPC number, and a slash (/) separates the FPC, PIC, and port numbers.

In the virtual part of the name, a period (.) separates the channel and logical unit numbers.

A colon (:) separates the physical and virtual parts of the interface name.

Channel Part of an Interface Name

The channel identifier part of the interface name is required only on channelized interfaces. For channelized interfaces, channel 0 identifies the first channelized interface. For channelized IQ and channelized IQE interfaces, channel 1 identifies the first channelized interface. A nonconcatenated (that is, channelized) SONET/SDH OC48 interface has four OC12 channels, numbered 0 through 3.

To determine which types of channelized PICs are currently installed in the router, use the **show chassis hardware** command from the top level of the command-line interface (CLI). Channelized IQ and IQE PICs are listed in the output with “intelligent queuing IQ” or “enhanced intelligent queuing IQE” in the description. For more information, see “Channelized Interfaces” on page 385.

For ISDN interfaces, you specify the B-channel in the form **bc-pim/0/port:n**. *n* is the B-channel ID and can be 1 or 2. You specify the D-channel in the form **dc-pim/0/port:0**.



NOTE: For ISDN, the B- and D-channel interfaces do not have any configurable parameters. However, when interface statistics are displayed, B- and D-channel interfaces have statistical values.




NOTE: In the JUNOS Software implementation, the term *logical interfaces* generally refers to interfaces you configure by including the **unit** statement at the **[edit interfaces interface-name]** hierarchy level. Logical interfaces have the *.logical* descriptor at the end of the interface name, as in **ge-0/0/0.1** or **t1-0/0/0.0.1**, where the logical unit number is 1.

Although channelized interfaces are generally thought of as logical or virtual, the JUNOS Software sees T3, T1, and NxDS0 interfaces within a channelized IQ or IQE PIC as physical interfaces. For example, both **t3-0/0/0** and **t3-0/0/0.1** are treated as physical interfaces by the JUNOS Software. In contrast, **t3-0/0/0.2** and **t3-0/0/0.1.2** are considered logical interfaces because they have the .2 at the end of the interface names.

Interface Naming for a Routing Matrix Based on a TX Matrix Router

A routing matrix based on a Juniper Networks TX Matrix Router is a multichassis architecture composed of one TX Matrix router and from one to four interconnected T640 routers. From the perspective of the user interface, the routing matrix appears as a single router. The TX Matrix router controls all the T640 routers, as shown in Figure 3 on page 58.

Figure 3: Routing Matrix

A TX Matrix router is also referred to as a *switch-card chassis* (SCC). The CLI uses `scc` to refer to the TX Matrix router. A T640 router in a routing matrix is also referred to as a *line-card chassis* (LCC). The CLI uses `lcc` as a prefix to refer to a specific T640 router.

LCCs are assigned numbers, 0 through 3, depending on the hardware setup and connectivity to the TX Matrix router. For more information, see the *TX Matrix Router Hardware Guide*. A routing matrix can have up to four T640 routers, and each T640 router has up to eight FPCs. Therefore, the routing matrix as a whole can have up to 32 FPCs (0 through 31).

In the JUNOS CLI, an interface name has the following format:

type-fpc/pic/port

When you specify the *fpc* number for a T640 router in a routing matrix, the JUNOS Software determines which T640 router contains the specified FPC based on the following assignment:

- On LCC 0, FPC hardware slots 0 through 7 are configured as 0 through 7.
- On LCC 1, FPC hardware slots 0 through 7 are configured as 8 through 15.
- On LCC 2, FPC hardware slots 0 through 7 are configured as 16 through 23.
- On LCC 3, FPC hardware slots 0 through 7 are configured as 24 through 31.

For example, the 1 in `se-1/0/0` refers to FPC hardware slot 1 on the T640 router labeled `lcc0`. The 11 in `t1-11/2/0` refers to FPC hardware slot 3 on the T640 router labeled `lcc1`. The 20 in `so-20/0/1` refers to FPC hardware slot 4 on the T640 router labeled `lcc2`. The 31 in `t3-31/1/0` refers to FPC hardware slot 7 on the T640 router labeled `lcc3`.

Table 4 on page 59 summarizes the FPC numbering for a T640 router in a routing matrix.

Table 4: FPC Numbering for T640 Routers in a Routing Matrix

LCC Numbers Assigned to the T640 Router	Configuration Numbers
0	0 through 7
1	8 through 15
2	16 through 23
3	24 through 31

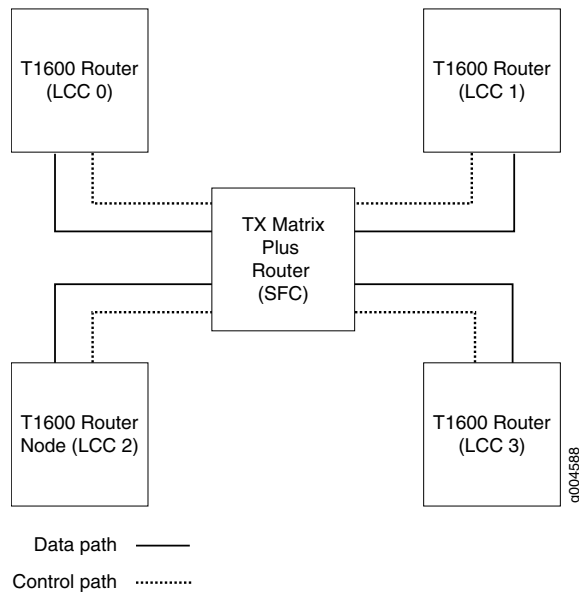
Table 5 on page 59 lists each FPC hardware slot and the corresponding configuration numbers for LCCs 0 through 3.

Table 5: One-to-One FPC Numbering for T640 Routers in a Routing Matrix

FPC Numbering	T640 Routers							
LCC 0								
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	0	1	2	3	4	5	6	7
LCC 1								
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	8	9	10	11	12	13	14	15
LCC 2								
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	16	17	18	19	20	21	22	23
LCC 3								
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	24	25	26	27	28	29	30	31

Interface Naming for a Routing Matrix Based on a TX Matrix Plus Router

A routing matrix based on a Juniper Networks TX Matrix Plus Router is a multichassis architecture composed of one TX Matrix Plus router and from one to four interconnected T1600 routers. From the perspective of the user interface, the routing matrix appears as a single router. The TX Matrix Plus router controls all the T1600 routers, as shown in Figure 4 on page 60.

Figure 4: Routing Matrix Based on a TX Matrix Plus Router

A TX Matrix Plus router is also referred to as a *switch-fabric chassis* (SFC). The CLI uses `sfc` to refer to the TX Matrix Plus router. A T1600 router in a routing matrix is also referred to as a *line-card chassis* (LCC). The CLI uses `lcc` as a prefix to refer to a specific T1600 router.

LCCs are assigned numbers, 0 through 3, depending on the hardware setup and connectivity to the TX Matrix Plus router. For more information, see the *TX Matrix Plus Router Hardware Guide*. A routing matrix based on a TX Matrix Plus router can have up to four T1600 routers, and each T1600 router has up to eight FPCs. Therefore, the routing matrix as a whole can have up to 32 FPCs (0 through 31).

In the JUNOS CLI, an interface name has the following format:

type-fpc/pic/port

When you specify the *fpc* number for a T1600 router in a routing matrix, the JUNOS Software determines which T1600 router contains the specified FPC based on the following assignment:

- On LCC 0, FPC hardware slots 0 through 7 are configured as 0 through 7.
- On LCC 1, FPC hardware slots 0 through 7 are configured as 8 through 15.
- On LCC 2, FPC hardware slots 0 through 7 are configured as 16 through 23.
- On LCC 3, FPC hardware slots 0 through 7 are configured as 24 through 31.

For example, the 1 in `se-1/0/0` refers to FPC hardware slot 1 on the T1600 router labeled `lcc0`. The 11 in `t1-11/2/0` refers to FPC hardware slot 3 on the T1600 router labeled `lcc1`. The 20 in `so-20/0/1` refers to FPC hardware slot 4 on the T1600 router labeled `lcc2`. The 31 in `t3-31/1/0` refers to FPC hardware slot 7 on the T1600 router labeled `lcc3`.

Table 6 on page 61 summarizes the FPC numbering for a routing matrix based on a TX Matrix Plus router.

Table 6: FPC Numbering for T1600 Routers in a Routing Matrix

LCC Numbers Assigned to the T1600 Router	Configuration Numbers
0	0 through 7
1	8 through 15
2	16 through 23
3	24 through 31

Table 7 on page 61 lists each FPC hardware slot and the corresponding configuration numbers for LCCs 0 through 3.

Table 7: One-to-One FPC Numbering for T1600 Routers in a Routing Matrix

FPC Numbering	T1600 Routers							
LCC 0								
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	0	1	2	3	4	5	6	7
LCC 1								
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	8	9	10	11	12	13	14	15
LCC 2								
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	16	17	18	19	20	21	22	23
LCC 3								
Hardware Slots	0	1	2	3	4	5	6	7
Configuration Numbers	24	25	26	27	28	29	30	31

Chassis Interface Naming

You configure some PIC properties, such as framing, at the [edit chassis] hierarchy level. Chassis interface naming varies depending on the routing hardware.

- To configure PIC properties for a standalone router, you must specify the FPC and PIC numbers, as follows:

```
[edit chassis]
fpc slot-number {
  pic pic-number {
    ...
  }
}
```

- To configure PIC properties for a T640 or T1600 router configured in a routing matrix, you must specify the LCC, FPC, and PIC numbers, as follows:

```
[edit chassis]
lcc lcc-number {
  fpc slot-number { # Use the hardware FPC slot number
    pic pic-number {
      ...
    }
  }
}
```


For the FPC slot in a T640 router in a routing matrix, specify the actual hardware slot number, as labeled on the T640 router chassis. Do not use the corresponding software FPC configuration numbers shown in Table 5 on page 59.

For the FPC slot in a T1600 router in a routing matrix, specify the actual hardware slot number, as labeled on the T1600 router chassis. Do not use the corresponding software FPC configuration numbers shown in Table 6 on page 61.

For more information about the [edit chassis] hierarchy, see the *JUNOS System Basics Configuration Guide*.

Examples: Interface Naming

This section provides examples of naming interfaces. For an illustration of where slots, PICs, and ports are located, see Figure 5 on page 63.

Figure 5: Interface Slot, PIC, and Port Locations

For an FPC in slot 1 with two OC3 SONET/SDH PICs in PIC positions 0 and 1, each PIC with two ports uses the following names:

```
so-1/0/0.0
so-1/0/1.0
so-1/1/0.0
so-1/1/1.0
```

An OC48 SONET/SDH PIC in slot 1 and in concatenated mode appears as a single FPC with a single PIC, which has a single port. If this interface has a single logical unit, it has the following name:

```
so-1/0/0.0
```

An OC48 SONET/SDH PIC in slot 1 and in channelized mode has a number for each channel. For example:

```
so-1/0/0:0
so-1/0/0:1
```

For an FPC in slot 1 with a Channelized OC12 PIC in PIC position 2, the DS3 channels have the following names:

```
t3-1/2/0:0
t3-1/2/0:1
t3-1/2/0:2
...
t3-1/2/0:11
```

For an FPC in slot 1 with four OC12 ATM PICs (the FPC is fully populated), the four PICs, each with a single port and a single logical unit, have the following names:

```
at-1/0/0.0
```

```
at-1/1/0.0
at-1/2/0.0
at-1/3/0.0
```

In a routing matrix on the T640 router labeled **lcc1**, for an FPC in slot 5 with four SONET OC192 PICs, the four PICs, each with a single port and a single logical unit, have the following names:

```
so-13/0/0.0
so-13/1/0.0
so-13/2/0.0
so-13/3/0.0
```

For an FPC in slot 1 with one 4-port BRI interface card, port 4 has the following name:

```
br-1/0/4
```

The first B channel, the second B channel, and the control channel have the following names:

```
bc-1/0/4:1
bc-1/0/4:2
dc-1/0/4:0
```

Displaying Interface Configurations

To display a configuration, use either the **show** command in configuration mode or the **show configuration** top-level command. Interfaces are listed in numerical order, from lowest to highest slot number, then from lowest to highest PIC number, and finally from lowest to highest port number.

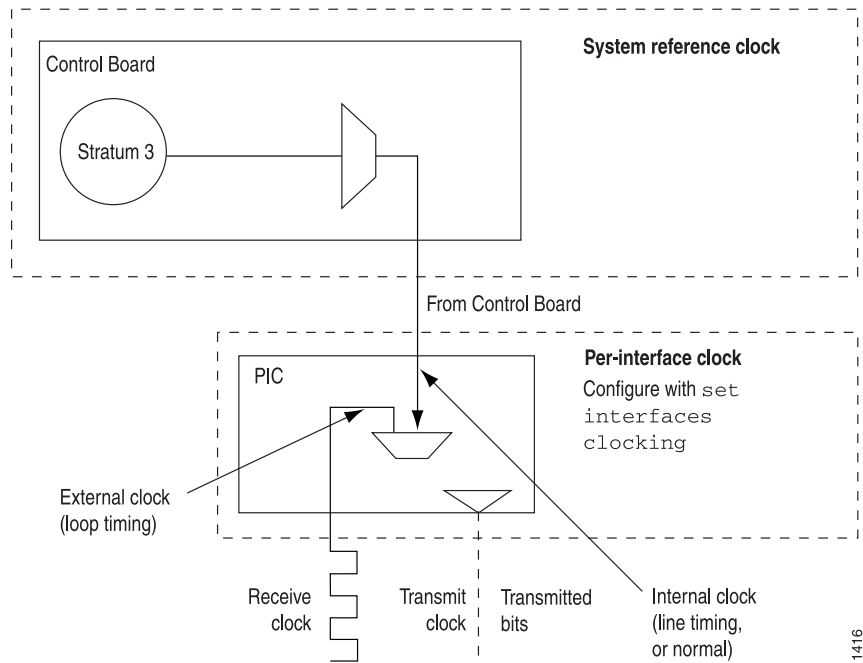
Interface and Router Clock Sources

When configuring the router, you can configure the *transmit clock* on each interface; the transmit clock aligns each outgoing packet transmitted over the router's interfaces. For both the router and interfaces, the clock source can be the router's internal Stratum 3 clock, which resides on the control board, or an external clock that is received from the interface you are configuring. For example, interface A can transmit on interface A's received clock (external, loop timing) or the Stratum 3 clock (internal, line timing). Interface A cannot use a clock from any other source.

By default, each interface uses the router's internal Stratum 3 clock. To configure the clock source of each interface, include the **clocking** statement at the **[edit interfaces interface-name]** hierarchy level:

```
[edit interfaces interface-name]
clocking (internal | external);
```

Figure 6 on page 65 illustrates the different clock sources.

Figure 6: Clock Sources

Configuring an External Synchronization Interface

The M Series 320 router supports an external synchronization interface that can be configured to synchronize the internal Stratum 3 clock to an external source, and then synchronize the chassis interface clock to the external source.

This feature can be configured for external primary and secondary interfaces that use Building Integrated Timing System (BITS) or SDH Equipment Timing Source (SETS) timing sources. When internal timing is set for SONET/SDH, Plesiochronous Digital Hierarchy (PDH), and digital hierarchy (DS1) interfaces on the Physical Interface Cards (PICs), the transmit clock of the interface is synchronized to BITS/SETS timing and traceable to timing within the network.

To configure external synchronization on the M Series 320 router, include the **synchronization** statement at the **[edit chassis]** hierarchy level.

For more information about the external synchronization interface, see the *JUNOS System Basics Configuration Guide*.

Chapter 3

Configuring Physical Interface Properties

The software driver for each network media type sets reasonable default values for general interface properties, such as the interface's maximum transmission unit (MTU) size, receive and transmit leaky bucket properties, link operational mode, and clock source.

This chapter discusses configuration of the following physical interface properties:

- Physical Interface Configuration Statements on page 68
- Physical Interfaces Properties Statements List on page 77
- Specifying an Aggregated Interface on page 92
- Specifying a USB Modem Interface on J Series Routers on page 93
- Specifying OC768-over-OC192 Mode on page 95
- Adding an Interface Description to the Configuration on page 96
- Configuring the Link Characteristics on page 97
- Configuring the Media MTU on page 98
- Configuring Interface Encapsulation on Physical Interfaces on page 106
- Configuring the PPP Challenge Handshake Authentication Protocol on page 112
- Configuring the PPP Password Authentication Protocol on page 114
- Monitoring a PPP Session on page 118
- Tracing Operations of the pppd Process on page 119
- Configuring PPP Address and Control Field Compression on page 120
- Configuring the PPP Protocol Field Compression on page 121
- Configuring the Interface Speed on page 122
- Configuring Keepalives on page 126
- Configuring the Clock Source on page 128
- Configuring the Router as a DCE on page 128
- Configuring Receive and Transmit Leaky Bucket Properties on page 129
- Configuring Accounting for the Physical Interface on page 130
- Interface Diagnostics on page 131
- Tracing Operations of an Individual Router Interface on page 137
- Damping Interface Transitions on page 138

- Configuring Multiservice Physical Interface Properties on page 138
- Enabling or Disabling SNMP Notifications on Physical Interfaces on page 139
- Enabling Unidirectional Traffic Flow on Physical Interfaces on page 139
- Disabling a Physical Interface on page 140

Physical Interface Configuration Statements

M Series, MX Series, T Series, and J Series routers are factory configured according to the specific router, its features, and its physical interfaces. This section includes a default configuration example showing the statements used to configure the physical interfaces properties. Additional statements are used to set properties for specific interface types and are described in “Physical Interfaces Properties Statements List” on page 77.

To modify any of the default general interface properties, include the appropriate statements at the [edit interfaces *interface-name*] hierarchy level:

```

interfaces {
  traceoptions {
    file filename <files number> <match regular-expression> <size size>
      <world-readable | no-world-readable>;
    flag flag <disable>;
  }
  interface-name {
    accounting-profile name;
    aggregated-ether-options {
      (flow-control | no-flow-control);
      lacp {
        (active | passive);
        link-protection {
          disable;
        }
        (revertive | non-revertive);
        periodic interval;
        system-priority priority;
      }
      link-protection;
      link-speed speed;
      (loopback | no-loopback);
      minimum-links number;
      source-address-filter {
        mac-address
      }
      (source-filtering | no-source-filtering);
    }
    aggregated-sonet-options {
      link-speed speed | mixed;
      minimum-links number;
    }
    atm-options {
      cell-bundle-size cells;
      ilmi;
      linear-red-profiles profile-name {
        high-plp-max-threshold percent;
      }
    }
  }
}

```

```

    low-plp-max-threshold percent;
    queue-depth cells high-plp-threshold percent low-plp-threshold percent;
}
mpls {
    pop-all-labels {
        required-depth number;
    }
}
pic-type (atm1 | atm2);
plp-to-clp;
promiscuous-mode {
    vpi vpi-identifier;
}
scheduler-maps map-name {
    forwarding-class class-name {
        epd-threshold cells plp1 cells;
        linear-red-profile profile-name;
        priority (high | low);
        transmit-weight (cells number | percent number);
    }
    vc-cos-mode (alternate | strict);
}
vpi vpi-identifier {
    maximum-vcs maximum-vcs;
    oam-liveness {
        up-count cells;
        down-count cells;
    }
    oam-period (seconds | disable);
    shaping {
        (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate
        sustained rate burst length);
        queue-length number;
    }
}
}
clocking clock-source;
data-input (system | interface interface-name);
dce;
serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
}

```

```

dsr-polarity (negative | positive);
dte-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
}
dtr-circuit (balanced | unbalanced);
dtr-polarity (negative | positive);
encoding (nrz | nrzi);
indication-polarity (negative | positive);
line-protocol protocol;
loopback mode;
rts-polarity (negative | positive);
tm-polarity (negative | positive);
transmit-clock invert;
}
description text;
dialer-options {
    pool pool-name <priority priority>;
}
disable;
ds0-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    byte-encoding (nx56 | nx64);
    fcs (16 | 32);
    idle-cycle-flag (flags | ones);
    invert-data;
    loopback payload;
    start-end-flag (filler | shared);
}
e1-options {
    bert-error-rate rate;
    bert-period seconds;
    fcs (16 | 32);
    framing (g704 | g704-no-crc4 | unframed);
    idle-cycle-flag (flags | ones);
    invert-data;
    loopback (local | remote);
    start-end-flag (filler | shared);
    timeslots time-slot-range;
}
e3-options {
    atm-encapsulation (direct | plcp);
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout feet;
    compatibility-mode (digital-link | kentrox | larscom) <subrate value>;

```

```

fcs (16 | 32);
framing (g.751 | g.832);
idle-cycle-flag (filler | shared);
invert-data;
loopback (local | remote);
(payload-scrambler | no-payload-scrambler);
start-end-flag (filler | shared);
(unframed | no-unframed);
}
encapsulation type;
es-options {
    backup-interface es-fpc/pic/port;
}
fastether-options {
    802.3ad aex;
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    ingress-rate-limit rate;
    (loopback | no-loopback);
    mpls {
        pop-all-labels {
            required-depth number;
        }
    }
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
}
flexible-vlan-tagging;
gigether-options {
    802.3ad aex;
    (asynchronous-notification | no-asynchronous-notification);
    (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
        local-interface-offline>;
    auto-reconnect seconds;
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    (loopback | no-loopback);
    mpls {
        pop-all-labels {
            required-depth number;
        }
    }
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
    ethernet-switch-profile {
        (mac-learn-enable | no-mac-learn-enable);
        tag-protocol-id [ tpids ];
        ethernet-policer-profile {
            input-priority-map {
                ieee802.1p premium [ values ];
            }
            output-priority-map {

```

```

        classifier {
            premium {
                forwarding-class class-name {
                    loss-priority (high | low);
                }
            }
        }
    }
    policer cos-policer-name {
        aggregate {
            bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
            burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
        }
        premium {
            bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
            burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
        }
    }
}
(gratuitous-arp-reply | no-gratuitous-arp-reply);
hold-time up milliseconds down milliseconds;
interface-set interface-set-name {
    interface ethernet-interface-name {
        (unit unit-number | vlan-tags-outer vlan-tag);
    }
}
}
isdn-options {
    bchannel-allocation (ascending | descending);
    calling-number number;
    pool pool-name <priority priority>;
    spid1 spid-string;
    spid2 spid-string;
    static-tei-val value;
    switch-type (att5e | etsi | ni1 | ntdms100 | ntt);
    t310 seconds;
    tei-option (first-call | power-up);
}
keepalives <down-count number> <interval seconds> <up-count number>;
link-mode mode;
lmi (Frame Relay) {
    lmi-type (ansi | itu);
    n391dte number;
    n392dce number;
    n392dte number;
    n393dce number;
    n393dte number;
    t391dte seconds;
    t392dce seconds;
}
lsq-failure-options {
    no-termination-request;
    [ trigger-link-failure interface-name ];
}
}
mac mac-address;

```



```

mlfr-uni-nni-bundle-options {
    acknowledge-retries number;
    acknowledge-timer milliseconds;
    action-red-differential-delay (disable-tx | remove-link);
    cisco-interoperability send-lip-remove-link-for-link-reject;
    drop-timeout milliseconds;
    fragment-threshold bytes;
    hello-timer milliseconds;
    link-layer-overhead percent;
    lmi-type (ansi | itu);
    minimum-links number;
    mrru bytes;
    n391 number;
    n392 number;
    n393 number;
    red-differential-delay milliseconds;
    t391 seconds;
    t392 seconds;
    yellow-differential-delay milliseconds;
    encapsulation type;
}
modem-options {
    dialin (console | routable);
    init-command-string initialization-command-string;
}
mtu bytes;
multiservice-options {
    (core-dump | no-core-dump);
    (syslog | no-syslog);
    (dump-on-flow-control);
}
native-vlan-id number;
no-gratuitous-arp-request;
no-keepalives;
no-partition {
    interface-type type;
}
optics-options {
    wavelength nm;
}
partition partition-number oc-slice oc-slice-range interface-type type;
timeslots time-slot-range;
passive-monitor-mode;
per-unit-scheduler;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
}
dynamic-profile profile-name;

```

```

no-termination-request;
pap {
    access-profile name;
    local-name name;
    local-password password;
    passive;
}
}
receive-bucket {
    overflow (discard | tag);
    rate percentage;
    threshold bytes;
}
redundancy-options {
    primary sp-fpc/pic/port;
    secondary sp-fpc/pic/port;
}
schedulers number;
serial-options {
    clock-rate rate;
    clocking-mode (dce | internal | loop);
    control-polarity (negative | positive);
    cts-polarity (negative | positive);
    dcd-polarity (negative | positive);
    dce-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dsr-polarity (negative | positive);
    dte-options {
        control-signal (assert | de-assert | normal);
        cts (ignore | normal | require);
        dcd (ignore | normal | require);
        dsr (ignore | normal | require);
        dtr signal-handling-option;
        ignore-all;
        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dtr-circuit (balanced | unbalanced);
    dtr-polarity (negative | positive);
    encoding (nrz | nrzi);
    indication-polarity (negative | positive);
    line-protocol protocol;
    loopback mode;
    rts-polarity (negative | positive);
    tm-polarity (negative | positive);
    transmit-clock invert;

```

```

}
services-options {
  inactivity-timeout seconds;
  open-timeout seconds;
  syslog {
    host hostname {
      facility-override facility-name;
      log-prefix prefix-number;
      services priority-level;
    }
  }
}
shdsl-options {
  annex (annex-a | annex-b);
  line-rate line-rate;
  loopback (local | remote);
  snr-margin {
    current margin;
    snext margin;
  }
}
sonet-options {
  aggregate asx;
  aps {
    advertise-interval milliseconds;
    authentication-key key;
    force;
    hold-time milliseconds;
    lockout;
    neighbor address;
    paired-group group-name;
    preserve-interface;
    protect-circuit group-name;
    request;
    revert-time seconds;
    switching-mode (bidirectional | unidirectional);
    working-circuit group-name;
  }
  bytes {
    c2 value;
    e1-quiet value;
    f1 value;
    f2 value;
    s1 value;
    z3 value;
    z4 value;
  }
  fcs (16 | 32);
  loopback (local | remote);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  path-trace trace-string;
  (payload-scrambler | no-payload-scrambler);
}

```

```

    rfc-2615;
    trigger {
        defect ignore;
        hold-time up milliseconds down milliseconds;
    }
    vtmapping (itu-t | klm);
    (z0-increment | no-z0-increment);
}
speed (10m | 100m | 1g | oc3 | oc12 | oc48);
stacked-vlan-tagging;
switch-options {
    switch-port port-number {
        (auto-negotiation | no-auto-negotiation);
        speed (10m | 100m | 1g);
        link-mode (full-duplex | half-duplex);
    }
}
t1-options {
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout value;
    byte-encoding (nx56 | nx64);
    crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);
    crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);
    fcs (16 | 32);
    framing (esf | sf);
    idle-cycle-flag (flags | ones);
    invert-data;
    line-encoding (ami | b8zs);
    loopback (local | payload | remote);
    remote-loopback-respond;
    start-end-flag (filler | shared);
    timeslots time-slot-range;
}
t3-options {
    atm-encapsulation (direct | plcp);
    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    buildout feet;
    (cbit-parity | no-cbit-parity);
    compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate
        value>;
    fcs (16 | 32);
    (feac-loop-respond | no-feac-loop-respond);
    idle-cycle-flag value;
    (long-buildout | no-long-buildout);
    (loop-timing | no-loop-timing);
    loopback (local | payload | remote);
    (mac | no-mac);
    (payload-scrambler | no-payload-scrambler);
    start-end-flag (filler | shared);
}
traceoptions {
    flag flag <flag-modifier> <disable>;

```

```

    }
    transmit-bucket {
        overflow discard;
        rate percentage;
        threshold bytes;
    }
    (traps | no-traps);
    unidirectional;
    vlan-tagging;
    vlan-vci-tagging;
    unit logical-unit-number {
        logical-interface-statements;
    }
}

```

For information about interface-specific physical properties, see “Physical Interfaces Properties Statements List” on page 77.

Physical Interfaces Properties Statements List

Table 8 on page 77 lists statements that you can use to configure physical interfaces.

Table 8: Statements for Physical Interface Properties

Statement	Interface Types	Usage Guidelines
802.3ad aex	Aggregated Ethernet interfaces	“Configuring Ethernet Link Aggregation” on page 625 or “Configuring Aggregated Ethernet Interfaces” on page 623
access-profile <i>name</i>	Interfaces with Point-to-Point Protocol (PPP) encapsulation	“Configuring the PPP Challenge Handshake Authentication Protocol” on page 112
accounting-profile <i>name</i>	All	“Configuring Accounting for the Physical Interface” on page 130
acfc	Interfaces with PPP encapsulation	“Identifying the Access Concentrator” on page 791
acknowledge-retries <i>number</i>	Link services and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
acknowledge-timer <i>milliseconds</i>	Link services and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
action-red-differential-delay (disable-tx remove-link)	Link services and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
advertise-interval <i>milliseconds</i>	SONET/SDH interfaces	“Configuring APS and MSP” on page 859

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
<code>aggregate</code>	Gigabit Ethernet intelligent queuing (IQ and IQE) interfaces and Gigabit Ethernet interfaces with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	“Configuring Gigabit Ethernet Policers” on page 757
<code>aggregate asx</code>	Aggregated SONET/SDH interfaces	“Configuring Aggregated SONET/SDH Interfaces” on page 881
<code>aggregated-ether-options</code>	Aggregated Ethernet interfaces	“Configuring Aggregated Ethernet Interfaces” on page 623
<code>aggregate-ports</code>	SONET/SDH interfaces	“Specifying OC768-over-OC192 Mode” on page 95
<code>aggregated-sonet-options</code>	Aggregated SONET/SDH interfaces	“Configuring Aggregated SONET/SDH Interfaces” on page 881
<code>annex (annex-a annex-b)</code>	ATM interfaces on J Series routers SONET interfaces using annex-b for MSP switching on M320 and M120 Routers	“Configuring SHDSL Operating Mode on an ATM Physical Interface” on page 364 Configuring APS and MSP
<code>aps</code>	SONET/SDH interfaces	“Configuring APS and MSP” on page 859
<code>atm-encapsulation (direct plcp)</code>	E3 and T3 traffic over Asynchronous Transfer Mode (ATM) interfaces	“Configuring E3 and T3 Parameters on ATM Interfaces” on page 337
<code>atm-options</code>	ATM1 and ATM2 IQ interfaces	“Configuring ATM Interfaces” on page 279
<code>authentication-key key</code>	SONET/SDH interfaces	“Configuring APS and MSP” on page 859
<code>backup-interface</code>	E1, E3, T1, T3 and Fast Ethernet	“Configuring an ISDN Dialer Interface as a Backup Interface” on page 826
<code>bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps</code>	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	“Configuring Gigabit Ethernet Policers” on page 757
<code>bchannel-allocation (ascending descending)</code>	J Series routers equipped with a Dual-Port Channelized T1/E1 PIM; for Integrated Services Digital Network Primary Rate Interfaces (ISDN PRI)	“Allocating B-Channels for Dialout” on page 513
<code>bert-algorithm algorithm</code>	E3, T1, T3, multichannel DS3, channelized interfaces (DS3, OC12, and STM1), and channelized IQ and IQE interfaces (E1 and DS3)	“Interface Diagnostics” on page 134
<code>bert-error-rate rate</code>	E1, E3, T1, T3, and channelized interfaces (DS3, OC3, OC12, and STM1)	“Interface Diagnostics” on page 134

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
bert-period <i>seconds</i>	E1, E3, T1, T3, and channelized interfaces (DS3, OC12, and STM1)	“Interface Diagnostics” on page 134
buildout <i>value</i>	T1 interfaces	“Configuring the T1 Buildout” on page 561
buildout <i>feet</i>	E3 and T3 traffic over ATM interfaces	“Configuring E3 and T3 Parameters on ATM Interfaces” on page 337
burst-size-limit (Policer for Gigabit Ethernet Interfaces) <i>bytes</i>	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	“Configuring Gigabit Ethernet Policers” on page 757
byte-encoding (nx56 nx64)	DS0 and T1 interfaces	“Configuring T1 Byte Encoding” on page 561
bytes [<i>values</i>	SONET/SDH interfaces	“Configuring SONET/SDH Header Byte Values” on page 849
cbit-parity no-cbit-parity	T3 interfaces	“Disabling T3 C-Bit Parity Mode” on page 571
cbr <i>rate</i>	ATM interfaces	“Defining the ATM Traffic-Shaping Profile” on page 319
cell-bundle-size <i>cells</i>	ATM2 IQ interfaces using ATM Layer 2 circuit cell-relay transport mode	“Configuring the Layer 2 Circuit Cell-Relay Cell Maximum” on page 313
chap	Interfaces with PPP encapsulation	“Configuring the PPP Challenge Handshake Authentication Protocol” on page 112
cisco-interoperability send-lip-remove-link-for-link-reject	link services IQ (lsq) interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
classifier	Gigabit Ethernet IQ interfaces	“Specifying an Output Priority Map” on page 759
clocking <i>clock-source</i>	ATM, DS0, E1, E3, SONET/SDH, T1, and T3 interfaces	“Configuring the Clock Source” on page 128
clocking-mode (dce internal loop)	Serial interfaces (EIA-530 and V.35)	“Configuring the Serial Clocking Mode” on page 269
clock-rate <i>rate</i>	Serial interfaces (EIA-530 and V.35)	“Configuring the DTE Clock Rate” on page 270
compatibility-mode <i>mode</i>	E3 and T3 interfaces	“Configuring the E3 CSU Compatibility Mode” on page 553 and “Configuring the T3 CSU Compatibility Mode” on page 572
compression	Interfaces with PPP encapsulation	“Configuring the PPP Protocol Field Compression” on page 121

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
control-polarity (negative positive)	Serial interfaces (X.21)	“Configuring Serial Signal Polarities” on page 274
control-signal (assert de-assert normal)	Serial interfaces (X.21)	“Configuring the Serial Signal Handling” on page 271
core-dump no-core-dump)	Adaptive services, monitoring services, and collector interfaces	“Configuring Multiservice Physical Interface Properties” on page 138
cts (ignore normal require)	Serial interfaces (EIA-530 and V.35)	“Configuring the Serial Signal Handling” on page 271
cts-polarity (negative positive)	Serial interfaces (EIA-530 and V.35)	“Configuring Serial Signal Polarities” on page 274
current <i>margin</i>	ATM interfaces on J Series routers	“Configuring SHDSL Operating Mode on an ATM Physical Interface” on page 364
dcd (ignore normal require)	Serial interfaces (EIA-530 and V.35)	“Configuring the Serial Signal Handling” on page 271
dcd-polarity (negative positive)	Serial interfaces (EIA-530 and V.35)	“Configuring Serial Signal Polarities” on page 274
dce	Interfaces with Frame Relay encapsulation	“Configuring the Router as a DCE” on page 128
dce-options	Serial interfaces (EIA-530, V.35, and X.21) on J Series routers	“Configuring the Serial Signal Handling” on page 271
default-chap-secret <i>name</i>	Interfaces with Point-to-Point Protocol (PPP) encapsulation	“Configuring a Default CHAP Secret” on page 113
description text	All	“Adding an Interface Description to the Configuration” on page 96
dialer-options	ISDN interfaces	“Configuring ISDN Physical Interface Properties” on page 821
disable	All	“Disabling a Physical Interface” on page 140, “Tracing Operations of an Individual Router Interface” on page 241, and “Tracing Operations of an Individual Router Interface” on page 137
dot1x	802.1x Port-Based Network Access Control	“Configuring IEEE 802.1x Port-Based Network Access Control” on page 741
down-count	ATM interfaces	“Configuring the ATM OAM F5 Loopback Cell Threshold” on page 329
drop-timeout <i>milliseconds</i>	Multilink, link services, and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
ds0-options	DS0 interfaces	“Channelized Interfaces” on page 385

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
<code>dsr</code> (ignore normal require)	Serial interfaces (EIA-530 and V.35)	“Configuring the Serial Signal Handling” on page 271
<code>dsr-polarity</code> (negative positive)	Serial interfaces (EIA-530 and V.35)	“Configuring Serial Signal Polarities” on page 274
<code>dte-options</code>	Serial interfaces (EIA-530, V.35, and X.21) on M Series and T Series routers	“Configuring the Serial Signal Handling” on page 271
<code>dtr signal-handling-option</code>	Serial interfaces (EIA-530 and V.35)	“Configuring the Serial Signal Handling” on page 271
<code>dtr-circuit</code> (balanced unbalanced)	Serial interfaces (EIA-530 and V.35)	“Configuring the Serial DTR Circuit” on page 274
<code>dtr-polarity</code> (negative positive)	Serial interfaces (EIA-530 and V.35)	“Configuring Serial Signal Polarities” on page 274
<code>e1-options</code>	E1 interfaces	“Configuring E1 Interfaces” on page 543
<code>e3-options</code>	E3 interfaces	“Configuring E3 Interfaces” on page 551
<code>encapsulation type</code>	All interfaces, except loopback and multicast tunnel	“Configuring Interface Encapsulation on Physical Interfaces” on page 106
<code>encoding</code> (nrz nrzi)	Serial interfaces (EIA-530, V.35, and X.21)	“Configuring Serial Line Encoding” on page 277
<code>epd-threshold cells</code>	ATM2 interfaces	“Configuring ATM2 IQ VC Tunnel CoS Components” on page 339
<code>es-options</code>	ES interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
<code>ethernet-policer-profile</code>	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC, and the built-in Gigabit Ethernet port on the M7i router)	“Configuring Gigabit Ethernet Policers” on page 757
<code>ethernet-switch-profile</code>	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC, Aggregated Ethernet with Gigabit Ethernet IQ interfaces, and the built-in Gigabit Ethernet port on the M7i router)	“Configuring Gigabit Ethernet Policers” on page 757, “Configuring MAC Address Filtering” on page 761, and Configuring the Management Ethernet Interface
<code>facility-override facility-name</code>	Adaptive services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
<code>fastether-options</code>	Fast Ethernet interfaces	“Configuring Ethernet Interfaces” on page 585

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
fcs (16 32)	E1/E3, SONET/SDH, and T1/T3 interfaces	“Configuring the E1 Frame Checksum” on page 545, “Configuring the E3 Frame Checksum” on page 554, “Configuring the SONET/SDH Frame Checksum” on page 851, “Configuring the T1 Frame Checksum” on page 563, and “Configuring the T3 Frame Checksum” on page 574
feac-loop-respond no-feac-loop-respond)	T3 interfaces	“Configuring the T3 FEAC Response” on page 575
flow-control no-flow-control)	Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces	“Configuring Flow Control” on page 594
force	SONET/SDH interfaces	“Configuring APS and MSP” on page 859
forwarding-class <i>class-name</i>	Gigabit Ethernet IQ and ATM2 interfaces	“Configuring ATM2 IQ VC Tunnel CoS Components” on page 339 and “Specifying an Output Priority Map” on page 759
fragment-threshold <i>bytes</i>	Multilink, link services, and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
framing <i>framing-type</i>	10-Gigabit Ethernet, SONET, E1, E3, and T1 interfaces	“Configuring E3 and T3 Parameters on ATM Interfaces” on page 337, “Configuring E1 Framing” on page 546, and “Configuring T1 Framing” on page 564, “Configuring 10-Gigabit Ethernet Framing” on page 781, and “Configuring SONET Options for 10-Gigabit Ethernet Interfaces” on page 872
gether-options	Gigabit Ethernet and Tri-Rate Ethernet copper interfaces	“Configuring Ethernet Interfaces” on page 585
(gratuitous-arp-reply no-gratuitous-arp-reply)	Ethernet interfaces	“Configuring Gratuitous ARP” on page 596
hello-timer <i>milliseconds</i>	Link services and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
high-plp-max-threshold	ATM2 interfaces	“Configuring Linear RED Profiles” on page 340
high-plp-threshold <i>percent</i>	ATM2 interfaces	= “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339
hold-time <i>milliseconds</i>	SONET/SDH interfaces	“Configuring APS and MSP” on page 859
hold-time up <i>milliseconds</i> down <i>milliseconds</i>	All interfaces, except aggregated SONET/SDH, generalized routing encapsulation (GRE) tunnel, and IP tunnel	“Damping Interface Transitions” on page 138 and “Configuring SONET/SDH Defect Triggers to Be Ignored” on page 855

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
host <i>hostname</i>	Adaptive services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
ieee802.1p premium [<i>values</i>]	Gigabit Ethernet IQ interfaces	“Specifying an Input Priority Map” on page 758
idle-cycle-flag <i>value</i>	E1, E3, T1, and T3 interfaces	“Configuring the E1 Idle Cycle Flag” on page 546, “Configuring the E3 Idle Cycle Flag” on page 555, “Configuring the T1 Idle Cycle Flag” on page 566, and “Configuring the T3 Idle Cycle Flag” on page 575
ignore-all	Serial interfaces (EIA-530, V.35, and X.21)	“Configuring the Serial Signal Handling” on page 271
ilmi	ATM interfaces	“Configuring Communication with Directly Attached ATM Switches and Routers” on page 291
inactivity-timeout <i>seconds</i>	Adaptive services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
indication (ignore normal require)	Serial interfaces (X.21)	“Configuring the Serial Signal Handling” on page 271
indication-polarity (negative positive)	Serial interfaces (X.21)	“Configuring Serial Signal Polarities” on page 274
ingress-rate-limit <i>rate</i>	8-port, 12-port, and 48-port Fast Ethernet interfaces	“Configuring the Ingress Rate Limit” on page 597
init-command-string <i>initialization-command-string;</i>	For USB ports (umd0) on J4350 and J6350 Services Routers	“Specifying a USB Modem Interface on J Series Routers” on page 93
input-priority-map	Gigabit Ethernet IQ interfaces	“Specifying an Input Priority Map” on page 758
interface-type <i>type</i>	Channelized IQ and IQE interfaces, ISDN interfaces	“Channelized Interfaces” on page 385 and “Configuring ISDN Physical Interface Properties” on page 821
invert-data	DS0, E1, E3, and T1 interfaces	“Configuring E1 Data Inversion” on page 546, “Configuring E3 Data Inversion” on page 555, and “Configuring T1 Data Inversion” on page 563
isdn-options	ISDN interfaces	“Configuring ISDN Logical Interface Properties” on page 823
keepalives <down-count <i>number</i> <interval <i>seconds</i> > <up-count <i>number</i> >	Aggregated SONET/SDH, DS0, E1, E3, SONET/SDH, T1, and T3 interfaces	“Configuring Keepalives” on page 126

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
<code>lACP mode</code>	Aggregated Ethernet interfaces	“Configuring Aggregated Ethernet LACP” on page 627
<code>line-encoding (ami b8zs)</code>	T1 interfaces	“Configuring T1 Line Encoding” on page 564
<code>line-protocol protocol</code>	Serial interfaces (EIA-530, V.35, and X.21)	“Configuring the Serial Line Protocol” on page 265
<code>line-rate line-rate</code>	ATM interfaces on J Series routers	“Configuring SHDSL Operating Mode on an ATM Physical Interface” on page 364
<code>linear-red-profile profile-name</code>	ATM2 interfaces	“Configuring ATM2 IQ VC Tunnel CoS Components” on page 339
<code>linear-red-profiles profile-name</code>	ATM2 interfaces	“Configuring ATM2 IQ VC Tunnel CoS Components” on page 339
<code>link-layer-overhead percent</code>	AS PIC link services IQ interfaces (lsq)	<i>JUNOS Services Interfaces Configuration Guide</i>
<code>link-mode mode</code>	Management Ethernet (fxp0 or em0) and Fast Ethernet interfaces	“Configuring the Link Characteristics” on page 97
<code>link-speed speed</code>	Aggregated Ethernet and aggregated SONET/SDH interfaces	“Configuring Aggregated Ethernet Link Speed” on page 634 and “Configuring Aggregated SONET/SDH Link Speed” on page 883
<code>master-only;</code>	Management Ethernet (fxp0 or em0) and Fast Ethernet interfaces	“Configuring a Consistent Management IP Address” on page 775
<code>lmi (Frame Relay) lmi-options</code>	Interfaces with Frame Relay encapsulation	“Configuring Tunable Keepalives for Frame Relay LMI” on page 378 and <i>JUNOS Services Interfaces Configuration Guide</i>
<code>lmi-type (ansi itu)</code>	Link services interfaces and interfaces with Frame Relay encapsulation	“Configuring Frame Relay Keepalives” on page 378
<code>lmi (Ethernet OAM)</code>	OAM CFM Ethernet Local Management Interface	“Configuring Ethernet Local Management Interface” on page 690
<code>local-name name</code>	Interfaces with PPP encapsulation	“Configuring the PPP Challenge Handshake Authentication Protocol” on page 112
<code>lockout</code>	SONET/SDH interfaces	“Configuring APS and MSP” on page 859
<code>log-prefix prefix-number</code>	Adaptive services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
<code>(long-buildout no-long-buildout)</code>	T3 interfaces	“Configuring the T3 Line Buildout” on page 575

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
(loop-timing no-loop-timing)	Channelized IQ interfaces	“Configuring the Channelized T3 Loop Timing” on page 576
loopback <i>mode</i>	DS0, E1, E3, T1, T3, Ethernet, SONET/SDH, ATM interfaces on J Series routers, serial interfaces (EIA-530, V.35, and X.21), and 10-Gigabit Ethernet interfaces in WAN PHY mode	“Configuring E1 Loopback Capability” on page 547, “Configuring E3 Loopback Capability” on page 555, “Configuring Ethernet Loopback Capability” on page 593, “Configuring Serial Loopback Capability” on page 275, “Configuring Channelized IQ and IQE SONET/SDH Loop Timing” on page 852, “Configuring T1 Loopback Capability” on page 565, “Configuring T3 Loopback Capability” on page 576, “Configuring SHDSL Operating Mode on an ATM Physical Interface” on page 364, and “Configuring SONET Options for 10-Gigabit Ethernet Interfaces” on page 872
(loopback no-loopback)	Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces	“Configuring Ethernet Loopback Capability” on page 593
loss-priority (high low)	Gigabit Ethernet IQ interfaces	“Specifying an Output Priority Map” on page 759
low-plp-max-threshold <i>percent</i>	ATM2 interfaces	“Configuring ATM2 IQ VC Tunnel CoS Components” on page 339
low-plp-threshold <i>percent</i>	ATM2 interfaces	“Configuring ATM2 IQ VC Tunnel CoS Components” on page 339
lsq-failure-options	Link services IQ (lsq) interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
mac <i>mac-address</i>	Management Ethernet interface (fxp0 or em0)	“Configuring the MAC Address on the Management Ethernet Interface” on page 777
(mac-learn-enable no-mac-learn-enable)	Gigabit Ethernet IQ and IQE, Tri-Rate Ethernet copper, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	“Configuring MAC Address Filtering” on page 761
maximum-vcs <i>maximum-vcs</i>	ATM interfaces	“Configuring the Maximum Number of ATM1 VCs on a VP” on page 300
minimum-links <i>number</i>	Multilink, link services, and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
mip-half-function	Connectivity Fault Management	“Configuring Maintenance Intermediate Points” on page 683
mlfr-uni-nni-bundle-options <i>bundle-options</i>	Link services and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
modem-options	For USB ports (umd0) on J4350 and J6350 Services Routers	"Specifying a USB Modem Interface on J Series Routers" on page 93
mpls	10-Gigabit Ethernet interfaces in WAN PHY mode and ATM and SONET/SDH interfaces in passive monitoring mode	"Removing MPLS Labels from Incoming Packets" on page 294 and "Removing MPLS Labels from Incoming Packets" on page 874 and "Configuring SONET Options for 10-Gigabit Ethernet Interfaces" on page 872
mrru bytes	Link services and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
mtu bytes	All interfaces, except management Ethernet (fxp0 or em0), loopback, multilink, and multicast tunnel	"Configuring the Media MTU" on page 98
multiservice-options	Adaptive services, monitoring services, and collector interfaces	"Configuring Multiservice Physical Interface Properties" on page 138
n391 number	Link services and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
n392 number	Link services and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
n393 number	Link services and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
neighbor address	SONET/SDH interfaces	"Configuring APS and MSP" on page 859
no-gratuitous-arp-request	Ethernet interfaces	"Configuring Gratuitous ARP" on page 596
no-keepalives	Interfaces with PPP, Frame Relay, or Cisco High-level Data Link Control (HDLC) encapsulation	"Configuring Keepalives" on page 126
no-partition	Channelized IQ interfaces	"Channelized Interfaces" on page 385
no-termination-request	link services IQ (lsq) interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
oam-liveness	ATM interfaces	"Configuring the OAM F4 Cell Flows" on page 315
oam-period (seconds disable)	ATM interfaces	"Defining the ATM OAM F5 Loopback Cell Period" on page 329
oc-slice oc-slice-range	Channelized OC12 IQ interfaces	"Configuring Channelized OC12/STM4 Interfaces" on page 423
open-timeout seconds	Adaptive services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
optics-options	Gigabit Ethernet dense wavelength-division multiplexing (DWDM) interfaces	“Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength” on page 779
output-priority-map	Gigabit Ethernet IQ interfaces	“Specifying an Output Priority Map” on page 759
overflow (discard tag)	All interfaces, except ATM, channelized E1, E1, Fast Ethernet, Gigabit Ethernet, and channelized IQ	“Configuring Receive and Transmit Leaky Bucket Properties” on page 129 and “Configuring Receive and Transmit Leaky Bucket Properties on SONET/SDH Interfaces” on page 876
paired-group <i>group-name</i>	SONET/SDH interfaces	“Configuring APS and MSP” on page 859
partition <i>partition-number</i>	Channelized IQ interfaces	“Channelized Interfaces” on page 385
passive	Interfaces with PPP encapsulation	“Configuring the PPP Challenge Handshake Authentication Protocol” on page 112
passive-monitor-mode	SONET/SDH interfaces	“Enabling Passive Monitoring on SONET/SDH Interfaces” on page 874
path-trace <i>trace-string</i>	10-Gigabit Ethernet interfaces in WAN PHY mode and SONET/SDH interfaces	“Configuring the SONET/SDH Path Trace Identifier” on page 853 and “Configuring SONET Options for 10-Gigabit Ethernet Interfaces” on page 872
(payload-scrambler no-payload-scrambler)	E3, SONET/SDH, and T3 interfaces	“Configuring E3 and T3 Parameters on ATM Interfaces” on page 337, “Configuring E3 HDLC Payload Scrambling” on page 557, “Configuring SONET/SDH HDLC Payload Scrambling” on page 854, “Configuring T3 HDLC Payload Scrambling” on page 578, and “Examples: Configuring T3 Interfaces” on page 579
periodic <i>interval</i>	Aggregated Ethernet interfaces	“Configuring Aggregated Ethernet LACP” on page 627
per-unit-scheduler	IQ interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
pfc	Interfaces with PPP encapsulation	“Configuring the PPP Protocol Field Compression” on page 121
pic-type (atm1 atm2)	ATM2 IQ interfaces	“Configuring the ATM PIC Type” on page 295
plp1 <i>cells</i>	ATM2 interfaces	“Configuring ATM2 IQ VC Tunnel CoS Components” on page 339

Table 8: Statements for Physical Interface Properties (*continued*)

Statement	Interface Types	Usage Guidelines
plp-to-clp	ATM2 IQ interfaces	“Enabling the PLP Setting to Be Copied to the CLP Bit” on page 348
policer <i>cos-policer-name</i>	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	“Configuring Gigabit Ethernet Policers” on page 757
pop-all-labels	ATM and SONET/SDH interfaces in passive monitoring mode	“Removing MPLS Labels from Incoming Packets” on page 294 and “Removing MPLS Labels from Incoming Packets” on page 874
ppp-options	Interfaces with PPP encapsulation	“Configuring the PPP Challenge Handshake Authentication Protocol” on page 112
premium	Gigabit Ethernet IQ interfaces	“Configuring Gigabit Ethernet Policers” on page 757 and “Specifying an Output Priority Map” on page 759
primary <i>sp-fpc/pic/port</i>	Redundant interfaces for adaptive services interfaces (<i>rsp</i>)	<i>JUNOS Services Interfaces Configuration Guide</i>
priority (high low)	ATM2 IQ interfaces	“Configuring ATM2 IQ VC Tunnel CoS Components” on page 339
priority <i>number</i>	ISDN interfaces	“Configuring an ISDN Dialer Interface as a Backup Interface” on page 826
promiscuous-mode	ATM2 IQ interfaces	“Configuring ATM Cell-Relay Promiscuous Mode” on page 296
protect-circuit <i>group-name</i>	SONET/SDH interfaces	“Configuring APS and MSP” on page 859
queue-depth <i>cells</i>	ATM2 interfaces	“Configuring ATM2 IQ VC Tunnel CoS Components” on page 339
queue-length <i>number</i>	ATM1 interfaces	“Configuring the ATM1 Queue Length” on page 325
rate <i>percentage</i>	All interfaces, except ATM, channelized E1, E1, Fast Ethernet, Gigabit Ethernet, and channelized IQ	“Configuring Receive and Transmit Leaky Bucket Properties” on page 129 and “Configuring Receive and Transmit Leaky Bucket Properties on SONET/SDH Interfaces” on page 876
receive-bucket	All interfaces, except ATM, Fast Ethernet, and Gigabit Ethernet	“Configuring Receive and Transmit Leaky Bucket Properties” on page 129
red-differential-delay <i>milliseconds</i>	Link services and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
redundancy-options	Redundant interfaces for adaptive services interfaces (rsp-)	<i>JUNOS Services Interfaces Configuration Guide</i>
remote-loopback-respond	T1 interfaces	“Configuring the T1 Remote Loopback Response” on page 564
request	SONET/SDH interfaces	“Configuring APS and MSP” on page 859
required-depth <i>number</i>	ATM and SONET/SDH interfaces in passive monitoring mode	“Removing MPLS Labels from Incoming Packets” on page 294 and “Removing MPLS Labels from Incoming Packets” on page 874
revert-time <i>seconds</i>	SONET/SDH interfaces	“Configuring APS and MSP” on page 859
rfc-2615	SONET/SDH interfaces	“Configuring SONET/SDH RFC 2615 Support” on page 855
rts (assert de-assert normal)	Serial interfaces (EIA-530 and V.35)	“Configuring the Serial Signal Handling” on page 271
rts-polarity (negative positive)	Serial interfaces (EIA-530 and V.35)	“Configuring Serial Signal Polarities” on page 274
rtvbr peak rate sustained <i>rate</i> burst <i>length</i>	ATM interfaces	“Configuring ATM2 IQ Real-Time VBR” on page 321
scheduler-maps <i>map-name</i>	ATM2 interfaces	“Configuring ATM2 IQ VC Tunnel CoS Components” on page 339
schedulers <i>number</i>	Ethernet IQ2 and IQ2-E PICs port interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
secondary sp-fpc/pic/port	Redundant interfaces for adaptive services interfaces (rsp-)	<i>JUNOS Services Interfaces Configuration Guide</i>
services-options	Services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
serial-options	Serial interfaces (EIA-530, V.35, and X.21)	“Configuring Serial Interfaces” on page 263
services priority-level	Adaptive services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
shdsl-options	ATM interfaces on J Series routers	“Configuring SHDSL Operating Mode on an ATM Physical Interface” on page 364
size	All	“Tracing Operations of the Interface Process” on page 241
shaping	ATM interfaces	“Defining the ATM Traffic-Shaping Profile” on page 319

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
<code>snext margin</code>	ATM interfaces on J Series routers	“Configuring SHDSL Operating Mode on an ATM Physical Interface” on page 364
<code>snr-margin</code>	ATM interfaces on J Series routers	“Configuring SHDSL Operating Mode on an ATM Physical Interface” on page 364
<code>sonet-options</code>	SONET/SDH interfaces	“Configuring SONET/SDH Physical Interface Properties” on page 844
<code>source-address-filter mac-address</code>	Aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, and Gigabit Ethernet interfaces	“Enabling Ethernet MAC Address Filtering” on page 591
<code>(source-filtering no-source-filtering)</code>	Aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, Gigabit Ethernet IQ and IQE, and Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	“Enabling Ethernet MAC Address Filtering” on page 591
<code>speed (10m 100m 1g oc3 oc12 oc48)</code>	Management Ethernet interface (<code>fxp0</code> or <code>em0</code>), Tri-Rate Ethernet copper interfaces, and 12-port and 48-port Fast Ethernet interfaces SONET/SDH PICs with SFP	“Configuring the Interface Speed” on page 122 and “SONET/SDH Interface” on page 124
<code>spid1spid2</code>	ISDN interfaces	“Configuring ISDN Physical Interface Properties” on page 821 and “Configuring an ISDN Dialer Interface as a Backup Interface” on page 826
<code>stacked-vlan-tagging</code>	Gigabit Ethernet IQ interfaces	“Configuring the Management Ethernet Interface” on page 775
<code>start-end-flag (filler shared)</code>	DS0, E1, E3, T1, and T3 interfaces	“Configuring E1 Start and End Flags” on page 548, “Configuring the E3 Start and End Flags” on page 557, “Configuring T1 Start and End Flags” on page 567, and “Configuring T3 Start and End Flags” on page 579
<code>switching-mode (bidirectional unidirectional)</code>	Unchannelized OC3, OC12, and OC48 SONET/SDH interfaces on T Series routers	“Configuring Switching Between the Working and Protect Circuits” on page 866
<code>syslog</code>	Adaptive services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
<code>(syslog no-syslog)</code>	Adaptive services, monitoring services, and collector interfaces	“Configuring Multiservice Physical Interface Properties” on page 138
<code>t1-options</code>	T1 interfaces	“Configuring T1 Interfaces” on page 559

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
t3-options	T3 interfaces	“Configuring T3 Interfaces” on page 569
t310	ISDN interfaces	“Configuring ISDN Physical Interface Properties” on page 821
t391 seconds	Link services and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
t392 number	Link services and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
tag-protocol-id (first-call power-up)	ISDN interfaces	“Configuring ISDN Physical Interface Properties” on page 821
threshold bytes	All interfaces, except ATM, channelized E1, E1, Fast Ethernet, Gigabit Ethernet, and channelized IQ	“Configuring Receive and Transmit Leaky Bucket Properties” on page 129 and “Configuring Receive and Transmit Leaky Bucket Properties on SONET/SDH Interfaces” on page 876
timeslots time-slot-range	Channelized T1 IQ and channelized E1 IQ interfaces	“Channelized Interfaces” on page 385
tm (ignore normal require)	Serial interfaces (EIA-530)	“Configuring the Serial Signal Handling” on page 271
tm-polarity (negative positive)	Serial interfaces (EIA-530)	“Configuring Serial Signal Polarities” on page 274
traceoptions	All	“Tracing Operations of an Individual Router Interface” on page 137 and “Tracing Operations of the Interface Process” on page 241
transmit-bucket	All interfaces, except ATM, Fast Ethernet, Tri-Rate Ethernet copper, and Gigabit Ethernet	“Configuring Receive and Transmit Leaky Bucket Properties” on page 129
transmit-clock invert	Serial interfaces (EIA-530, V.35, and X.21)	“Configuring the Serial Clocking Mode” on page 269
transmit-weight (cells number percent number)	ATM2 IQ interfaces	“Configuring ATM2 IQ VC Tunnel CoS Components” on page 339
(traps no-traps)	All	“Enabling or Disabling SNMP Notifications on Physical Interfaces” on page 139
trigger defect ignore defect hold-time up milliseconds down milliseconds;	10-Gigabit Ethernet interfaces in WAN PHY mode and ATM over SONET/SDH and SONET/SDH interfaces	“Configuring SONET/SDH Defect Triggers to Be Ignored” on page 855 and “Configuring SONET Options for 10-Gigabit Ethernet Interfaces” on page 872

Table 8: Statements for Physical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
(unframed no-unframed)	E3 IQ interfaces	“Configuring E3 IQ and IQE Unframed Mode” on page 558
unidirectional	10-Gigabit Ethernet interfaces on: <ul style="list-style-type: none"> ■ MX960 4-Port 10-Gigabit Ethernet DPC ■ T Series 10-Gigabit Ethernet IQ2 PIC ■ T Series 10-Gigabit Ethernet IQ2E PIC 	“Enabling Unidirectional Traffic Flow on Physical Interfaces” on page 139
vbr <i>peak rate</i> sustained <i>rate burst length</i>	ATM interfaces	“Defining the ATM Traffic-Shaping Profile” on page 319
vc-cos-mode (alternate strict)	ATM2 interfaces	“Configuring ATM2 IQ VC Tunnel CoS Components” on page 339
vlan-tagging	Fast Ethernet, Tri-Rate Ethernet copper, and Gigabit Ethernet interfaces	“Configuring 802.1Q VLANs” on page 599
vlan-vci-tagging	Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces	“Configuring ATM-to-Ethernet Interworking” on page 229
vpi <i>vpi-identifier</i>	ATM interfaces	“Configuring ATM Cell-Relay Promiscuous Mode” on page 296 and “Configuring the Maximum Number of ATM1 VCs on a VP” on page 300
vtmapping	Channelized STM1 interfaces	“Configuring Virtual Tributary Mapping of Channelized STM1 Interfaces” on page 472
wavelength <i>nm</i>	Gigabit Ethernet dense wavelength-division multiplexing (DWDM) interfaces	“Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength” on page 779
working-circuit <i>group-name</i>	SONET/SDH interfaces	“Configuring APS and MSP” on page 859
yellow-differential-delay <i>milliseconds</i>	Link services and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
(z0-increment no-z0-increment)	SONET/SDH interfaces	“Configuring an Incrementing STM ID” on page 850

Specifying an Aggregated Interface

The M Series, MX Series, and T Series routers support aggregated interfaces.

You specify aggregated interfaces by assigning a number for the aggregated interface. For aggregated Ethernet interfaces, configure **aex** as in the following example:

```
[edit interfaces]
ae0 {
  ...
}
```

For aggregated SONET/SDH interfaces, configure **asx** as in the following example:

```
[edit interfaces]
as0 {
  ...
}
```

The maximum number of aggregated Ethernet interfaces is 128, and the assigned number can be from 0 through 127. The maximum number of aggregated Ethernet interfaces (LAG bundles) on all MX Series routers is 480, and the assigned number can be from 0 through 479. The maximum number of aggregated SONET interfaces is 16, and the assigned number can be from 0 through 15. You should not mix SONET and SDH modes on the same aggregated interface.



NOTE: SONET/SDH aggregation is proprietary to the JUNOS Software and might not work with other software.

If you are configuring VLANs for aggregated Ethernet interfaces, you must include the **vlan-tagging** statement at the **[edit interfaces aex]** hierarchy level to complete the association.

For more information, see “Configuring Aggregated Ethernet Interfaces” on page 623 and “Configuring Aggregated SONET/SDH Interfaces” on page 881.

Specifying a USB Modem Interface on J Series Routers

The J Series routers contain two USB ports controlled by a single USB controller. One USB port can support USB devices, while the other one can act as a USB modem.

The USB modem provides a dial-in remote management interface, and supports dialer interface features by sharing the same dial pool as a dialer interface. The dial pool allows the logical dialer interface (**dln**) and the physical interface (**umd0**) to be bound together dynamically on a per-call basis.

The following dialer interface features are supported by the USB modem interface:

- Encapsulation PPP
- CoS
- NAT
- Interface statistics
- Packet capture
- GRE tunnel

- Stateful firewall
- Traffic sampling

To configure a USB modem interface, include the following statements at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
umd0 {
  dialer-options {
    pool pool-name <priority priority>;
  }
  modem-options {
    dialin (console | routable);
    init-command-string initialization-command-string;
  }
}
```

The pool name specified at the **[edit interfaces umd0 dialer-options pool]** hierarchy level must be the same as the pool name specified at the **[edit interfaces dln unit *logical-unit-number* dialer-options pool]** hierarchy level.

Configure the USB modem to operate as a dial-in WAN backup interface by including the **dialin** statement and specifying the **routable** option. If the USB modem is to be used as a dial-in console, specify the **console** option in the **dialin** statement.

When the Services Router applies the modem AT commands configured in the **init-command-string** statement or the default sequence of initialization commands to the modem, it compares them to the initialization commands already configured on the modem and makes the following changes:

- If the commands are the same, the router overrides the existing modem values that do not match. For example, if the initialization commands on the modem include **S0 = 0** and the router's **init-command-string** configuration includes **S0 = 2**, the Services Router applies **S0 = 2**.
- If the initialization commands on the modem do not include a command in the router's **init-command-string** statement configuration, the router adds it. For example, if the **init-command-string** statement includes the command **L2**, but the modem commands do not include it, the router adds **L2** to the initialization commands configured on the modem.

Include the following statements at the **[edit interfaces dln]** hierarchy level to support a minimum configuration for a dialer interface connected to a USB modem:

```
[edit interfaces dln]
encapsulation ppp;
unit logical-unit-number;
dialer-options {
  dial-string dial-string-numbers;
  pool pool-name <priority priority>;
}
ppp-options {
  chap;
  access-profile name;
```

```

    local-name name;
    passive;
  }
  family inet {
    mtu bytes;
    address address {
      destination address;
    }
  }
}

```

For more information about configuring dial-in, see “Configuring Dial-In and Callback” on page 832.

Specifying OC768-over-OC192 Mode

The T Series routers support OC768-over-OC192 mode on the 4-port OC192c PIC. In OC768-over-OC192 mode, four OC192 links are aggregated into one OC768 link with one logical interface. This single interface achieves data rates of approximately 40 Gbps. OC768 optics are expensive, and most long-distance networks currently use fiber optics and regenerators that cannot carry OC768 SONET. When you create an OC768 pipe as a large data pipe running over existing infrastructures, you transfer network traffic without link bonding or load sharing over parallel links. Load sharing is automatically accomplished in the JUNOS Software using a proprietary method, and does not need to be manually configured.

The following limitations apply to OC768-over-OC192 mode:

- The maximum difference in delay between all links in the bundle is 8 μ (microseconds), equivalent to approximately 1.5 km maximum difference in length between the longest and shortest fiber pairs.
- If a single link in the bundle fails, the whole bundle fails. If link redundancy is required, implement an aggregated SONET/SDH bundle instead.
- Only routers that contain 4-port OC192 PICs can operate in OC768-over-OC192 mode.

To configure the 4-port OC192 PIC to operate in OC768-over-OC192 mode on a TX Matrix router, include the **aggregate-ports** statement at the **[edit chassis lcc *lcc-number* fpc *slot-number* pic *pic-number*]** hierarchy level:

```

[edit chassis]
lcc lcc-number {
  fpc slot-number {
    pic pic-number {
      aggregate-ports;
    }
  }
}
...

```

To configure the 4-port OC192 PIC to operate in OC768-over-OC192 mode on a T640 router, include the **aggregate-ports** statement at the **[edit chassis fpc *slot-number* pic *pic-number*]** hierarchy level:

```
[edit chassis]
fpc slot-number {
  pic pic-number {
    aggregate-ports;
  }
}
...
```

When you configure the 4-port OC192 PIC for OC768-over-OC192 mode, only port 0 (the first port) needs be configured as the OC768 port.

To display logical and physical interface information, use the operational mode command `show interfaces so-fpc/pic/port extensive`. When this command is used for the 4-port OC192 PIC configured for OC768-over-OC192 mode, only port 0 (`so-fpc/pic/0`) is displayed. This port is displayed as **OC768**.

Adding an Interface Description to the Configuration

You can include a text description of each physical interface in the configuration file. Any descriptive text you include is displayed in the output of the `show interfaces` commands, and is also exposed in the `ifAlias` Management Information Base (MIB) object. It has no impact on the interface's configuration. To add a text description, include the `description` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]
description text;
```

The description can be a single line of text. If the text contains spaces, enclose it in quotation marks.



NOTE: You can configure the extended DHCP relay to include the interface description in the option 82 Agent Circuit ID suboption. See *Enabling and Disabling Insertion of Option 82 Information* in the *JUNOS Subscriber Access Configuration Guide*.

For information about describing logical units, see “Adding a Logical Unit Description to the Configuration” on page 156.

Example: Adding an Interface Description to the Configuration

Add a description to a Fast Ethernet interface:

```
[edit interfaces]
user@host#
set fe-0/0/1 description "Backbone connection to PHL01"
[edit interfaces]
user@host#
show
fe-0/0/1 {
  description "Backbone connection to PHL01";
  unit 0 {
    family inet {
      address 192.168.0.1/30;
    }
  }
}
```



```

    }
  }
}

```

To display the description from the router CLI, use the **show interfaces** command:

```

user@host>
show interfaces fe-0/0/1
Physical interface: fe-0/0/1, Enabled, Physical link is Up
  Interface index: 129, SNMP ifIndex: 23
  Description: Backbone connection to PHL01
  ...

```

To display the interface description from the interfaces MIB, use the **snmpwalk** command from a server. To isolate information for a specific interface, search for the interface index shown in the **SNMP ifIndex** field of the **show interfaces** command output. The **ifAlias** object is in **ifXTable**.

```

user-server>snmpwalk host-fxp0.mylab public ifXTable | grep -e '\.23'
snmpwalk host-fxp0.mylab public ifXTable | grep -e '\.23'
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifName.23 = fe-0/0/1
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifInMulticastPkts.23 = Counter32: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifInBroadcastPkts.23 = Counter32: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifOutMulticastPkts.23 = Counter32: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifOutBroadcastPkts.23 = Counter32: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCInOctets.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCInUcastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCInMulticastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCInBroadcastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCOctets.23 = Counter64: 42
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCOUcastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCOMulticastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHCOBroadcastPkts.23 = Counter64: 0
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifLinkUpDownTrapEnable.23 = enabled(1)
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifHighSpeed.23 = Gauge32: 100
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifPromiscuousMode.23 = false(2)
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifConnectorPresent.23 = true(1)
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifAlias.23 = Backbone connection to PHL01
ifMIB.ifMIBObjects.ifXTable.ifXEntry.ifCounterDiscontinuityTime.23 = Timeticks:
(0) 0:00:00.00

```

Configuring the Link Characteristics

By default, the router's management Ethernet interface, **fxp0** or **em0**, autonegotiates whether to operate in full-duplex or half-duplex mode. Fast Ethernet and J Series router Gigabit Ethernet interfaces can operate in either full-duplex or half-duplex mode, and all other interfaces can operate only in full-duplex mode. For Gigabit Ethernet, the link partner must also be set to full duplex.



NOTE: When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.



NOTE: When you manually configure Fast Ethernet interfaces on the M Series and T Series routers, link mode and speed must both be configured. If both these values are not configured, the router uses autonegotiation for the link and ignores the user-configured settings.

To explicitly configure an Ethernet interface to operate in either full-duplex or half-duplex mode, include the `link-mode` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
link-mode (full-duplex | half-duplex);
```

Configuring the Media MTU

The default media MTU size used on a physical interface depends on the encapsulation used on that interface. In some cases, the default IP Protocol MTU depends on whether the protocol used is IP version 4 (IPv4) or International Organization for Standardization (ISO). Table 9 on page 98 through Table 15 on page 102 list the media and protocol MTU sizes by interface type, and Table 18 on page 104 lists the encapsulation overhead by encapsulation type.

Table 9: Media MTU Sizes by Interface Type for M5, M7i with CFEB, M10, M10i with CFEB, M20, and M40 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	9192	1500
E3/T3	4474	9192	4470
Fast Ethernet	1514	9192 (4-port) 1532 (8-port) 1532 (12-port)	1500 (IPv4) 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4) 1497 (ISO)
Serial	1504	9192	1500 (IPv4) 1497 (ISO)
SONET/SDH	4474	9192	4470

Table 10: Media MTU Sizes by Interface Type for M40e Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	4500	1500
E3/T3	4474	4500	4470
		9192 (4-port)	
E3/DS3 IQ	4474	9192	4470
Fast Ethernet	1514	4500	1500 (IPv4) 1497 (ISO)
Gigabit Ethernet	1514	9192 (1- or 2-port)	1500 (IPv4) 1497 (ISO)
		9192 (4-port)	
Serial	1504	9192	1500 (IPv4) 1497 (ISO)
SONET/SDH	4474	4500 (1-port nonconcatenated)	4470
		9192 (4-port OC3)	
		9192 (4-port OC3c)	
		4500 (1-port OC12)	
		4500 (4-port OC12)	
		4500 (4-port OC12c)	
		4500 (1-port OC48)	
		9192 (2-port OC3)	
		9192 (2-port OC3c)	
		9192 (1-port OC12c)	
		9192 (1-port OC48c)	
		4500 (1-port OC192)	
		9192 (1-port OC192c)	

Table 11: Media MTU Sizes by Interface Type for M160 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Adaptive Services (MTU size not configurable)	9192	N/A	N/A
ATM	4482	9192	4470
E1/T1	1504	4500	1500
E3/T3	4474	4500	4470
E3/DS3 IQ	4474	9192	4470
Fast Ethernet	1514	4500	1500 (IPv4)1497 (ISO)
Gigabit Ethernet	1514	9192 (1- or 2-port) 4500 (4-port)	1500 (IPv4)1497 (ISO)
Serial	1504	9192	1500 (IPv4)1497 (ISO)
SONET/SDH	4474	4500 (1-port nonconcatenated) 9192 (1- or 2-port) 4500 (4-port)	4470

Table 12: Media MTU Sizes by Interface Type for M7i with CFEB-E, M10i with CFEB-E, M320 and M120 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM2 IQ	4482	9192	4470
Channelized DS3 IQ	4471	4500	4470
Channelized E1 IQ	1504	4500	1500
Channelized OC12 IQ	4474	9192	4470
Channelized STM1 IQ	4474	9192	4470
DS3	4471	4500	4470
E1	1504	4500	1500
E3 IQ	4471	4500	4470

Table 12: Media MTU Sizes by Interface Type for M7i with CFEB-E, M10i with CFEB-E, M320 and M120 Routers *(continued)*

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Fast Ethernet	1514	9192 (4-port) 1532 (8-, 12- and 48-port)	1500 (IPv4) 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4) 1497 (ISO)
SONET/SDH	4474	9192	4470
T1	1504	4500	1500

Table 13: Media MTU Sizes by Interface Type for T320 Routers

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM	4482	9192	4470
ATM2 IQ	4482	9192	4470
Channelized OC12 IQ	4474	9192	4470
Channelized STM1 IQ	4474	9192	4470
DS3	4471	4500	4470
Fast Ethernet	1514	4500 (4-port) 1532 (12- and 48-port)	1500 (IPv4) 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4) 1497 (ISO)
SONET/SDH	4474	9192	4470

Table 14: Media MTU Sizes by Interface Type for T640 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ATM2 IQ	4482	9192	4470
48-port Fast Ethernet	1514	1532	1500 (IPv4) 1497 (ISO)
Gigabit Ethernet	1514	9192	1500 (IPv4) 1497 (ISO)
SONET/SDH	4474	9192	4470

Table 15: Media MTU Sizes by Interface Type for J2300 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Fast Ethernet (10/100)	1514	9192	1500
G.SHDSL	4482	9150	4470
ISDN BRI	1504	4092	1500
Serial	1504	9150	1500
T1 or E1	1504	9150	1500

Table 16: Media MTU Sizes by Interface Type for J4300 and J6300 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
ADSL2 + PIM	4482	9150	4470
Dual-port Fast Ethernet (10/100) PIM	1514	9192	1500
Dual-port Serial PIM	1504	9150	1500
Dual-port T1 or E1 PIM	1504	9150	1500
Dual-port Channelized T1/E1 PIM (channelized to DS0s)	1504	4500	1500
Dual-port Channelized T1/E1 PIM (clear channel T1 or E1)	1504	9150	1500
Fast Ethernet (10/100) built-in interface	1514	9192	1500
G.SHDSL PIM	4482	9150	4470
4-port ISDN BRI PIM	1504	4092	1500
T3 (DS3) or E3 PIM	4474	9192	4470

Table 17: Media MTU Sizes by Interface Type for J4350 and J6350 Platforms

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
4-port ISDN BRI PIM	1504	4092	1500
ADSL2 + PIM	4482	9150	4470

Table 17: Media MTU Sizes by Interface Type for J4350 and J6350 Platforms *(continued)*

Interface Type	Default Media MTU (Bytes)	Maximum MTU (Bytes)	Default IP Protocol MTU (Bytes)
Dual-port Fast Ethernet (10/100) PIM	1514	9192	1500
Dual-port Serial PIM	1504	9150	1500
Dual-port T1 or E1 PIM	1504	9150	1500
Dual-port Channelized T1/E1 PIM (channelized to DS0s)	1504	4500	1500
Dual-port Channelized T1/E1 PIM (clear channel T1 or E1)	1504	9150	1500
4-port Fast Ethernet (10/100) ePIM	1518	1518	1500
Gigabit Ethernet (10/100/1000) built-in interface	1514	9018	1500
Gigabit Ethernet (10/100/1000) Enhanced Physical Interface Module (ePIM)	1514	9018	1500
Gigabit Ethernet (10/100/1000) SFP ePIM	1514	9018	1500
G.SHDSL PIM	4482	9150	4470
T3 (DS3) or E3 PIM	4474	9192	4470



NOTE: On Gigabit Ethernet ePIMs in J4350 and J6350 Services Routers, you can configure a maximum transmission unit (MTU) size of only 9018 bytes even though the CLI indicates that you can configure an MTU of up to 9192 bytes. If you configure an MTU greater than 9018 bytes, the router does not accept the configuration and generates a system log error message similar to the following:

```
/kernel: ge-0/0/0: Illegal media change. MTU invalid: 9192. Max MTU supported on
this PIC: 9018
```

On 4-port Fast Ethernet ePIMs in J4350 and J6350 Services Routers, you can configure a maximum transmission unit (MTU) size of only 1518 bytes even though the CLI indicates that you can configure an MTU of up to 9192 bytes. If you configure an MTU greater than 1518 bytes, the router does not accept the configuration and generates a system log error message similar to the following:

```
/kernel: fe-3/0/1: Illegal media change. MTU invalid: 9192. Max MTU supported on
this PIC: 1518
```

Table 18: Encapsulation Overhead by Encapsulation Type

Interface Encapsulation	Encapsulation Overhead (Bytes)
802.1Q/Ethernet 802.3	21
802.1Q/Ethernet Subnetwork Access Protocol (SNAP)	26
802.1Q/Ethernet version 2	18
ATM Cell Relay	4
ATM permanent virtual connection (PVC)	12
Cisco HDLC	4
Ethernet 802.3	17
Ethernet circuit cross-connect (CCC) and virtual private LAN service (VPLS)	4
Ethernet over ATM	32
Ethernet SNAP	22
Ethernet translational cross-connect (TCC)	18
Ethernet version 2	14
Extended virtual local area network (VLAN) CCC and VPLS	4
Extended VLAN TCC	22
Frame Relay	4
PPP	4

Table 18: Encapsulation Overhead by Encapsulation Type (*continued*)

Interface Encapsulation	Encapsulation Overhead (Bytes)
VLAN CCC	4
VLAN VPLS	4
VLAN TCC	22

The default media MTU is calculated as follows:

$$\text{Default media MTU} = \text{Default IP MTU} + \text{encapsulation overhead}$$

When you are configuring point-to-point connections, the MTU sizes on both sides of the connections must be the same. Also, when you are configuring point-to-multipoint connections, all interfaces in the subnet must use the same MTU size.



NOTE: The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the media MTU. For example, the media MTU for a Gigabit Ethernet Version 2 interface is specified as 1514 bytes, but the largest possible frame size is actually 1518 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

The physical MTU for Ethernet interfaces does not include the 4-byte frame check sequence (FCS) field of the Ethernet frame.

A SONET/SDH interface operating in concatenated mode has a “c” added to the rate descriptor. For example, a concatenated OC48 interface is referred to as OC48c.

If you do not configure an MPLS MTU, the JUNOS Software derives the MPLS MTU from the physical interface MTU. From this value, the software subtracts the encapsulation-specific overhead and space for the maximum number of labels that might be pushed in the Packet Forwarding Engine. Currently, the software provides for three labels of four bytes each, for a total of 12 bytes.

In other words, the formula used to determine the MPLS MTU is the following:

$$\text{MPLS MTU} = \text{physical interface MTU} - \text{encapsulation overhead} - 12$$

If you configure an MTU value by including the `mtu` statement at the `[edit interfaces interface-name unit logical-unit-number family mpls]` hierarchy level, the configured value is used.

For information about configuring the encapsulation on an interface, see “Configuring Interface Encapsulation on Physical Interfaces” on page 106.

To modify the default media MTU size for a physical interface, include the `mtu` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]
mtu bytes;
```

If you change the size of the media MTU, you must ensure that the size is equal to or greater than the sum of the protocol MTU and the encapsulation overhead.



NOTE: Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

You configure the protocol MTU by including the `mtu` statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

Because tunnel services interfaces are considered logical interfaces, you cannot configure the MTU setting for the physical interface. This means you cannot include the `mtu` statement at the [edit interfaces *interface-name*] hierarchy level for the following interface types: generic routing encapsulation (**gr-**), IP-IP (**ip-**), loopback (**lo-**), link services (**ls-**), multilink services (**ml-**), and multicast (**pe-**, **pd-**). You can, however, configure the protocol MTU on tunnel interfaces, as described in “Setting the Protocol MTU” on page 191.

Configuring Interface Encapsulation on Physical Interfaces

Point-to-Point Protocol (PPP) encapsulation is the default encapsulation type for physical interfaces. You need not configure encapsulation for any physical interfaces that support PPP encapsulation. If you do not configure encapsulation, PPP is used by default. For physical interfaces that do not support PPP encapsulation, you must configure an encapsulation to use for packets transmitted on the interface.

You can optionally configure an encapsulation on a logical interface, which is the encapsulation used within certain packet types. For more information about logical interface encapsulation, see “Configuring the Encapsulation on a Logical Interface” on page 160.

This section contains the following topics:

- Configuring the Encapsulation on a Physical Interface on page 106
- Encapsulation Capabilities on page 110

Configuring the Encapsulation on a Physical Interface

By default, PPP is the encapsulation type for physical interfaces. To configure the encapsulation on a physical interface, include the `encapsulation` statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
```

```
encapsulation (atm-ccc-cell-relay | atm-pvc | cisco-hdlc | cisco-hdlc-ccc | cisco-hdlc-tcc
| ethernet-ccc | ethernet-over-atm | ethernet-tcc | ethernet-vpls |
extended-frame-relay-ccc | extended-frame-relay-ether-type-tcc |
extended-frame-relay-tcc | extended-vlan-ccc | extended-vlan-tcc | extended-vlan-vpls
| flexible-ethernet-services | flexible-frame-relay | frame-relay | frame-relay-ccc |
frame-relay-ether-type | frame-relay-ether-type-tcc | frame-relay-port-ccc | frame-relay-tcc
| multilink-frame-relay-uni-nni | ppp | ppp-ccc | ppp-tcc | vlan-ccc | vlan-vpls);
```

The physical interface encapsulation can be one of the following:

- **ATM CCC cell relay**—Connects two remote virtual circuits or ATM physical interfaces with a label-switched path (LSP). Traffic on the circuit is ATM cells.

You can configure an ATM1 Physical Interface Card (PIC) to use cell-relay accumulation mode (CAM). In this mode, the incoming cells (1 to 8 cells) are packaged into a single packet and forwarded to the LSP. Cell-relay accumulation mode is not supported on ATM2 PICs. You configure CAM as shown in the following example:

```
[edit chassis]
fpc 1 {
  pic 0 {
    atm-cell-relay-accumulation;
  }
}
```

For more information, see the *JUNOS System Basics Configuration Guide*.

- **ATM PVC**—Defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*. When you configure physical ATM interfaces with ATM PVC encapsulation, an RFC 2684-compliant ATM Adaptation Layer 5 (AAL5) tunnel is set up to route the ATM cells over a Multiprotocol Label Switching (MPLS) path that is typically established between two MPLS-capable routers using the Label Distribution Protocol (LDP).
- **Cisco HDLC**—E1, E3, SONET/SDH, T1, and T3 interfaces can use Cisco HDLC encapsulation. Two related versions are supported:
 - **CCC version (cisco-hdlc-ccc)**—The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the **ccc** family only.
 - **TCC version (cisco-hdlc-tcc)**—Similar to CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- **Ethernet over ATM**—As defined in RFC 1483, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, this encapsulation type allows ATM interfaces to connect to devices that support only bridged-mode protocol data units (BPDUs). The JUNOS Software does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet logical link control (LLC)/SNAP frames with IP or Address Resolution Protocol (ARP) in the payload, and drops the rest. For packets destined to the Ethernet local area network (LAN), a route lookup is done using the destination IP address. If the route lookup yields a full

address match, the packet is encapsulated with an LLC/SNAP and media access control (MAC) header, and the packet is forwarded to the ATM interface.

- Ethernet cross-connect—Ethernet interfaces without VLAN tagging can use Ethernet CCC encapsulation. Two related versions are supported:
 - CCC version (**ethernet-ccc**)—Ethernet interfaces with standard Tag Protocol ID (TPID) tagging can use Ethernet CCC encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.
 - TCC version (**ethernet-tcc**)—Similar to CCC, but used for circuits with different media on either side of the connection.

For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.

- VLAN CCC (**vlan-ccc**)—Ethernet interfaces with VLAN tagging enabled can use VLAN CCC encapsulation. VLAN CCC encapsulation supports TPID 0x8100 only. When you use this encapsulation type, you can configure the **ccc** family only.
- Extended VLAN cross-connect—Gigabit Ethernet interfaces with VLAN 802.1Q tagging enabled can use extended VLAN cross-connect encapsulation. (Ethernet interfaces with standard TPID tagging can use VLAN CCC encapsulation.) Two related versions of extended VLAN cross-connect are supported:
 - CCC version (**extended-vlan-ccc**)—Extended VLAN CCC encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. When you use this encapsulation type, you can configure the **ccc** family only.
 - TCC version (**extended-vlan-tcc**)—Similar to CCC, but used for circuits with different media on either side of the connection.

For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC and extended VLAN TCC are not supported.

- Ethernet VPLS (**ethernet-vpls**)—Ethernet interfaces with VPLS enabled can use Ethernet VPLS encapsulation. For more information about VPLS, see the *JUNOS VPNs Configuration Guide* and the *JUNOS Feature Guide*.
- Ethernet VLAN VPLS (**vlan-vpls**)—Ethernet interfaces with VLAN tagging and VPLS enabled can use Ethernet VLAN VPLS encapsulation. For more information about VPLS, see the *JUNOS VPNs Configuration Guide* and the *JUNOS Feature Guide*.
- Extended VLAN VPLS (**extended-vlan-vpls**)—Ethernet interfaces with VLAN 802.1Q tagging and VPLS enabled can use Ethernet Extended VLAN VPLS encapsulation. (Ethernet interfaces with standard TPID tagging can use Ethernet VLAN VPLS encapsulation.) Extended Ethernet VLAN VPLS encapsulation supports TPIDs 0x8100, 0x9100, and 0x9901. For more information about VPLS, see the *JUNOS VPNs Configuration Guide* and the *JUNOS Feature Guide*.
- Flexible Ethernet services (**flexible-ethernet-services**)—Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) can use flexible Ethernet services encapsulation. Aggregated Ethernet bundles can use this encapsulation type. You use this encapsulation type when you want to configure multiple per-unit Ethernet encapsulations. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs),

and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

- **Flexible Frame Relay (`flexible-frame-relay`)**—IQ and IQE interfaces can use flexible Frame Relay encapsulation. You use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any data-link connection identifier (DLCI) value from 1 through 1022.
- **Frame Relay (`frame-relay`)**—Defined in RFC 1490, *Multiprotocol Interconnect over Frame Relay*. E1, E3, link services, SONET/SDH, T1, T3, and voice services interfaces can use Frame Relay encapsulation. Five related versions are supported:
 - **CCC version (`frame-relay-ccc`)**—The same as standard Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to CCC. The logical interface must also have `frame-relay-ccc` encapsulation. When you use this encapsulation type, you can configure the `ccc` family only.
 - **TCC version (`frame-relay-tcc`)**—Similar to Frame Relay CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
 - **Extended CCC version (`extended-frame-relay-ccc`)**—This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC. The logical interface must have `frame-relay-ccc` encapsulation. When you use this encapsulation type, you can configure the `ccc` family only.
 - **Extended TCC version (`extended-frame-relay-tcc`)**—Similar to extended Frame Relay CCC, this encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC, which is used for circuits with different media on either side of the connection.
 - **Port CCC version (`frame-relay-port-ccc`)**—Defined in the IETF document *Frame Relay Encapsulation over Pseudo-Wires* (expired December 2002). This encapsulation type allows you to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. The connection between the two CE routers can be either user-to-network interface (UNI) or network-to-network interface (NNI); this is completely transparent to the PE routers. The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the `ccc` family only.
- **Frame Relay Ether Type (`frame-relay-ether-type`)**—Physical interfaces can use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. IETF frame relay encapsulation identifies the payload format using NLPID and SNAP formats. Cisco-compatible Frame Relay encapsulation uses the Ethernet type to identify the type of payload. Two related versions are supported:
 - **TCC version (`frame-relay-ether-type-tcc`)**—Cisco-compatible Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to TCC. This encapsulation is used for circuits with different media on either side of the connection.

- Extended TCC version (**extended-frame-relay-ether-type-tcc**)—This encapsulation allows you to dedicate Cisco-compatible Frame Relay TCC for DLCIs 1 through 1022. This encapsulation is used for circuits with different media on either side of the connection.
- Multilink Frame Relay (MLFR) UNI and NNI (**multilink-frame-relay-uni-nni**)—Link services and voice services interfaces functioning as FRF.16 bundles can use multilink Frame Relay UNI NNI encapsulation. This encapsulation is also used on link services and voice services interfaces' constituent T1, E1, or NxDS0 interfaces.
- PPP—Defined in RFC 1661, *The Point-to-Point Protocol (PPP) for the Transmission of Multiprotocol Datagrams over Point-to-Point Links*. PPP is the default encapsulation type for physical interfaces. E1, E3, SONET/SDH, T1, and T3 interfaces can use PPP encapsulation. Two related versions are supported:
 - Circuit cross-connect (CCC) version (**ppp-ccc**)—The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the **ccc** family only.
 - Translational cross-connect (TCC) version (**ppp-tcc**)—Similar to CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.



NOTE: When the encapsulation type is set to Cisco-compatible Frame Relay encapsulation, ensure that the LMI type is set to ANSI or Q933-A.

Encapsulation Capabilities

When you configure a point-to-point encapsulation (such as PPP or Cisco HDLC) on a physical interface, the physical interface can have only one logical interface (that is, only one **unit** statement) associated with it. When you configure a multipoint encapsulation (such as Frame Relay), the physical interface can have multiple logical units, and the units can be either point-to-point or multipoint.

Ethernet CCC encapsulation for Ethernet interfaces with standard TPID tagging requires that the physical interface have only a single logical interface. Ethernet interfaces in VLAN mode can have multiple logical interfaces.

For Ethernet interfaces in VLAN mode, VLAN IDs are applicable as follows:

- VLAN ID 0 is reserved for tagging the priority of frames.
- For encapsulation type **vlan-ccc**, VLAN IDs 1 through 511 are reserved for normal VLANs. VLAN IDs 512 and above are reserved for VLAN CCCs.
- For encapsulation type **vlan-vpls**, VLAN IDs 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for VPLS VLANs.
- For Gigabit Ethernet interfaces and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can configure flexible Ethernet services encapsulation

on the physical interface. For interfaces with **flexible-ethernet-services** encapsulation, all VLAN IDs are valid. VLAN IDs from 1 through 511 are not reserved.

- For encapsulation types **extended-vlan-ccc** and **extended-vlan-vpls**, all VLAN IDs are valid.

The upper limits for configurable VLAN IDs vary by interface type. For more information, see “Configuring 802.1Q VLANs” on page 599.

When you configure a TCC encapsulation, some modifications are needed to handle VPN connections over unlike Layer 2 and Layer 2.5 links and terminate the Layer 2 and Layer 2.5 protocol locally.

The router performs the following media-specific changes:

- PPP TCC—Both Link Control Protocol (LCP) and Network Control Protocol (NCP) are terminated on the router. Internet Protocol Control Protocol (IPCP) IP address negotiation is not supported. The JUNOS Software strips all PPP encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to PPP encapsulation.
- Cisco HDLC TCC—Keepalive processing is terminated on the router. The JUNOS Software strips all Cisco HDLC encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to Cisco HDLC encapsulation.
- Frame Relay TCC—All Local Management Interface (LMI) processing is terminated on the router. The JUNOS Software strips all Frame Relay encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to Frame Relay encapsulation.
- ATM—Operation, Administration, and Maintenance (OAM) and Interim Local Management Interface (ILMI) processing is terminated at the router. Cell relay is not supported. The JUNOS Software strips all ATM encapsulation data from incoming frames before forwarding them. For output, the next hop is changed to ATM encapsulation.

Example: Configuring the Encapsulation on a Physical Interface

Configure PPP encapsulation on a SONET/SDH interface. The second and third **family** statements allow Intermediate System-to-Intermediate System (IS-IS) and MPLS to run on the interface.

```
[edit interfaces]
so-7/0/0 {
  encapsulation ppp;
  unit 0 {
    point-to-point;
    family inet {
      address 192.168.1.113/32 {
        destination 192.168.1.114;
      }
    }
  }
  family iso;
```

```

        family mpls;
    }
}

```

Configuring the PPP Challenge Handshake Authentication Protocol

For interfaces with PPP encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP), as defined in RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*. When you enable CHAP on an interface, the interface can authenticate its peer and can be authenticated by its peer.

By default, PPP CHAP is disabled. If CHAP is not explicitly enabled, the interface makes no CHAP challenges and denies all incoming CHAP challenges. To enable CHAP, you must create an access profile, and you must configure the interfaces to use CHAP.

To configure a CHAP access profile, include the **profile** statement and specify a profile name at the **[edit access]** hierarchy level:

```

[edit access]
profile profile-name {
    client name chap-secret data;
}

```

For more information about configuring access profiles, see the *JUNOS System Basics Configuration Guide*.

When you configure an interface to use CHAP, you must assign an access profile to the interface. When an interface receives CHAP challenges and responses, the access profile in the packet is used to look up the shared secret, as defined in RFC 1994.

If no matching access profile is found for the CHAP challenge that was received by the interface, the optionally configured default CHAP secret is used. The default CHAP secret is useful if the CHAP name of the peer is unknown, or if the CHAP name changes during PPP link negotiation.

To configure PPP CHAP on an interface with PPP encapsulation, include the **chap** statement at the **[edit interfaces interface-name ppp-options]** hierarchy level:

```

[edit interfaces interface-name ppp-options]
chap {
    access-profile name;
    default-chap-secret name;
    local-name name;
    passive;
}

```

On each interface with PPP encapsulation, you can configure the following PPP CHAP properties:

- Assigning an Access Profile to an Interface on page 113
- Configuring a Default CHAP Secret on page 113

- Configuring the Local Name on page 113
- Configuring Passive Mode on page 114
- Example: Configuring the PPP Challenge Handshake Authentication Protocol on page 114

When you configure PPP over ATM or Multilink PPP over ATM encapsulation, you can enable CHAP on the logical interface. For more information, see “Configuring PPP over ATM2 Encapsulation” on page 334.

Assigning an Access Profile to an Interface

To assign an access profile to an interface, include the `access-profile` statement at the `[edit interfaces interface-name ppp-options chap]` hierarchy level:

```
[edit interfaces interface-name ppp-options chap]
access-profile name;
```

You must include the `access-profile` statement when you configure the CHAP authentication method. If an interface receives a CHAP challenge or response from a peer that is not in the applied access profile, the link is immediately dropped unless a default CHAP secret has been configured. For information about configuring the default CHAP secret, see “Configuring a Default CHAP Secret” on page 113.

Configuring a Default CHAP Secret

To configure a default CHAP secret for an interface, include the `default-chap-secret` statement at the `[edit interfaces interface-name ppp-options chap]` hierarchy level:

```
[edit interfaces interface-name ppp-options chap]
default-chap-secret name;
```

The default CHAP secret is used when no matching CHAP access profile exists, or if the CHAP name changes during PPP link negotiation.

Configuring the Local Name

By default, when CHAP is enabled on an interface, the interface uses the router’s system hostname as the name sent in CHAP challenge and response packets.

To configure the name the interface uses in CHAP challenge and response packets, include the `local-name` statement at the `[edit interfaces interface-name ppp-options chap]` hierarchy level:

```
[edit interfaces interface-name ppp-options chap]
local-name name;
```

The local name is any string from 1 to 250 characters in length, starting with an alphanumeric or underscore character, and including only the following characters:

```
a-z A-Z 0-9 % @ # / \ . _ -
```

Configuring Passive Mode

By default, when CHAP is enabled on an interface, the interface always challenges its peer and responds to challenges from its peer.

You can configure the interface not to challenge its peer, and only respond when challenged. To configure the interface not to challenge its peer, include the **passive** statement at the [edit interfaces *interface-name* ppp-options chap] hierarchy level:

```
[edit interfaces interface-name ppp-options chap]
passive;
```

Example: Configuring the PPP Challenge Handshake Authentication Protocol

Configure CHAP:

```
[edit access]
profile pe-A-ppp-clients;
client cpe-1 chap-secret "$1$dQYsZ$B5ojUeUjDsUo.yKwcCZ0";
    # SECRET-DATA
client cpe-2 chap-secret "$1$kdAsfaDAfkdjDsASxfafKdFKJ";
    # SECRET-DATA
[edit interfaces so-1/2/0]
encapsulation ppp;
ppp-options {
    chap {
        access-profile pe-A-ppp-clients;
        default-chap-secret "$9$mPafafhdsaiufhyrv1Rxd";
        local-name "pe-A-so-1/1/1";
    }
}
[edit interfaces so-1/1/2]
encapsulation ppp;
ppp-options {
    chap {
        access-profile pe-A-ppp-clients;
        default-chap-secret "$9$mPafafhdsaiufhyrv1Rxd";
        local-name "pe-A-so-1/1/2";
    }
}
```

Configuring the PPP Password Authentication Protocol

For interfaces with PPP encapsulation, you can configure interfaces to support the Password Authentication Protocol (PAP), as defined in RFC 1334, *PAP Authentication Protocols*. If authentication is configured, the PPP link negotiates using CHAP or PAP protocol for authentication during the Link Control Protocol (LCP) negotiation phase. PAP is only performed after the link establishment phase (LCP up) portion of the authentication phase.

During authentication, the PPP link sends a PAP authentication-request packet to the peer with an ID and password. The authentication-request packet is sent every

2 seconds, similar to the CHAP challenge, until a response is received (acknowledgment packet, nonacknowledgment packet). If an acknowledgment packet is received, the PPP link transitions to the next state, the network phase. If a nonacknowledgment packet is received, an LCP terminate request is sent, and the PPP link goes back to the link establishment phase. If no response is received, and an optional retry counter is set to **true**, a new request acknowledgment packet is resent. If the retry counter expires, the PPP link transitions to the LCP negotiate phrase.

You can configure the PPP link with PAP in passive mode. By default, when PAP is enabled on an interface, the interface expects authenticate-request packets from the peer. However, if a peer does not support bidirectional authentication, you can configure PAP to operate in passive mode. In PAP passive mode, the interface sends the authenticate-request packets to the peer only if the interface receives the PAP option from the peer—in passive mode, the interface does not authenticate the peer.

Both CHAP and PAP authentication can be configured on a PPP interface. If both are configured, CHAP is negotiated first. If CHAP authentication fails, PAP authentication is negotiated.

To enable PAP, you must create an access profile, and you must configure the interfaces to use PAP.

To configure a PAP access profile, include the **profile** statement and specify a profile name at the [edit access] hierarchy level:

```
[edit access]
profile profile-name {
  client name;
  pap-password password;
}
```

For more information about configuring access profiles, see the *JUNOS System Basics Configuration Guide*.

When you configure an interface to use PAP, you must assign an access profile to the interface. When an interface receives PAP authentication requests, the access profile in the packet is used to look up the password.

If no matching access profile is found for the PAP authentication request that was received by the interface, the optionally configured default PAP password is used. For information about configuring the default PAP password, see “Configuring PPP PAP Authentication” on page 164.

To configure PPP PAP on a physical interface with PPP encapsulation, include the **pap** statement at the [edit interfaces interface-name ppp-options] hierarchy level:

```
[edit interfaces interface-name ppp-options]
pap {
  access-profile name;
  local-name name;
  local-password password;
  passive;
}
```

To configure PPP PAP on a logical interface with PPP encapsulation, include the `pap` statement with options:

```
pap {
  default-pap-password password;
  local-name name;
  local-password password;
  passive;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For more information about configuring PAP for logical interfaces, see “Configuring PPP PAP Authentication” on page 164. For information about configuring tracing operations for PPP, see “Tracing Operations of the pppd Process” on page 119.

On each physical interface with PPP encapsulation, you can perform one of the following tasks:

- Configuring the Local Name on page 116
- Configuring the Local Password on page 116
- Configuring Passive Mode on page 117
- Example: Configuring PAP Authentication Protocol on page 117

Configuring the Local Name

By default, when PAP is enabled on an interface, the interface uses the router’s system hostname as the name sent in PAP request and response packets.

To configure the name the interface uses in PAP request and response packets, include the `local-name` statement at the [edit interfaces *interface-name* ppp-options `pap`] hierarchy level:

```
[edit interfaces interface-name ppp-options pap]
local-name name;
```

Configuring the Local Password

You need to configure the password to be used for authentication. To configure the host password for sending PAP requests, include the `local-password` statement at the [edit interfaces *interface-name* ppp-options `pap`] hierarchy level:

```
local-password password;
```

Configuring Passive Mode

By default, when PAP is enabled on an interface, the interface expects authenticate-request packets from the peer. However, if a peer does not support bidirectional authentication, you can configure PAP to operate in passive mode. In PAP passive mode, the interface sends the authenticate-request packets to the peer only if the interface receives the PAP option from the peer—in passive mode, the interface does not authenticate the peer.

To configure the interface to authenticate with PAP in passive mode, include the `passive` statement at the `[edit interfaces interface-name ppp-options pap]` hierarchy level:

```
[edit interfaces interface-name ppp-options pap]
passive;
```

Example: Configuring PAP Authentication Protocol

Configure a PAP access profile, the physical and logical interfaces, and tracing operations for PPP.

For PAP authentication, a username and password for the peer is configured in the access profile, along with a PAP password. Each user can have either a PAP password or a CHAP secret.

```
[edit access]
profile userlist1;
client {
  papuser {
    pap-password "%^***"; # SECRET-DATA;
  }
  chapuser {
    chap-secret "%^***"; # SECRET-DATA;
  }
}
```

To configure the same name for the PAP password and the CHAP secret, configure the client with two different access profiles:

```
[edit access]
profile chap-profile;
client {
  sjcrouter {
    chap-secret "%^***"; # SECRET-DATA;
  }
  boston {
    chap-secret "%^***"; # SECRET-DATA;
  }
}
profile pap-profile;
client {
  sjcrouter {
    pap-password "%^***"; # SECRET-DATA;
```

```

    }
    boston {
        pap-password "%@^***"; # SECRET-DATA;
    }
}

```

Configure the physical interface, including the access profile name to be used for PPP authentication:

```

[edit interfaces so-0/0/0]
ppp-options {
    pap {
        access-profile "pap-profile";
        local-name "rtrnum1";
        local-password "XXXXXXX"; #SECRET-DATA
        passive;
    }
}

```

Configure the logical interface, including the default PAP password to be used, should the access profile not be located during authentication:

```

[edit interfaces so-0/0/0]
encapsulation frame-relay;
unit 0 {
    dlci 100;
    encapsulation frame-relay-ppp;
    ppp-options {
        pap {
            local-name "rtrnum1";
            local-password "XXXXXXX"; #SECRET-DATA
            default-pap-password "XXXXX"; #SECRET-DATA
            passive;
        }
    }
}

```

Include the **pap** statement to trace PPP protocol operations:

```

[edit protocols]
ppp {
    traceoptions {
        flag {
            pap;
        }
    }
}

```

Monitoring a PPP Session

You can monitor PPP packet exchanges. When monitoring is enabled, packets exchanged during a session are logged by default to `/var/log/pppd`, or to the file specified in the **traceoptions** statement.

To configure PPP packet monitoring, include the `monitor-session` statement at the `[edit protocols ppp]` hierarchy level:

```
[edit protocols ppp]
monitor-session (interface-name | all);
```

When monitoring is configured, the operational mode commands `show ppp summary` and `show ppp interface` display a **Monitored** flag in the **Session flags** column or line.

Tracing Operations of the pppd Process

To trace the operations of the router's pppd process, include the `traceoptions` statement at the `[edit protocols ppp]` hierarchy level:

```
[edit protocols ppp]
traceoptions {
  file filename <files number> <match regular-expression> <size size> <world-readable |
    no-world-readable>;
  flag flag;
  level severity-level;
  no-remote-trace;
}
```

To specify more than one tracing operation, include multiple `flag` statements.

You can specify the following flags in the `traceoptions` statement:

- `access`—Access code
- `address-pool`—Address pool code
- `all`—All areas of code
- `auth`—Authentication code
- `chap`—Challenge Handshake Authentication Protocol (CHAP) code
- `config`—Configuration code
- `ifdb`—Interface database code
- `lcp`—LCP state machine code
- `memory`—Memory management code
- `message`—Message processing code
- `ncp`—NCP state machine code
- `pap`—Password Authentication Protocol (PAP) code
- `ppp`—PPP protocol processing code
- `radius`—RADIUS processing code
- `rtsock`—Routing socket code
- `session`—Session management code
- `signal`—Signal handling code

- **timer**—Timer code
- **ui**—User interface code

For general information about tracing, see the tracing and logging information in the *JUNOS System Basics Configuration Guide*.

Configuring PPP Address and Control Field Compression

For interfaces with PPP, PPP CCC, or PPP TCC encapsulation, you can configure compression of the Data Link Layer address and control fields, as defined in RFC 1661, *The Point-to-Point Protocol (PPP)*. By default, the address and control fields are not compressed. This means PPP-encapsulated packets are transmitted with two 1-byte fields (0xff and 0x03). If you configure address and control field compression (ACFC) and ACFC is successfully negotiated with the local router's peer, the local router transmits packets without these 2 bytes. ACFC allows you to conserve bandwidth by transmitting less data.

On M320, M120, and T Series routers, ACFC is not supported for any ISO family protocols. Do not include the **acfc** statement at the [edit interfaces *interface-name* **ppp-options** **compression**] hierarchy level when you include the **family iso** statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level.



NOTE: The address and control fields cannot be compressed in Link Control Protocol (LCP) packets.

The PPP session restarts when you configure or modify compression options.

To configure ACFC, include the **compression** statement at the [edit interfaces *interface-name* **ppp-options**] hierarchy level, and specify **acfc**:

```
[edit interfaces interface-name ppp-options]
compression acfc;
```

This configuration causes the local router to try to negotiate ACFC with its peer. If ACFC is successfully negotiated, the local router sends packets with compressed address and control fields. When you include the **compression acfc** statement in the configuration, the PPP session restarts, and the local router sends the ACFC option in the LCP Configure-Request packet. The ACFC option informs the local router's peer that the local router can receive packets with compression. If the peer indicates that it, too, can receive packets with compression, then ACFC is negotiated. If ACFC is successfully negotiated, the local router can receive packets with or without the address and control bytes included.

To monitor the configuration, issue the **show interfaces *interface-name*** command. Configured options are displayed in the **link flags** field for the physical interface. Successfully negotiated options are displayed in the **flags** field for the logical interface. In this example, both ACFC and PFC are configured, but neither compression feature has been successfully negotiated.


```

user@router# run show interfaces so-0/1/1
Physical interface: so-0/1/1, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 27
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
  Loopback: None, FCS: 16
    Payload scrambler: Enabled
    Device flags   : Present Running
    Interface flags: Point-To-Point SNMP-Traps 16384
    Link flags     : No-Keepalives ACFC PFC
    LCP state: Opened
    NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
    CHAP state: Not-configured
    CoS queues   : 4 supported
    Last flapped : 2004-12-29 10:49:32 PST (00:18:35 ago)
    Input rate    : 0 bps (0 pps)
    Output rate   : 0 bps (0 pps)
    SONET alarms  : None
    SONET defects : None
  Logical interface so-0/1/1.0 (Index 68) (SNMP ifIndex 169)
    Flags: Point-To-Point SNMP-Traps ACFC Encapsulation: PPP
    Protocol inet, MTU: 4470
      Flags: None
      Addresses, Flags: Is-Preferred Is-Primary
      Destination: 3.3.3/24, Local: 3.3.3.2, Broadcast: 3.3.3.255

```

Configuring the PPP Protocol Field Compression

For interfaces with PPP, PPP CCC, or PPP TCC encapsulation, you can configure protocol field compression. By default, the protocol field is not compressed. This means PPP-encapsulated packets are transmitted with a two-byte protocol field. For example, IPv4 packets are transmitted with the protocol field set to 0x0021, and MPLS packets are transmitted with the protocol field set to 0x0281.

For all protocols with identifiers in the range 0x0000 through 0x00ff, you can configure the router to compress the protocol field to one byte, as defined in RFC 1661, *The Point-to-Point Protocol (PPP)*. Protocol field compression (PFC) allows you to conserve bandwidth by transmitting less data.



NOTE: The protocol field cannot be compressed in Link Control Protocol (LCP) packets.

The PPP session restarts when you configure or modify compression options.

To configure PFC, include the **compression** statement at the [edit interfaces *interface-name* ppp-options] hierarchy level, and specify **pfc**:

```

[edit interfaces interface-name ppp-options]
  compression pfc;

```

This configuration causes the local router to try to negotiate PFC with its peer. If PFC is successfully negotiated, the local router sends packets with compressed protocol fields. When you include the **compression pfc** statement in the configuration, the PPP session restarts, and the local router sends the PFC option in the LCP

Configure-Request packet. The PFC option informs the local router's peer that the local router can receive packets with compression. If the peer indicates that it, too, can receive packets with compression, then PFC is negotiated. If PFC is successfully negotiated, the local router can receive packets with either 2-byte (uncompressed) or 1-byte (compressed) protocol fields.

To monitor the configuration, issue the `show interfaces interface-name` command. Configured options are displayed in the `link flags` field for the physical interface. Successfully negotiated options are displayed in the `flags` field for the logical interface. In this example, both ACFC and PFC are configured, but neither compression feature has been successfully negotiated.

```
user@router# run show interfaces so-0/1/1
Physical interface: so-0/1/1, Enabled, Physical link is Up
  Interface index: 133, SNMP ifIndex: 27
  Link-level type: PPP, MTU: 4474, Clocking: Internal, SONET mode, Speed: OC3,
Loopback: None, FCS: 16,
  Payload scrambler: Enabled
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps 16384
  Link flags     : No-Keepalives ACFC PFC
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
  CHAP state: Not-configured
  CoS queues   : 4 supported
  Last flapped : 2004-12-29 10:49:32 PST (00:18:35 ago)
  Input rate    : 0 bps (0 pps)
  Output rate   : 0 bps (0 pps)
  SONET alarms  : None
  SONET defects : None
Logical interface so-0/1/1.0 (Index 68) (SNMP ifIndex 169)
  Flags: Point-To-Point SNMP-Traps ACFC Encapsulation: PPP
  Protocol inet, MTU: 4470
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 3.3.3/24, Local: 3.3.3.2, Broadcast: 3.3.3.255
```

Configuring the Interface Speed

You can configure the interface speed on the following interfaces:

- Management Ethernet Interface on M Series and T Series routers on page 122
- Gigabit Ethernet Interfaces on J Series Routers on page 123
- Fast Ethernet Interface on page 123
- Tri-Rate Ethernet Copper Interface on page 124
- SONET/SDH Interface on page 124

Management Ethernet Interface on M Series and T Series routers

By default, the M Series and T Series routers management Ethernet interface autonegotiates whether to operate at 10 megabits per second (Mbps) or 100 Mbps. All other interfaces automatically choose the correct speed based on the PIC type

and whether the PIC is configured to operate in multiplexed mode (using the `no-concatenate` statement in the `[edit chassis]` configuration hierarchy, as described in the *JUNOS System Basics Configuration Guide*).



NOTE: For M Series, MX Series, and most T Series routers, the management Ethernet interface is `fxp0`. For TX Matrix Plus routers and T1600 routers configured in a routing matrix, the management Ethernet interface is `em0`.



NOTE: Automated scripts that you have developed for standalone T1600 routers (T1600 routers that are not in a routing matrix) might contain references to the `fxp0` management Ethernet interface. Before reusing the scripts on T1600 routers in a routing matrix, edit the command lines that reference the `fxp0` management Ethernet interface so that the commands reference the `em0` management Ethernet interface instead.

To configure the management Ethernet interface to operate at 10 Mbps or 100 Mbps, include the `speed` statement at the `[edit interfaces fxp0]` or `[edit interfaces em0]` hierarchy level:

```
[edit interfaces (fxp0 | em0)]
  speed (10m | 100m);
```

For information about configuring the link mode, see “Configuring the Link Characteristics on Ethernet Interfaces” on page 595.



NOTE: The `fxp0` interface does not support CoS.

Gigabit Ethernet Interfaces on J Series Routers

By default, Gigabit Ethernet interfaces (both built-in and PIMs) for J Series routers autonegotiate whether to operate at 10 megabits per second (Mbps), 100 Mbps, or 1000 Mbps.

To configure a J Series Gigabit Ethernet interface to operate at 10 Mbps, 100 Mbps, or 1000 Mbps, include the `speed` statement at the `[edit interfaces ge-pim/O/port]` hierarchy level:

```
[edit interfaces ge-pim/O/port]
  speed (10m | 100m | 1g);
```

For information about configuring the link mode, see “Configuring the Link Characteristics on Ethernet Interfaces” on page 595.

Fast Ethernet Interface

By default, both of the built-in Fast Ethernet ports on the M7i router FIC autonegotiate whether to operate at 10 Mbps or 100 Mbps. All other interfaces automatically choose

the correct speed based on the PIC type and whether the PIC is configured to operate in multiplexed mode (using the `no-concatenate` statement at the `[edit chassis]` hierarchy level, as described in the *JUNOS System Basics Configuration Guide*).

If the link partner does not support autonegotiation, configure either Fast Ethernet port manually to match its link partner's speed and link mode. When the link mode is configured, autonegotiation is disabled.



NOTE: When you manually configure Fast Ethernet interfaces on the M Series and T Series routers, link mode and speed must both be configured. If both these values are not configured, the router uses autonegotiation for the link and ignores the user-configured settings.

To configure a Fast Ethernet port on the FIC to operate at 10 Mbps or 100 Mbps, include the `speed` statement at the `[edit interfaces fe-fpc/pic/port]` hierarchy level:

```
[edit interfaces fe-fpc/pic/port]
speed (10m | 100m);
```

For information about configuring the link mode, see “Configuring the Link Characteristics on Ethernet Interfaces” on page 595.

Tri-Rate Ethernet Copper Interface

By default, the Tri-Rate Ethernet copper interfaces on MX Series routers operate at 1 Gbps. Tri-Rate Ethernet copper interfaces can also be configured to operate at 10 Mbps, 100 Mbps, or 1 Gbps.



NOTE: When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.



NOTE: Half-duplex mode is not supported on Tri-Rate Ethernet copper interfaces.

To configure a Tri-Rate Ethernet copper interface to operate at 10 Mbps, 100 Mbps, or 1 Gbps, include the `speed` statement at the `[edit interfaces ge-fpc/pic/port]` hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
speed (10m | 100m | 1g);
```

For information about configuring the link mode, see “Configuring the Link Characteristics on Ethernet Interfaces” on page 595.

SONET/SDH Interface

You can configure the speed of SONET/SDH interfaces on next-generation SONET/SDH Type 1 and Type 2 PICs with SFP. The speed you select is dependent upon whether

the PIC is in concatenated or nonconcatenated mode. In concatenated mode, the bandwidth of the interface is in a single channel. In nonconcatenated mode, the PIC operates in channelized (multiplexed) mode.

Table 19 on page 125 shows the mode combinations for the next-generation SONET/SDH Type 1 PICs with SFP.

Table 19: Type 1 PIC Mode Combinations

PIC	Mode	Speed Configuration	Default Mode
2-port OC3	2xOC3 concatenated	<i>fpc/pic/0 speed oc3</i>	concatenated
4-port OC3	1xOC12 concatenated	<i>fpc/pic/0 speed oc12</i>	—
	1xOC12 nonconcatenated	<i>fpc/pic/0:0 speed oc3</i>	nonconcatenated
	4xOC3 concatenated	<i>fpc/pic/port speed oc3</i>	concatenated
1-port OC12	1xOC12 concatenated	<i>fpc/pic/0 speed oc12</i>	concatenated
	1xOC12 nonconcatenated	<i>fpc/pic/0:0 speed oc3</i>	nonconcatenated
	1xOC3 concatenated	<i>fpc/pic/0 speed oc3</i>	—

Table 20 on page 125 shows the mode combinations for the next-generation SONET/SDH Type 2 PICs with SFP.

Table 20: Type 2 PIC Mode Combinations

PIC	Mode	Speed Configuration	Default Mode
1-port OC48, IQ and IQE	1xOC48 concatenated	<i>fpc/pic/0 speed oc48</i>	concatenated
	1xOC48 nonconcatenated	<i>fpc/pic/0:0 speed oc12</i>	nonconcatenated
	1xOC12 concatenated	<i>fpc/pic/0 speed oc12</i>	—
	1xOC12 nonconcatenated	<i>fpc/pic/0 0 speed oc3</i>	—
	1xOC3 concatenated	<i>fpc/pic/0 speed oc3</i>	—

Table 20: Type 2 PIC Mode Combinations (*continued*)

PIC	Mode	Speed Configuration	Default Mode
4-port OC12, IQ and IQE	1xOC48 concatenated	<i>fpc/pic/0 speed oc48</i>	—
	1xOC48 nonconcatenated	<i>fpc/pic/0:0 speed oc12</i>	nonconcatenated
	1xOC12 nonconcatenated	<i>fpc/pic/0 speed oc3</i>	—
	4xOC12 concatenated	<i>fpc/pic/port speed oc3 oc12</i>	concatenated
4-port OC3, IQ and IQE	1xOC12 concatenated	<i>fpc/pic/0 speed oc12</i>	—
	1xOC12 nonconcatenated	<i>fpc/pic/0:0 speed oc3</i>	nonconcatenated
	4xOC3 concatenated	<i>fpc/pic/port speed oc3</i>	concatenated

By default, SONET/SDH PICs operate in concatenated mode. To specify interface speed in concatenated mode, include the **speed** statement with options at the [edit interfaces *so-fpc/pic/port*] hierarchy level:

```
[edit interfaces so-fpc/pic/port]
speed (oc3 | oc12 | oc48);
```

For example, each port of the 4-port OC12 PIC can be configured to be in OC3 or OC12 speed independently when this PIC is in 4xOC12 concatenated mode.

To specify interface speed in nonconcatenated mode, include the **speed** statement at the [edit interfaces *so-fpc/pic/port.channel*] hierarchy level:

```
[edit interfaces so-fpc/pic/port.channel]
speed (oc3 | oc12);
```

To configure the PIC to operate in channelized (multiplexed) mode, include the **no-concatenate** statement at the [edit chassis *fpc slot-number pic pic-number*] hierarchy level.

For more information about using the **no-concatenate** statement, see the *JUNOS System Basics Configuration Guide*.

Configuring Keepalives

By default, physical interfaces configured with Cisco HDLC or PPP encapsulation send keepalive packets at 10-second intervals. The Frame Relay term for keepalives is LMI packets; the JUNOS Software supports both ANSI T1.617 Annex D LMIs and ITU Q933 Annex A LMIs. On ATM networks, OAM cells perform the same function. You configure OAM cells at the logical interface level; for more information, see “Defining the ATM OAM F5 Loopback Cell Period” on page 329.

To disable the sending of keepalives on a physical interface, include the **no-keepalives** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]  
no-keepalives;
```

To disable the sending of keepalives on a physical interface configured with Cisco HDLC encapsulation for a translational cross-connection, include the **no-keepalives** statement at the [edit interfaces *interface-name* encapsulation cisco-hdlc-tcc] hierarchy level:

```
[edit interfaces interface-name]  
encapsulation cisco-hdlc-tcc {  
    no-keepalives;  
}
```

For more information about translation cross-connections, see “Configuring Circuit and Translational Cross-Connects” on page 223.

When you configure PPP over ATM or Multilink PPP over ATM encapsulation, you can enable or disable keepalives on the logical interface. For more information, see “Configuring PPP over ATM2 Encapsulation” on page 334.

To explicitly enable the sending of keepalives on a physical interface, include the **keepalives** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]  
keepalives;
```

To change one or more of the default keepalive values, include the appropriate option at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]  
keepalives <interval seconds> <down-count number> <up-count number>;
```

On interfaces configured with Cisco HDLC or PPP encapsulation, you can include the following three keepalive statements; note that Frame Relay encapsulation is not affected by these statements:

- **interval seconds**—The time in seconds between successive keepalive requests. The range is from 1 second through 32767 seconds, with a default of 10 seconds.
- **down-count number**—The number of keepalive packets a destination must fail to receive before the network takes a link down. The range is from 1 through 255, with a default of 3.
- **up-count number**—The number of keepalive packets a destination must receive to change a link’s status from down to up. The range is from 1 through 255, with a default of 1.



WARNING: If interface keepalives are configured on an interface that does not support the **keepalives** configuration statement (for example, 10-Gigabit Ethernet), the link

layer may go down when the PIC is restarted. Avoid configuring the `keepalives` on interfaces that do not support the `keepalives` configuration statement.

For information about Frame Relay keepalive settings, see “Configuring Frame Relay Keepalives” on page 378.

Configuring the Clock Source

For both the router and interfaces, the clock source can be the router’s internal Stratum 3 clock, which resides on the System Control Board (SCB), the System and Switch Board (SSB), the Forwarding Engine Board (FEB), or the Miscellaneous Control Subsystem (MCS) (depending on the router model), or an external clock that is received on the interface. By default, the 19.44-MHz Stratum 3 reference clock generates the clock signal for all serial PICs (SONET/SDH) and Plesiochronous Digital Hierarchy (PDH) PICs. PDH PICs include DS3, E3, T1, and E1 PICs.

For example, interface A can transmit on interface A’s received clock (external, loop timing) or the Stratum 3 clock (internal, line timing or normal timing). Interface A cannot use a clock from any other source. For interfaces such as SONET/SDH that can use different clock sources, you can configure the source of the transmit clock on each interface.

To set the clock source, include the `clocking` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]
  clocking (external | internal);
```

For information about clocking on channelized interfaces, see “Clock Sources on Channelized Interfaces” on page 390. Also see “Configuring Channelized IQ and IQE SONET/SDH Loop Timing” on page 852 and “Configuring the Channelized T3 Loop Timing” on page 576. For information about configuring an external synchronization interface that can be used to synchronize the internal Stratum 3 clock to an external source on the M320 and M120 routers, see “Configuring an External Synchronization Interface” on page 65.

Configuring the Router as a DCE

By default, when you configure an interface with Frame Relay encapsulation, the router is assumed to be data terminal equipment (DTE). That is, the router is assumed to be at a terminal point on the network. To configure the router to be data circuit-terminating equipment (DCE), include the `dce` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]
  dce;
```

When you configure the router to be a DCE, keepalives are disabled by default.

For back-to-back Frame Relay connections, either disable the sending of keepalives on both sides of the connection, or configure one side of the connection as a DTE (the default JUNOS configuration) and the other as a DCE.

Configuring Receive and Transmit Leaky Bucket Properties

Congestion control is particularly difficult in high-speed networks with high volumes of traffic. When congestion occurs in such a network, it is usually too late to react. You can avoid congestion by regulating the flow of packets into your network. Smoother flows prevent bursts of packets from arriving at (or being transmitted from) the same interface and causing congestion.

For all interface types except ATM, channelized E1, E1, Fast Ethernet, Gigabit Ethernet, and channelized IQ, you can configure leaky bucket properties, which allow you to limit the amount of traffic received on and transmitted by a particular interface. You effectively specify what percentage of the interface's total capacity can be used to receive or transmit packets. You might want to set leaky bucket properties to limit the traffic flow from a link that is known to transmit high volumes of traffic.



NOTE: Instead of configuring leaky bucket properties, you can limit traffic flow by configuring policers. Policers work on all interfaces. For more information, see “Applying Policers” on page 194 and the *JUNOS Policy Framework Configuration Guide*.

The leaky bucket is used at the host-network interface to allow packets into the network at a constant rate. Packets might be generated in a bursty manner, but after they pass through the leaky bucket, they enter the network evenly spaced. In some cases, you might want to allow short bursts of packets to enter the network without smoothing them out. By controlling the number of packets that can accumulate in the bucket, the **threshold** property controls burstiness. The maximum number of packets entering the network in t time units is $\text{threshold} + \text{rate} * t$.

By default, leaky buckets are disabled, and the interface can receive and transmit packets at the maximum line rate.

To configure leaky bucket properties, include one or both of the **receive-bucket** and **transmit-bucket** statements at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
receive-bucket {
    overflow (discard | tag);
    rate percentage;
    threshold bytes;
}
transmit-bucket {
    overflow discard;
    rate percentage;
    threshold bytes;
}
```

In the **rate** statement, specify the percentage of the interface line rate that is available to receive or transmit packets. The percentage can be a value from 0 (none of the

interface line rate is available) to 100 (the maximum interface line rate is available). For example, when you set the line rate to 33, the interface receives or transmits at one-third of the maximum line rate.

In the **threshold** statement, specify the bucket threshold, which controls the burstiness of the leaky bucket mechanism. The larger the value, the more bursty the traffic, which means that over a very short time the interface can receive or transmit close to line rate, but the average over a longer time is at the configured bucket rate. The threshold can be a value from 0 through 65,535 bytes. For ease of entry, you can enter *number* either as a complete decimal number or as a decimal number followed by the abbreviation k (1,000). For example, the entry **threshold 2k** corresponds to a threshold of 2,000 bytes.

In the **overflow** statement, specify how to handle packets that exceed the threshold:

- **tag** (receive bucket only)—Tag, count, and process received packets that exceed the threshold.
- **discard**—Discard received packets that exceed the threshold. No counting is done.

Configuring Accounting for the Physical Interface

Juniper Networks routers can collect various kinds of data about traffic passing through the router. You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- The fields used in the accounting records
- The number of files that the router retains before discarding, and the number of bytes per file
- The polling period that the system uses to record the data

You configure the profiles and define a unique name for each profile using statements at the [edit **accounting-options**] hierarchy level. There are two types of accounting profiles: interface profiles and filter profiles. You configure interface profiles by including the **interface-profile** statement at the [edit **accounting-options**] hierarchy level. You configure filter profiles by including the **filter-profile** statement at the [edit **accounting-options**] hierarchy level. For more information, see the *JUNOS Network Management Configuration Guide*.

You apply filter profiles by including the **accounting-profile** statement at the [edit **firewall filter filter-name**] and [edit **firewall family family filter filter-name**] hierarchy levels. For more information, see the *JUNOS Policy Framework Configuration Guide*.

Applying an Accounting Profile to the Physical Interface

To enable accounting on an interface, include the **accounting-profile** statement at the [edit **interfaces interface-name**] hierarchy level:

```
[edit interfaces interface-name]
  accounting-profile name;
```

You can also reference profiles by logical unit; for more information, see “Configuring Accounting for the Logical Interface” on page 157.

Example: Applying an Accounting Profile to the Physical Interface

Configure an accounting profile for an interface and apply it to a physical interface:

```
[edit]
accounting-options {
  file if_stats {
    size 4m files 10 transfer-interval 15;
    archive-sites {
      "ftp://login:password@host/path";
    }
  }
  interface-profile if_profile {
    interval 15;
    file if_stats {
      fields {
        input-bytes;
        output-bytes;
        input-packets;
        output-packets;
        input-errors;
        output-errors;
      }
    }
  }
}
[edit interfaces ge-1/0/1]
accounting-profile if_profile;
```

Interface Diagnostics

You can use two diagnostic tools to test the physical layer connections of interfaces: Loopback testing and bit error rate test (BERT) testing. Loopback testing enables you to verify the connectivity of a circuit. BERT testing enables you to identify poor signal quality on a circuit. This section contains the following topics:

- Configuring Loopback Testing on page 131
- Interface Diagnostics on page 134

Configuring Loopback Testing

Loopback testing allows you to verify the connectivity of a circuit. You can configure any of the following interfaces to execute a loopback test: Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, E1, E3, NxDS0, serial, SONET/SDH, T1, and T3.

The physical path of a network data circuit usually consists of segments interconnected by devices that repeat and regenerate the transmission signal. The transmit path on one device connects to the receive path on the next device. If a circuit fault occurs in the form of a line break or a signal corruption, you can isolate the problem by

using a loopback test. Loopback tests allow you to isolate segments of the circuit and test them separately.

To do this, configure a *line loopback* on one of the routers. Instead of transmitting the signal toward the far-end device, the line loopback sends the signal back to the originating router. If the originating router receives back its own data link layer packets, you have verified that the problem is beyond the originating router. Next, configure a line loopback farther away from the local router. If this originating router does not receive its own data link layer packets, you can assume the problem is on one of the segments between the local router and the remote router's interface card. In this case, the next troubleshooting step is to configure a line loopback closer to the local router to find the source of the problem.

There are several types of loopback testing supported by the JUNOS Software, as follows:

- DCE local—Loops packets back on the local DCE.
- DCE remote—Loops packets back on the remote DCE.
- Local—Useful for troubleshooting physical PIC errors. A local loopback loops packets, including both data and timing information, back on the local router's PIC. When you configure a local loopback, the interface transmits packets to the channel services unit (CSU) built into the interface. These packets are transmitted onto the circuit toward the far-end device. The PIC receives back its own transmission and ignores any data sent from the physical circuit and the CSU. To test a local loopback, issue the **show interfaces *interface-name*** command. If PPP keepalives transmitted on the interface are received by the PIC, the **Device Flags** field contains the output **Loop-Detected**.
- Payload—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A payload loopback loops data only (without clocking information) on the remote router's PIC. With payload loopback, overhead is recalculated.
- Remote—Useful for troubleshooting the physical circuit problems between the local router and the remote router. A remote loopback loops packets, including both data and timing information, back on the remote router's interface card. A router at one end of the circuit initiates a remote loopback toward its remote partner. When you configure a remote loopback, the packets received from the physical circuit and CSU are received by the interface. Those packets are then retransmitted by the PIC back toward the CSU and the circuit. This loopback tests all the intermediate transmission segments.

Table 21 on page 132 shows the loopback modes supported on the various interface types.

Table 21: Loopback Modes by Interface Type

Interface	Loopback Modes	Usage Guidelines
Aggregated Ethernet, Fast Ethernet, Gigabit Ethernet	Local	"Configuring Ethernet Loopback Capability" on page 593

Table 21: Loopback Modes by Interface Type (*continued*)

Interface	Loopback Modes	Usage Guidelines
Circuit Emulation E1	Local and remote	“Configuring E1 Loopback Capability” on page 547
Circuit Emulation T1	Local and remote	“Configuring T1 Loopback Capability” on page 565
E1 and E3	Local and remote	“Configuring E1 Loopback Capability” on page 547 and “Configuring E3 Loopback Capability” on page 555
NxDS0	Payload	“Configuring NxDS0 IQ and IQE Interfaces” on page 502, “Configuring T1 and NxDS0 Interfaces” on page 458, “Configuring NxDS0 Interfaces” on page 431, “Configuring an NxDS0 IQ Interface” on page 468, and “Configuring an NxDS0 IQ Interface” on page 481
Serial (V.35 and X.21)	Local and remote	“Configuring Serial Loopback Capability” on page 275
Serial (EIA-530)	DCE local, DCE remote, local, and remote	“Configuring Serial Loopback Capability” on page 275
SONET/SDH	Local and remote	“Configuring Channelized IQ and IQE SONET/SDH Loop Timing” on page 852
T1 and T3	Local, payload, and remote	“Configuring T1 Loopback Capability” on page 565 and “Configuring T3 Loopback Capability” on page 576 See also “Configuring the T1 Remote Loopback Response” on page 564

To configure loopback testing, include the **loopback** statement:

```
loopback mode;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* ds0-options]
- [edit interfaces *interface-name* e1-options]
- [edit interfaces *interface-name* e3-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]
- [edit interfaces *interface-name* serial-options]
- [edit interfaces *interface-name* sonet-options]

- [edit interfaces *interface-name* t1-options]
- [edit interfaces *interface-name* t3-options]

Interface Diagnostics

BERT allows you to troubleshoot problems by checking the quality of links. You can configure any of the following interfaces to execute a BERT when the interface receives a request to run this test: E1, E3, T1, T3; the channelized DS3, OC3, OC12, and STM1 interfaces; and the channelized DS3 IQ, E1 IQ, and OC12 IQ interfaces.

A BERT test requires a line loop to be in place on either the transmission devices or the far-end router. The local router generates a known bit pattern and sends it out the transmit path. The received pattern is then verified against the sent pattern. The higher the bit error rate of the received pattern, the worse the noise is on the physical circuit. As you move the position of the line loop increasingly downstream toward the far-end router, you can isolate the troubled portion of the link.

To configure BERT, you must configure the duration of the test, the bit pattern to send on the transmit path, and the error rate to monitor when the inbound pattern is received.

To configure the duration of the test, the pattern to send in the bit stream, and the error rate to include in the bit stream, include the **bert-period**, **bert-algorithm**, and **bert-error-rate** statements, respectively, at the [edit interfaces *interface-name interface-type-options*] hierarchy level:

```
[edit interfaces interface-name interface-type-options]
bert-algorithm algorithm;
bert-error-rate rate;
bert-period seconds;
```

By default, the BERT period is 10 seconds. You can configure the BERT period to last from 1 through 239 seconds on some PICs and from 1 through 240 seconds on other PICs.

rate is the bit error rate. This can be an integer from 0 through 7, which corresponds to a bit error rate from 10^{-0} (1 error per bit) to 10^{-7} (1 error per 10 million bits).

algorithm is the pattern to send in the bit stream. For a list of supported algorithms, enter a ? after the **bert-algorithm** statement; for example:

```
[edit interfaces t1-0/0/0 t1-options]
user@host# set bert-algorithm ?
Possible completions:
pseudo-2e11-o152    Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151    Pattern is 2^15 - 1 (per 0.152 standard)
pseudo-2e20-o151    Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153    Pattern is 2^20 - 1 (per 0.153 standard)
...
```

For specific hierarchy information, see the individual interface types.



NOTE: The 4-port E1 PIC supports only the following algorithms:

pseudo-2e11-o152	Pattern is $2^{11} - 1$ (per 0.152 standard)
pseudo-2e15-o151	Pattern is $2^{15} - 1$ (per 0.151 standard)
pseudo-2e20-o151	Pattern is $2^{20} - 1$ (per 0.151 standard)
pseudo-2e23-o151	Pattern is 2^{23} (per 0.151 standard)

When you issue the **help** command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



NOTE: The 12-port T1/E1 Circuit Emulation (CE) PIC supports only the following algorithms:

all-ones-repeating	Repeating one bits
all-zeros-repeating	Repeating zero bits
alternating-double-ones-zeros	Alternating pairs of ones and zeros
alternating-ones-zeros	Alternating ones and zeros
pseudo-2e11-o152	Pattern is $2^{11} - 1$ (per 0.152 standard)
pseudo-2e15-o151	Pattern is $2^{15} - 1$ (per 0.151 standard)
pseudo-2e20-o151	Pattern is $2^{20} - 1$ (per 0.151 standard)
pseudo-2e7	Pattern is $2^7 - 1$
pseudo-2e9-o153	Pattern is $2^9 - 1$ (per 0.153 standard)
repeating-1-in-4	1 bit in 4 is set
repeating-1-in-8	1 bit in 8 is set
repeating-3-in-24	3 bits in 24 are set

When you issue the **help** command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.



NOTE: The IQE PICs support only the following algorithms:

all-ones-repeating	Repeating one bits
all-zeros-repeating	Repeating zero bits
alternating-double-ones-zeros	Alternating pairs of ones and zeros
alternating-ones-zeros	Alternating ones and zeros
pseudo-2e9-o153	Pattern is $2^9 - 1$ (per 0.153 (511 type) standard)
pseudo-2e11-o152	Pattern is $2^{11} - 1$ (per 0.152 and 0.153 (2047 type) standards)
pseudo-2e15-o151	Pattern is $2^{15} - 1$ (per 0.151 standard)
pseudo-2e20-o151	Pattern is $2^{20} - 1$ (per 0.151 standard)
pseudo-2e20-o153	Pattern is $2^{20} - 1$ (per 0.153 standard)
pseudo-2e23-o151	Pattern is $2^{23} - 1$ (per 0.151 standard)
repeating-1-in-4	1 bit in 4 is set
repeating-1-in-8	1 bit in 8 is set
repeating-3-in-24	3 bits in 24 are set

When you issue the **help** command from the CLI, all BERT algorithm options are displayed, regardless of the PIC type, and no commit check is available. Unsupported patterns for a PIC type can be viewed in system log messages.

Table 22 on page 136 shows the BERT capabilities for various interface types.

Table 22: BERT Capabilities by Interface Type

Interface	T1 BERT	T3 BERT	Comments
12-port T1/E1 Circuit Emulation	Yes (ports 0–11)		■ Limited algorithms
4-port Channelized OC3/STM1 Circuit Emulation	Yes (port 0–3)		■ Limited algorithms
E1 or T1	Yes (port 0–3)	Yes (port 0–3)	■ Single port at a time ■ Limited algorithms
E3 or T3	Yes (port 0–3)	Yes (port 0–3)	■ Single port at a time
Channelized OC12	N/A	Yes (channel 0–11)	■ Single channel at a time ■ Limited algorithms ■ No bit count
Channelized STM1	Yes (channel 0–62)	N/A	■ Multiple channels ■ Only one algorithm ■ No error insert ■ No bit count
Channelized T3 and Multichannel T3	Yes (channel 0–27)	Yes (port 0–3 on channel 0)	■ Multiple ports and channels ■ Limited algorithms for T1 ■ No error insert for T1 ■ No bit count for T1

These limitations do not apply to channelized IQ interfaces. For information about BERT capabilities on channelized IQ interfaces, see “Channelized IQ and IQE Interfaces Properties” on page 393.

Starting and Stopping a BERT Test

Before you can start the BERT test, you must disable the interface. To do this, include the `disable` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
disable;
```

After you configure the BERT properties and commit the configuration, begin the test by issuing the `test interface interface-name interface-type bert start` operational mode command:

```
user@host> test interface interface-name interface-type bert start
```


The test runs for the duration you specify with the **bert-period** statement. If you wish to terminate the test sooner, issue the **test interface *interface-name* *interface-type*-bert-stop** command:

```
user@host> test interface interface-name interface-type-bert-stop
```

For example:

```
user@host> test interface t3-1/2/0 t3-bert-start
user@host> test interface t3-1/2/0 t3-bert-stop
```

To view the results of the BERT test, issue the **show interfaces extensive | find BERT** command:

```
user@host> show interfaces interface-name extensive | find BERT
```

For more information about running and evaluating the results of the BERT procedure, see the *JUNOS System Basics and Services Command Reference*.



NOTE: To exchange BERT patterns between a local router and a remote router, include the **loopback remote** statement in the interface configuration at the remote end of the link. From the local router, issue the **test interface** command.

Example: Configuring Bit Error Rate Testing

Configure a BERT test on a T3 interface. In this example, the run duration lasts for 120 seconds. The configured error rate is 0, which corresponds to a bit error rate of 10^{-0} (1 error per bit). The configured bit pattern of **all-ones-repeating** means that every bit the interface sends is a set to a value of 1.

```
[edit interfaces]
t3-1/2/0 {
  t3-options {
    bert algorithm all-ones-repeating;
    bert-error-rate 0;
    bert-period 120;
  }
}
```

Tracing Operations of an Individual Router Interface

To trace the operations of individual router interfaces, include the **traceoptions** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
traceoptions {
  flag flag <disable>;
}
```

You can specify the following interface tracing flags:

- **all**—Trace all interface operations.
- **event**—Trace all interface events.
- **ipc**—Trace all interface interprocess communication (IPC) messages.
- **media**—Trace all interface media changes.

The interfaces **traceoptions** statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system **syslog** files.

For more information about trace operations, see “Tracing Operations of the Interface Process” on page 241.

Damping Interface Transitions

By default, when an interface changes from being up to being down, or from down to up, this transition is advertised immediately to the hardware and the JUNOS Software. In some situations—for example, when an interface is connected to an add-drop multiplexer (ADM) or wavelength-division multiplexer (WDM), or to protect against SONET/SDH framer holes—you might want to damp interface transitions. This means not advertising the interface’s transition until a certain period of time has passed, called the *hold-time*. When you have damped interface transitions and the interface goes from up to down, the interface is not advertised to the rest of the system as being down until it has remained down for the hold-time period. Similarly when an interface goes from down to up, it is not advertised as being up until it has remained up for the hold-time period.

To damp interface transitions, include the **hold-time** statement at the **[edit interfaces interface-name]** hierarchy level:

```
[edit interfaces interface-name]
hold-time up milliseconds down milliseconds;
```

The time can be a value from 0 through 4,294,967,295 milliseconds. The default value is 0, which means that interface transitions are not damped. The JUNOS Software advertises the transition within 100 milliseconds of the time value you specify.

For most Ethernet interfaces, hold timers are implemented using a 1-second polling algorithm. For 1-port, 2-port, and 4-port Gigabit Ethernet interfaces with small form-factor pluggable transceivers (SFPs), hold timers are interrupt-driven.



NOTE: The hold-time option is not available for controller interfaces.

Configuring Multiservice Physical Interface Properties

The adaptive services (AS), collector, monitoring services, and monitoring services II interfaces are multiservice interfaces specifically designed to enable IP services.

To configure multiservice physical interface properties on the collector, monitoring services, and AS interfaces, include the `multiservice-options` statement:

```
multiservice-options {
  (core-dump | no-core-dump);
  (syslog | no-syslog);
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *cp-fpc/pic/port*]
- [edit interfaces *mo-fpc/pic/port*]
- [edit interfaces *sp-fpc/pic/port*]

For more information about the services interfaces, see the *JUNOS Services Interfaces Configuration Guide*.

Enabling or Disabling SNMP Notifications on Physical Interfaces

By default, Simple Network Management Protocol (SNMP) notifications are sent when the state of an interface or a connection changes. To explicitly enable these notifications on the physical interface, include the `traps` statement at the [edit interfaces *interface-name*] hierarchy level. To disable these notifications on the physical interface, include the `no-traps` statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
(traps | no-traps);
```



NOTE: Gigabit Ethernet interfaces on J Series routers do not support SNMP.

Enabling Unidirectional Traffic Flow on Physical Interfaces

By default, physical interfaces are bidirectional; that is, they both transmit and receive traffic. You can configure unidirectional link mode on a 10-Gigabit Ethernet interface that creates two new physical interfaces that are unidirectional. The new transmit-only and receive-only interfaces operate independently, but both are subordinate to the original parent interface.

The unidirectional interfaces enable the configuration of a unidirectional link topology. Unidirectional links are useful for applications such as broadband video services where almost all traffic flow is in one direction, from the provider to the user. Unidirectional link mode conserves bandwidth by enabling it to be differentially dedicated to transmit and receive interfaces. In addition, unidirectional link mode conserves ports for such applications because the transmit-only and receive-only interfaces act independently. Each can be connected to different routers, for example, reducing the total number of ports required.

To enable unidirectional link mode on a physical interface, include the **unidirectional** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]  
unidirectional;
```



NOTE: Unidirectional link mode is currently supported on only the following hardware:

- 4-Port 10-Gigabit Ethernet DPC on the MX960 router
- 10-Gigabit Ethernet IQ2 PIC and 10-Gigabit Ethernet IQ2E PIC on the T Series router

The transmit-only interface is always operationally up. The operational status of the receive-only interface depends only on local faults; it is independent of remote faults and of the status of the transmit-only interface.

On the parent interface, you can configure attributes common to both interfaces, such as clocking, framing, *gigether-options*, and *sonet-options*. On each of the unidirectional interfaces, you can configure encapsulation, MAC address, MTU size, and logical interfaces.

Unidirectional interfaces support IP and IPv6. Packet forwarding takes place by means of static routes and static ARP entries, which you can configure independently on both unidirectional interfaces.

Only transmit statistics are reported on the transmit-only interface (and shown as zero on the receive-only interface). Only receive statistics are reported on the receive-only interface (and shown as zero on the transmit-only interface). Both transmit and receive statistics are reported on the parent interface.

Disabling a Physical Interface

You can disable a physical interface, marking it as being down, without removing the interface configuration statements from the configuration. To do this, include the **disable** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]  
disable;
```

Example: Disabling a Physical Interface

Disable a physical interface:

```
[edit interfaces]
so-1/1/0 {
  mtu 8000;
  clocking internal;
  encapsulation ppp;
  sonet-options {
    fcs 16;
  }
  unit 0 {
    family inet {
      address 172.16.0.0/12 {
        destination 172.16.0.4;
      }
    }
  }
}
[edit interfaces]
user@host# set so-1/1/0 disable
[edit interfaces]
user@host# show so-1/1/0
so-1/1/0 {
  disable;# Interface is marked as disabled
  mtu 8000;
  clocking internal;
  encapsulation ppp;
  sonet-options {
    fcs 16;
  }
  unit 0 {
    family inet {
      address 172.16.0.0 {
        destination 172.16.0.3;
      }
    }
  }
}
```


Chapter 4

Configuring Logical Interface Properties

For a physical interface device to function, you must configure at least one logical interface on that device. For each logical interface, you must specify the protocol family that the interface supports. You can also configure other logical interface properties. These vary by Physical Interface Card (PIC) and encapsulation type, but include the IP address of the interface, and whether the interface supports multicast traffic, data-link connection identifiers (DLCIs), virtual channel identifiers (VCIs) and virtual path identifiers (VPIs), and traffic shaping.

This chapter describes the configuration of logical interface properties:

- Logical Interfaces Configuration Statements on page 144
- Logical Interfaces Statements List on page 147
- Specifying the Logical Interface Number on page 155
- Configuring Logical System Interface Properties on page 155
- Adding a Logical Unit Description to the Configuration on page 156
- Configuring a Point-to-Point Connection on page 157
- Configuring a Multipoint Connection on page 157
- Configuring Accounting for the Logical Interface on page 157
- Configuring the Interface Bandwidth on page 159
- Enabling or Disabling SNMP Notifications on Logical Interfaces on page 159
- Configuring Interface Encapsulation on Logical Interfaces on page 160
- Configuring the LCP Configure-Request Maximum Sent on page 161
- Configuring the NCP Configure-Request Maximum Sent on page 161
- Configuring the PPP Restart Timers on page 162
- Configuring the PPP Clear Loop Detected Timer on page 162
- Configuring Dynamic Profiles for PPP on page 163
- Configuring PPP CHAP Authentication on page 163
- Configuring PPP PAP Authentication on page 164
- Configuring Dynamic Call Admission Control on page 166
- Disabling a Logical Interface on page 167

Logical Interfaces Configuration Statements

To configure logical interface properties, include the following statements:

```

unit logical-unit-number {
  accept-source-mac {
    mac-address mac-address {
      policer {
        input cos-policer-name;
        output cos-policer-name;
      }
    }
  }
  accounting-profile name;
  allow-any-vci;
  atm-scheduler-map (map-name | default);
  backup-options {
    interface interface-name;
  }
  bandwidth rate;
  cell-bundle-size cells;
  clear-dont-fragment-bit;
  compression {
    rtp {
      f-max-period number;
      queues [ queue-numbers ];
      port {
        minimum port-number;
        maximum port-number;
      }
    }
  }
  compression-device interface-name;
  copy-tos-to-outer-ip-header;
  demux-destination family;
  demux-source family;
  demux-options {
    underlying-interface interface-name;
  }
  description text;
  dial-options {
    l2tp-interface-id name;
    (dedicated | shared);
  }
  dialer-options {
    activation-delay seconds;
    callback;
    callback-wait-period time;
    deactivation-delay seconds;
    dial-string [ dial-string-numbers ];
    idle-timeout seconds;
    incoming-map {
      caller (caller-id | accept-all);
      initial-route-check seconds;
    }
  }
}

```



```

        load-interval seconds;
        load-threshold number;
        pool pool-name;
        redial-delay time;
        watch-list {
            [ routes ];
        }
    }
}
disable;
disable-mlppp-inner-ppp-pfc;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
    activation-priority priority;
    bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold plp1 cells;
filter filter-name;
fragment-threshold bytes;
inner-vlan-id-range start start-id end end-id;
input-vlan-map {
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
}
interleave-fragments;
inverse-arp;
link-layer-overhead percent;
layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
    up-count cells;
    down-count cells;
}
oam-period (seconds | disable);
output-vlan-map {
    inner-tag-protocol-id;
    inner-vlan-id;
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    tag-protocol-id tpid;
    vlan-id number;
}

```

```

passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
    dynamic-profile profile-name;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-restart-timer milliseconds;
    pap {
        default-pap-password password;
        local-name name;
        local-password password;
        passive;
    }
    pppoe-options {
        access-concentrator name;
        auto-reconnect seconds;
        (client | server);
        service-name name;
        underlying-interface interface-name;
    }
    proxy-arp;
    service-domain (inside | outside);
    shaping {
        (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
         rate burst length);
        queue-length number;
    }
    short-sequence;
    transmit-weight number;
    (traps | no-traps);
    trunk-bandwidth rate;
    trunk-id number;
    tunnel {
        backup-destination address;
        destination address;
        key number;
        routing-instance {
            destination routing-instance-name;
        }
        source source-address;
        ttl number;
    }
    vci vpi-identifier.vci-identifier;
    vci-range start start-vci end end-vci;

```

```

vpi vpi-identifier;
vlan-id number;
vlan-id-range number-number;
vlan-tags (Stacked VLAN Tags) inner tpid.vlan-id outer tpid.vlan-id;
family family {
    [ family-statements ];
}

```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name*]
- [edit logical-systems *logical-system-name* interfaces *interface-name*]

For information about interface-specific logical properties, see “Logical Interfaces Statements List” on page 147.

Logical Interfaces Statements List

Table 23 on page 147 lists statements that you can use to configure logical interfaces.

Table 23: Statements for Logical Interface Properties

Statement	Interface Types	Usage Guidelines
access-profile <i>name</i>	ATM2 IQ interfaces	“Configuring PPP PAP Authentication” on page 164
accept-source-mac	Gigabit Ethernet intelligent queuing (IQ) interfaces	“Configuring MAC Address Filtering” on page 761
accounting-profile <i>name</i>	All	“Configuring Accounting for the Logical Interface” on page 157
ack-delay-time <i>time</i>	Ethernet interfaces configured for DLSw	“Configuring LLC2 Options” on page 180
ack-max <i>count</i>	Ethernet interfaces configured for DLSw	“Configuring LLC2 Options” on page 180
activation-delay <i>seconds</i>	ISDN interfaces	“Configuring ISDN Interfaces” on page 819
activation-priority <i>priority</i>	Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation on J4350 and J6350 Services Routers supporting voice over IP with the TGM550 media gateway module	“Configuring Dynamic Call Admission Control” on page 166
adaptive-shapers <i>adaptive-shaper-name</i>	Frame Relay interfaces on J Series routers	<i>JUNOS Class of Service Configuration Guide</i>
allow-any-vci	Asynchronous Transfer Mode (ATM) interfaces	“Configuring ATM Interface Encapsulation” on page 330

Table 23: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
atm-scheduler-map (<i>map-name</i> default)	ATM2 IQ interfaces	“Configuring ATM2 IQ VC Tunnel CoS Components” on page 339
backup-destination <i>address</i>	Encryption interfaces	<i>JUNOS Class of Service Configuration Guide</i>
backup-options	J Series routers ISDN interfaces	“Configuring an ISDN Dialer Interface as a Backup Interface” on page 826
bandwidth <i>rate</i>	All interfaces, except multilink and aggregated	“Configuring the Interface Bandwidth” on page 159
bearer-bandwidth-limit <i>kilobits-per-second</i>	Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation on J4350 and J6350 Services Routers supporting voice over IP with the TGM550 media gateway module	“Configuring Dynamic Call Admission Control” on page 166
cbr <i>rate</i>	ATM interfaces	“Defining the ATM Traffic-Shaping Profile” on page 319
cell-bundle-size <i>cells</i>	ATM2 IQ interfaces	“Configuring the Layer 2 Circuit Cell-Relay Cell Maximum” on page 313
clear-dont-fragment-bit	Adaptive services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
compression	AS PIC or MultiServices PIC link services IQ interfaces (<i>lsq</i>) and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
compression-device <i>interface-name</i>	J Series routers E1 and T1 interfaces.	<i>JUNOS Services Interfaces Configuration Guide</i>
copy-tos-to-outer-ip-header	GRE tunnel interfaces	<i>JUNOS Class of Service Configuration Guide</i>
deactivation-delay <i>seconds</i>	ISDN interfaces	“Configuring ISDN Interfaces” on page 819
demux-destination <i>family</i>	IP demux interfaces	“Configuring an IP Demux Underlying Interface” on page 252
demux-options <i>family</i>	IP demux interfaces	“Specifying the Demux Underlying Interface” on page 253
demux-source <i>family</i>	IP demux interfaces	“Configuring an IP Demux Underlying Interface” on page 252
description <i>text</i>	All	“Adding a Logical Unit Description to the Configuration” on page 156
destination (<i>address</i> <i>routing-instance-name</i>)	Encryption generic routing encapsulation (GRE) tunnel, and IP tunnel interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>

Table 23: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
dialer-options	Adaptive services interfaces on M7i routers J Series routers ISDN interfaces	<i>JUNOS Services Interfaces Configuration Guide</i> “Configuring ISDN Physical Interface Properties” on page 821
disable	All	“Disabling a Logical Interface” on page 167
disable-mlppp-inner-ppp-pfc	MLPPP interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
dlci <i>dlci-identifier</i>	Point-to-point interfaces with Frame Relay encapsulation	“Configuring Frame Relay DLCIs” on page 380
drop-timeout <i>milliseconds</i>	Multilink interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
dynamic-call-admission-control	Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation on J4350 and J6350 Services Routers supporting voice over IP with the TGM550 media gateway module	“Configuring Dynamic Call Admission Control” on page 166
dynamic-profile <i>profile-name</i>	1-Gigabit Ethernet and 10-Gigabit Ethernet interfaces configured with PPP over Ethernet on M120 and M320 routers	<i>JUNOS Subscriber Access Configuration Guide</i>
encapsulation <i>type</i>	All interfaces, except aggregated SONET/SDH and loopback	“Configuring the Encapsulation on a Logical Interface” on page 160
epd-threshold <i>cells</i>	ATM2 IQ interfaces	“Configuring the ATM2 IQ EPD Threshold” on page 326
f-max-period <i>number</i>	AS PIC or MultiServices link services IQ interfaces (<i>lsq</i>) and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
family	All	“Configuring the Protocol Family” on page 172
fragment-threshold <i>bytes</i>	Multilink interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
frame-relay <i>map-name</i> default)	Frame Relay Interfaces on J Series routers	<i>JUNOS Services Interfaces Configuration Guide</i> and <i>JUNOS Class of Service Configuration Guide</i>
idle-time <i>time</i>	Ethernet interfaces configured for DLSw	“Configuring LLC2 Options” on page 180
idle-timeout	ISDN interfaces	“Configuring Bandwidth on Demand” on page 829

Table 23: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
initial-route-check <i>seconds</i>	ISDN interfaces	“Configuring ISDN Logical Interface Properties” on page 823
inner-tag-protocol-id	Gigabit Ethernet IQ interfaces	“Configuring 802.1Q VLANs” on page 599
inner-vlan-id	Gigabit Ethernet IQ interfaces	“Configuring 802.1Q VLANs” on page 599
inner-vlan-id-range	Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet IQ interfaces	“Configuring ATM-to-Ethernet Interworking” on page 229
input	AS PIC or MultiServices link services	<i>JUNOS Services Interfaces Configuration Guide</i>
input-policer <i>policer-name</i>	For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers	<i>JUNOS Services Interfaces Configuration Guide</i> and “Applying a Policer” on page 765
input-three-color <i>policer-name</i>	For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers	<i>JUNOS Class of Service Configuration Guide</i> and “Applying a Policer” on page 765
input-vlan-map	Gigabit Ethernet IQ interfaces	“Configuring the Management Ethernet Interface” on page 775
interleave-fragments	Link services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
inverse-arp	Interfaces with ATM and Frame Relay encapsulation	“Configuring Inverse ATM1 or ATM2 ARP” on page 318 and “Configuring Inverse Frame Relay ARP” on page 379
key <i>number</i>	GRE tunnel interfaces on Adaptive Services PICs	<i>JUNOS Services Interfaces Configuration Guide</i>
layer2-policer	1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces	“Applying a Policer” on page 765
lcp-restart-timer	Interfaces with PPP encapsulation	“Configuring the PPP Restart Timers” on page 162
l2tp-interface-id <i>name</i>	Adaptive services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
link-layer-overhead <i>percent</i>	AS PIC or MultiServices link services IQ interfaces (lsq)	<i>JUNOS Services Interfaces Configuration Guide</i>
llc2	Ethernet interfaces configured for DLSw on J Series routers	“Configuring LLC2 Options” on page 180
load-threshold <i>number</i>	ISDN interfaces	“Configuring Bandwidth on Demand” on page 829

Table 23: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
local-name <i>name</i>	ATM2 IQ interfaces	“Configuring PPP CHAP Authentication” on page 163 and “Configuring PPP PAP Authentication” on page 164
local-window <i>count</i>	Ethernet interfaces configured for DLSw	“Configuring LLC2 Options” on page 180
loss-priority-maps	Frame Relay interfaces on J Series routers	<i>JUNOS Services Interfaces Configuration Guide</i> and <i>JUNOS Class of Service Configuration Guide</i>
mac-address <i>mac-address</i>	Gigabit Ethernet interfaces and Gigabit Ethernet IQ and IQE interfaces with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	“Configuring MAC Address Filtering” on page 761
max-retry <i>count</i>	Ethernet interfaces configured for DLSw	“Configuring LLC2 Options” on page 180
minimum-links <i>number</i>	Multilink interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
mrru <i>bytes</i>	Multilink interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
multicast-dlci <i>dlci-identifier</i>	Point-to-multipoint Frame Relay interfaces	“Configuring a Multicast-Capable Frame Relay Connection” on page 381
multicast-vc vpi-identifier vci-identifier	Point-to-multipoint ATM1 and ATM2 IQ interfaces	“Configuring the ATM OAM F5 Loopback Cell Threshold” on page 329
multilink-max-classes <i>number</i>	AS PIC or MultiServices link services IQ interfaces (lsq-)	<i>JUNOS Services Interfaces Configuration Guide</i>
multipoint	All	“Configuring a Multipoint Connection” on page 157
ncp-restart-timer	Interfaces with PPP encapsulation	“Configuring the PPP Restart Timers” on page 162
oam-liveness	ATM1 and ATM2 IQ interfaces	“Configuring the ATM OAM F5 Loopback Cell Threshold” on page 329
oam-period (disable seconds)	ATM1 and ATM2 IQ interfaces	“Defining the ATM OAM F5 Loopback Cell Period” on page 329
output	All	<i>JUNOS Services Interfaces Configuration Guide</i>
output-policer <i>policer-name</i>	For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers	<i>JUNOS Class of Service Configuration Guide</i> and “Applying a Policer” on page 765

Table 23: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
output-three-color <i>policer-name</i>	For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers	<i>JUNOS Class of Service Configuration Guide</i> and “Applying a Policer” on page 765
output-vlan-map	Gigabit Ethernet IQ interfaces	“Configuring the Management Ethernet Interface” on page 775
passive	ATM2 IQ interfaces	“Configuring PPP CHAP Authentication” on page 163 and “Configuring PPP PAP Authentication” on page 164
passive-monitor-mode	SONET/SDH interfaces	“Enabling Passive Monitoring on SONET/SDH Interfaces” on page 874
p-bit-timeout <i>time</i>	Ethernet interfaces configured for DLSw	“Configuring LLC2 Options” on page 180
peer-unit <i>unit-number</i>	Logical tunnel interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
pfc	Interfaces with PPP, PPP CCC, or PPP TCC encapsulation	“Configuring the PPP Protocol Field Compression” on page 121
plp1 <i>cells</i>	ATM2 IQ interfaces	“Configuring the ATM2 IQ EPD Threshold” on page 326
plp-to-clp	ATM2 IQ interfaces	“Enabling the PLP Setting to Be Copied to the CLP Bit” on page 348
point-to-point	All	“Configuring a Point-to-Point Connection” on page 157
policer	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router)	“Configuring MAC Address Filtering” on page 761
pop	Gigabit Ethernet IQ interfaces	“Removing a VLAN Tag” on page 649
pop-pop	Gigabit Ethernet IQ interfaces	“Removing the Outer and Inner VLAN Tags” on page 649
pop-swap	Gigabit Ethernet IQ interfaces	“Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag” on page 650
port	AS PIC or MultiServices or MultiServices link services IQ interfaces (<i>lsq</i>) and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
ppp-options	Interfaces with PPP, PPP CCC, or PPP TCC encapsulation	“Configuring PPP CHAP Authentication” on page 163 and “Configuring PPP PAP Authentication” on page 164

Table 23: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
proxy-arp	Ethernet interfaces	“Configuring Unrestricted Proxy ARP” on page 671
push	Gigabit Ethernet IQ interfaces	“Stacking a VLAN Tag” on page 648
push-push	Gigabit Ethernet IQ interfaces	“Stacking Two VLAN Tags” on page 651
queue-length <i>number</i>	ATM1 interfaces	“Configuring the ATM1 Queue Length” on page 325
queues [<i>queue-numbers</i>]	AS PIC or MultiServices link services IQ interfaces (lsq) and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
routing-instance	GRE tunnel and IP tunnel interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
rtp	AS PIC or MultiServices link services IQ interfaces (lsq) and voice services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
rtvbr peak <i>rate</i> sustained <i>rate</i> burst <i>length</i>	ATM2 interfaces	“Configuring ATM2 IQ Real-Time VBR” on page 321
service-domain (inside outside)	Adaptive services interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
shaping	ATM1 and ATM2 IQ interfaces	“Defining the ATM Traffic-Shaping Profile” on page 319
short-sequence	Multilink interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
source <i>source-address</i>	Encryption, GRE tunnel, and IP tunnel interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
swap	Gigabit Ethernet IQ interfaces	“Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames” on page 644
swap-push	Gigabit Ethernet IQ interfaces	“Rewriting a VLAN Tag and Adding a New Tag” on page 655
swap-swap	Gigabit Ethernet IQ interfaces	“Rewriting the Inner and Outer VLAN Tags” on page 655
t1-time <i>time</i>	Ethernet interfaces configured for DLSw	“Configuring LLC2 Options” on page 180
t2-time <i>time</i>	Ethernet interfaces configured for DLSw	“Configuring LLC2 Options” on page 180

Table 23: Statements for Logical Interface Properties (continued)

Statement	Interface Types	Usage Guidelines
tag-protocol-id <i>tpid</i>	Gigabit Ethernet and Gigabit Ethernet IQ and IQE PICs with SFPs (except the 10-port Gigabit Ethernet PIC, Aggregated Ethernet with Gigabit Ethernet IQ interfaces, and the built-in Gigabit Ethernet port on the M7i router)	“Rewriting the VLAN Tag on Tagged Frames” on page 651
transmit-weight <i>number</i>	ATM2 IQ interfaces	“Configuring the ATM2 IQ Transmission Weight” on page 329
trej-time <i>time</i>	Ethernet interfaces configured for DLSw	“Configuring LLC2 Options” on page 180
(traps no-traps)	All	“Enabling or Disabling SNMP Notifications on Logical Interfaces” on page 159
trunk-bandwidth <i>rate</i>	ATM2 IQ interfaces	“Configuring Layer 2 Circuit Trunk Mode Scheduling” on page 309
trunk-id <i>number</i>	ATM2 IQ interfaces	“Configuring Layer 2 Circuit Transport Mode” on page 300
ttl <i>number</i>	GRE tunnel and IP tunnel interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
tunnel	Encryption, GRE tunnel, and IP tunnel interfaces	<i>JUNOS Services Interfaces Configuration Guide</i>
underlying-interface	IP demux interfaces	“Specifying the Demux Underlying Interface” on page 253
vbr peak <i>rate</i> sustained <i>rate</i> burst <i>length</i>	ATM interfaces	“Defining the ATM Traffic-Shaping Profile” on page 319
vci <i>vpi-identifier</i> <i>vci-identifier</i>	ATM1 and ATM2 IQ point-to-point interfaces	“Configuring a Point-to-Point ATM1 or ATM2 IQ Connection” on page 316
vci-range	ATM2 IQ interfaces	“Configuring ATM-to-Ethernet Interworking” on page 229
vpi <i>vpi-identifier</i>	ATM1 and ATM2 IQ point-to-point interfaces	“Configuring a Point-to-Point ATM1 or ATM2 IQ Connection” on page 316
vlan-id <i>number</i>	Fast Ethernet, Gigabit Ethernet, and Gigabit Ethernet IQ interfaces and aggregated Ethernet using Gigabit Ethernet IQ interfaces	“Binding VLAN IDs to Logical Interfaces” on page 604 and “Rewriting the VLAN Tag on Tagged Frames” on page 651
vlan-tags (Stacked VLAN Tags) inner <i>tpid</i> <i>vlan-id</i> outer <i>tpid</i> <i>vlan-id</i>	Gigabit Ethernet IQ interfaces	“Configuring Dual VLAN Tags” on page 645
watch-list	ISDN interfaces	“Configuring Dialer Watch” on page 835

Specifying the Logical Interface Number

Each logical interface must have a logical unit number. The logical unit number corresponds to the logical unit part of the interface name. For more information, see “Interface Naming” on page 51.

Point-to-Point Protocol (PPP), Cisco High-level Data Link Control (HDLC), and Ethernet circuit cross-connect (CCC) encapsulations support only a single logical interface, whose logical unit number must be 0. Frame Relay and ATM encapsulations support multiple logical interfaces, so you can configure one or more logical unit numbers.

You specify the logical unit number by including the **unit** statement:

```
unit logical-unit-number {
  ...
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name*]
- [edit logical-systems *logical-system-name* interfaces *interface-name*]

The logical unit number can be from 0 through 16,384.

Configuring Logical System Interface Properties

With JUNOS Software, you can partition a single physical router into multiple logical devices that perform independent routing tasks. Because logical systems perform a subset of the tasks once handled by the physical router, logical systems offer an effective way to maximize the use of a single router.

You can include the following logical system statements:

```
[edit logical-systems logical-system-name]
interfaces interface-name {
  unit logical-unit-number {
    logical-interface-statements;
  }
}
policy-options {
  policy-options-statements;
}
protocols {
  protocols-statements;
}
routing-instances {
  routing-instances-statements;
}
routing-options {
  routing-options-statements;
}
```

For an overview of logical systems, see the *JUNOS Feature Guide*. For detailed information about logical system configuration, see the *JUNOS Routing Protocols Configuration Guide*. For information about configuring peer relationships between logical systems, see *JUNOS Services Interfaces Configuration Guide*.

To configure interface properties of a logical system, you must include the following statements at the [edit logical-systems *logical-system-name*] hierarchy level:

```
[edit logical-systems logical-system-name]
interfaces interface-name {
  unit logical-unit-number {
    logical-interface-statements;
  }
}
```

Example: Configuring Logical System Interface Properties

Configure a logical system's interface properties:

```
[edit interfaces t3-0/0/1]
description "Physical interface to be partitioned into multiple logical systems";
[edit logical-systems 1-on-t3-0/0/1]
interfaces t3-0/0/1 {
  unit 1 {
    family inet {
      address 10.0.0.1/32 {
        destination 10.0.0.2;
      }
    }
  }
}
```

Adding a Logical Unit Description to the Configuration

You can include a text description of each logical unit in the configuration file. Any descriptive text you include is displayed in the output of the `show interfaces` commands, and is also exposed in the ifAlias Management Information Base (MIB) object. It has no impact on the interface's configuration. To add a text description, include the `description` statement:

```
description text;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The description can be a single line of text. If the text contains spaces, enclose it in quotation marks.



NOTE: You can configure the extended DHCP relay to include the interface description in the option 82 Agent Circuit ID suboption. See Enabling and Disabling Insertion of Option 82 Information in the *JUNOS Subscriber Access Configuration Guide*.

For information about describing physical interfaces, see “Adding an Interface Description to the Configuration” on page 96.

Configuring a Point-to-Point Connection

By default, all interfaces are assumed to be point-to-point connections. You must ensure that the maximum transmission unit (MTU) sizes on both sides of the connection are the same.

For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, you can explicitly configure an interface to be a point-to-point connection by including the `point-to-point` statement:

```
point-to-point;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Configuring a Multipoint Connection

By default, all interfaces are assumed to be point-to-point connections. To configure an interface to be a multipoint connection, include the `multipoint` statement:

```
multipoint;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Configuring Accounting for the Logical Interface

Juniper Networks routing platforms can collect various kinds of data about traffic passing through the routing platform. You can set up one or more *accounting profiles* that specify some common characteristics of this data, including the following:

- The fields used in the accounting records
- The number of files that the router retains before discarding, and the number of bytes per file

- The period that the system uses to record the data

You configure the profiles and define a unique name for each profile using statements at the `[edit accounting-options]` hierarchy level. There are two types of accounting profiles: interface profiles and filter profiles. You configure interface profiles by including the `interface-profile` statement at the `[edit accounting-options]` hierarchy level. You configure filter profiles by including the `filter-profile` statement at the `[edit accounting-options]` hierarchy level. For more information, see the *JUNOS Network Management Configuration Guide*.

You apply filter profiles by including the `accounting-profile` statement at the `[edit firewall filter filter-name]` and `[edit firewall family family filter filter-name]` hierarchy levels. For more information, see the *JUNOS Policy Framework Configuration Guide*.

Applying an Accounting Profile to the Logical Interface

To enable accounting on a logical interface, include the `accounting-profile` statement:

```
accounting-profile name;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number]`

You can also reference profiles for the physical interface; for more information, see “Configuring Accounting for the Physical Interface” on page 130.

Example: Applying an Accounting Profile to the Logical Interface

Configure an accounting profile for an interface and apply it to a logical interface:

```
[edit]
accounting-options {
  file if_stats {
    size 4m files 10 transfer-interval 15;
    archive-sites {
      "ftp://login:password@host/path";
    }
  }
  interface-profile if_profile {
    interval 15;
    file if_stats {
      fields {
        input-bytes;
        output-bytes;
        input-packets;
        output-packets;
        input-errors;
        output-errors;
      }
    }
  }
}
```

```
[edit interfaces ge-1/0/1 unit 1]
  accounting-profile if_profile;
```

To reference profiles by physical interface, see “Applying an Accounting Profile to the Physical Interface” on page 130. For information about configuring a firewall filter accounting profile, see the *JUNOS Policy Framework Configuration Guide*.

Configuring the Interface Bandwidth

By default, the JUNOS Software uses the physical interface’s speed for the MIB-II object, `ifSpeed`. You can configure the logical unit to populate the `ifSpeed` variable by configuring a bandwidth value for the logical interface. The `bandwidth` statement sets an informational-only parameter; you cannot adjust the actual bandwidth of an interface with this statement.

To configure the bandwidth value for a logical interface, include the `bandwidth` statement:

```
bandwidth rate;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

rate is the peak rate, in bps or cps. You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation c; values expressed in cells per second are converted to bits per second using the formula 1 cps = 384 bps. The value can be any positive integer. The `bandwidth` statement is valid for all logical interfaces, except multilink and aggregated interfaces.

Enabling or Disabling SNMP Notifications on Logical Interfaces

By default, Simple Network Management Protocol (SNMP) notifications are sent when the state of an interface or a connection changes. To explicitly enable these notifications on the logical interface, include the `traps` statement; to disable these notifications on the logical interface, include the `no-traps` statement:

```
(traps | no-traps);
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]



NOTE: Gigabit Ethernet interfaces on J Series routers do not support SNMP.

Configuring Interface Encapsulation on Logical Interfaces

PPP encapsulation is the default encapsulation type for physical interfaces. You need not configure encapsulation for any physical interfaces that support PPP encapsulation. If you do not configure encapsulation, PPP is used by default. For physical interfaces that do not support PPP encapsulation, you must configure an encapsulation to use for packets transmitted on the interface. For more information about physical interface encapsulation, see “Configuring the Encapsulation on a Physical Interface” on page 106.

You can optionally configure an encapsulation on a logical interface, which is the encapsulation used within certain packet types.

Configuring the Encapsulation on a Logical Interface

Generally, you configure an interface’s encapsulation at the [edit interfaces *interface-name*] hierarchy level. However, for some encapsulation types, such as Frame Relay, ATM, and Ethernet virtual local area network (VLAN) encapsulations, you can also configure the encapsulation type that is used inside the Frame Relay, ATM, or VLAN circuit itself. To do this, include the **encapsulation** statement:

```
encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-tcc-vc-mux | atm-cisco-nlpid |
  atm-mlppp-llc | atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap |
  atm-vc-mux | ether-over-atm-llc | ether-vpls-over-atm-llc | ethernet |
  frame-relay-ether-type | frame-relay-ether-type-tcc | frame-relay-ccc | frame-relay-tcc
  | multilink-frame-relay-end-to-end | multilink-ppp | vlan-ccc | vlan-tcc | vlan-vpls);
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Some of the ATM encapsulations are defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*.

The following restrictions apply to logical interface encapsulation:

- With the atm-nlpid, atm-cisco-nlpid, and atm-vc-mux encapsulations, you can configure the inet family only.
- With the CCC circuit encapsulations, you cannot configure a family on the logical interface.
- A logical interface cannot have frame-relay-ccc encapsulation unless the physical device also has frame-relay-ccc encapsulation.
- A logical interface cannot have frame-relay-tcc encapsulation unless the physical device also has frame-relay-tcc encapsulation. In addition, you must assign this logical interface a DLCI from 512 through 1022 and configure it as point-to-point.

- A logical interface cannot have frame-relay-ether-type or frame-relay-ether-type-tcc encapsulation unless the physical interface has flexible-frame-relay encapsulation and is on an IQ or IQE PIC.
- For frame-relay-ether-type-tcc encapsulation, you must assign this logical interface a DLCI from 512 through 1022.
- For interfaces that carry IP version 6 (IPv6) traffic, you cannot configure ether-over-atm-llc encapsulation.
- When you use ether-over-atm-llc encapsulation, you cannot configure multipoint interfaces.
- A logical interface cannot have vlan-ccc or vlan-vpls encapsulation unless the physical device also has vlan-ccc or vlan-vpls encapsulation, respectively. In addition, you must assign this logical interface a VLAN ID from 512 through 1023; if the VLAN ID is 511 or lower, it is subject to the normal destination filter lookups in addition to source address filtering. For more information, see “Configuring VLAN Encapsulation” on page 609.
- You can create an ATM cell-relay circuit by configuring an entire ATM physical device or an individual virtual circuit (VC). When you configure an entire device, only cell-relay encapsulation is allowed on the logical interfaces. For more information, see “Configuring an ATM1 Cell-Relay Circuit” on page 332.

For more information about ATM encapsulations, see “Configuring ATM Interface Encapsulation” on page 330.

For more information about Frame Relay encapsulations, see “Configuring Frame Relay Interface Encapsulation” on page 372.

For more information about multilink encapsulations, see the *JUNOS Services Interfaces Configuration Guide*.

Configuring the LCP Configure-Request Maximum Sent

Link Control Protocol (LCP) Configure-Request is used to establish a link. You can configure the maximum number of LCP Configure-Requests to send. The router stops sending LCP Configure-Requests after the specified maximum number is sent. To configure the LCP Configure-Request maximum, use the `lcp-max-conf-req` statement at the `[edit interfaces interface-name unit number ppp-options]` hierarchy level. The *number* range is from 0 to 65,535; where 0 specifies no limit and the LCP Configure-Request is sent indefinitely.

Configuring the NCP Configure-Request Maximum Sent

Network Control Protocol (NCP) Configure-Request is used to establish a link. You can configure the maximum number of NCP Configure-Requests to send. The router stops sending NCP Configure-Requests after the specified maximum number is sent. To configure the NCP Configure-Request maximum, use the `ncp-max-conf-req` statement at the `[edit interfaces interface-name unit number ppp-options]` hierarchy level. The *number* range is from 0 to 65,535; where 0 specifies no limit and NCP Configure-Request is sent indefinitely.

Configuring the PPP Restart Timers

You can configure a restart timer for the Link Control Protocol (LCP) and Network Control Protocol (NCP) components of a PPP session. You can configure the LCP restart timer on interfaces with PPP, PPP TCC, PPP over Ethernet, PPP over ATM, and PPP over Frame Relay encapsulations. You can configure the NCP restart timer on interfaces with PPP and PPP TCC encapsulations and on multilink PPP bundle interfaces.

To configure the restart timer for the NCP component of a PPP session, include the `ncp-restart-timer` statement, and specify the number of milliseconds.

To configure the restart timer for the LCP component of a PPP session, include the `lcp-restart-timer` statement, and specify the number of milliseconds:

```
lcp-restart-timer milliseconds;
ncp-restart-timer milliseconds;
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* ppp-options]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* ppp-options]

To monitor the configuration, issue the `show interfaces interface-name` command. Configured options are displayed in the PPP parameters field for the physical interface.

```
user@host> run show interfaces t1-0/0/0:1:1.0 detail
Logical interface t1-0/0/0:1:1.0 (Index 67) (SNMP ifIndex 40)
(Generation 156)
Flags: Hardware-Down Device-Down Point-To-Point SNMP-Traps 0x4000
Encapsulation: PPP
PPP parameters:
  LCP restart timer: 2000 msec
  NCP restart timer: 2000 msec
Protocol inet, MTU: 1500, Generation: 163, Route table: 0
Flags: Protocol-Down
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 1.1.1/24, Local: 1.1.1.2, Broadcast: 1.1.1.255,
```

Configuring the PPP Clear Loop Detected Timer

When a Point-to-Point Protocol (PPP) session detects a loop, the loop detected flag is set. If the flag is not cleared by the protocol after the loopback is cleared, the clear loop detected timer clears the flag after the specified time has elapsed.

To configure the clear loop detected timer for the LCP component of a PPP session, include the `loopback-clear-timer` statement, and specify the number of seconds.

```
loopback-clear-timer seconds;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* ppp-options]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* ppp-options]

To monitor the configuration, issue the `show interfaces interface-name extensive` command.

Configuring Dynamic Profiles for PPP

A dynamic profile acts as a template that enables you to create, update, or remove a configuration that includes attributes for client access (for example, interface or protocol) or service (for example, IGMP). Using these profiles you can consolidate all of the common attributes of a client (and eventually a group of clients) and apply the attributes simultaneously.

After they are created, the profiles reside in a profile library on the router. You can then use the `dynamic-profile` statement to attach profiles to interfaces. To assign a dynamic profile to a PPP interface, you can include the `dynamic-profile` statement at the [edit interfaces *interface-name* unit *logical-unit-number* ppp-options] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number ppp-options]
dynamic-profile profile-name;
```

To monitor the configuration, issue the `show interfaces interface-name` command.

For information about dynamic profiles, see Dynamic Profiles Overview in the *JUNOS Subscriber Access Configuration Guide*.

For information about creating dynamic profiles, see Configuring a Basic Dynamic Profile in the *JUNOS Subscriber Access Configuration Guide*.

For information about assigning a dynamic profile to a PPP interface, see Attaching Dynamic Profiles to PPP Subscriber Interfaces in the *JUNOS Subscriber Access Configuration Guide*.



NOTE: Dynamic profiles for PPP subscribers are supported only on PPPoE interfaces for this release.

Configuring PPP CHAP Authentication

For interfaces with PPP encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP), as defined in RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*. When you enable CHAP on an interface, the interface can authenticate its peer and can be authenticated by its peer.

For information about configuring CHAP on logical interfaces, see the *JUNOS System Basics Configuration Guide*.

For information about configuring CHAP on physical interfaces, see “Configuring the PPP Challenge Handshake Authentication Protocol” on page 112

Configuring PPP PAP Authentication

The Password Authentication Protocol (PAP) provides a simple method for the peer to establish its identity using a 2-way handshake. This is done only upon initial link establishment.

After the link is established, an ID and password pair is repeatedly sent by the peer to the authenticator until authentication is acknowledged or the connection is terminated.

To configure PAP, you must create an access profile, configure tracing operations, and configure the logical and physical interfaces.

To configure PAP on a logical interface with PPP encapsulation, include the `pap` statement with options:

```
pap {
  default-pap-password password;
  local-name name;
  local-password password;
  passive;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For more information about configuring PAP for physical interfaces, see “Configuring the PPP Password Authentication Protocol” on page 114. For information about configuring tracing operations for the PPP protocol, see “Tracing Operations of the pppd Process” on page 119.

On each logical interface with PPP encapsulation, you can perform the following tasks:

- Configuring a Default PAP Password on page 165
- Configuring the Local Name on page 165
- Configuring the Local Password on page 165
- Configuring Passive Mode on page 166

Configuring a Default PAP Password

The default PAP password is used when no matching PAP access profile exists, or if the PAP access profile name changes during PPP link negotiation.

To configure a default PAP password for an interface, include the `default-pap-password` statement:

```
default-pap-password password;
```

You can include the statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* ppp-options pap]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* ppp-options pap]

Configuring the Local Name

By default, when PAP is enabled on an interface, the interface uses the router's system hostname as the name sent in PAP request and response packets.

To configure the name the interface uses in PAP request and response packets, include the `local-name` statement:

```
local-name name;
```

You can include the statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* ppp-options pap]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* ppp-options pap]

Configuring the Local Password

You need to configure the password to be used for authentication.

To configure the host password for sending PAP requests, include the `local-password` statement:

```
local-password password;
```

You can include the statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* ppp-options pap]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* ppp-options pap]

Configuring Passive Mode

By default, when PAP is enabled on an interface, the interface expects authenticate-request packets from the peer. However, if a peer does not support bidirectional authentication, you can configure PAP to operate in passive mode. In PAP passive mode, the interface sends the authenticate-request packets to the peer only if the interface receives the PAP option from the peer—in passive mode, the interface does not authenticate the peer.

To configure the interface to authenticate with PAP in passive mode, include the `passive` statement:

```
passive;
```

You can include the statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* ppp-options pap]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* ppp-options pap]

Configuring Dynamic Call Admission Control

Dynamic call admission control (CAC) provides enhanced control over WAN bandwidth. You can configure dynamic CAC on J4350 and J6350 Services Routers supporting voice over IP through the TGM550 media gateway module. It can be used with the following interfaces:

- Fast Ethernet and Gigabit Ethernet interfaces
- ISDN BRI interfaces
- Serial interfaces with PPP or Frame Relay encapsulation

When dynamic CAC is configured on an interface responsible for providing call bandwidth, the TGM550 informs the Media Gateway Controller (MGC) of the bandwidth limit available for voice packets on the interface and requests the MGC to block new calls when the bandwidth is exhausted.

Dynamic CAC is useful when a primary link becomes unavailable and a backup link with less bandwidth takes its place. Without dynamic CAC, the MGC cannot detect the switchover to the backup link or the resulting changes in network topology and available bandwidth. The MGC would continue to admit calls at the bandwidth of the primary link, causing network congestion and possible jitter, delay, and loss of calls.

To configure dynamic CAC for a logical interface, include the `dynamic-call-admission-control` statement, with options:

```
dynamic-call-admission-control {
  activation-priority priority;
  bearer-bandwidth-limit kilobits-per-second;
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

bearer-bandwidth-limit *kilobits-per-second* is the dynamic CAC bearer bandwidth limit (BBL)—the maximum bandwidth available for voice traffic on the interface. The TGM550 reports the BBL to the MGC. When the call bandwidth exceeds the BBL, the MGC blocks new calls and alerts the user with a busy tone. The BBL range is from 0 through 9999. The default BBL is -1, which indicates that dynamic CAC is not configured on an interface.

activation-priority *priority* specifies the order in which interfaces are used for providing call bandwidth. The interface with the highest activation priority value is used as the primary link for providing call bandwidth. If the primary link becomes unavailable, the TGM550 switches to the next active interface with the highest activation priority value, and so on. The activation priority value range is from 0 through 255. The default is 50.



NOTE: Dynamic CAC works in conjunction with the Avaya Communication Manager (CM) Call Admission Control: Bandwidth Limitation (CAC-BL) feature. If you configure dynamic CAC on WAN interfaces, you must also configure CAC-BL on Avaya CM. For more information about configuring CAC-BL, see the *Administrator Guide for Avaya Communication Manager*.

Example: Configuring Dynamic CAC

Configure dynamic CAC on a logical interface:

```
[edit]
interfaces {
  t1-4/0/0 {
    unit 0 {
      dynamic-call-admission-control {
        bearer-bandwidth-limit 900 kbps;
        activation-priority 75;
      }
    }
  }
}
```

Disabling a Logical Interface

You can unconfigure a logical interface, effectively disabling that interface, without removing the logical interface configuration statements from the configuration. To do this, include the **disable** statement:

```
disable;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

When an interface is disabled, a route (pointing to the reserved target “REJECT”) with the IP address of the interface and a 32-bit subnet mask is installed in the routing table. See *Routing Protocols*.

Chapter 5

Configuring Protocol Family and Interface Address Properties

This chapter describes the configuration of the interface protocol and address properties:

- Protocol Family Configuration and Interface Address Statements on page 169
- Configuring the Protocol Family on page 172
- Configuring the Interface Address on page 174
- Configuring ICCP for MC-LAG on page 177
- Configuring IPCP Options on page 177
- Configuring LLC2 Options on page 180
- Configuring an Unnumbered Interface on page 185
- Setting the Protocol MTU on page 191
- Disabling the Removal of Address and Control Bytes on page 192
- Disabling the Transmission of Redirect Messages on an Interface on page 192
- Configuring Default, Primary, and Preferred Addresses and Interfaces on page 192
- Applying Policers on page 194
- Applying a Filter to an Interface on page 203
- Configuring Unicast RPF on page 208
- Enabling Source Class and Destination Class Usage on page 214

Protocol Family Configuration and Interface Address Statements

For each logical interface, you must configure one or more protocol families. You can also configure interface address properties. To do this, include the following statements:

```
family family {  
    accounting {  
        destination-class-usage;  
        source-class-usage {  
            direction;  
        }  
    }  
}  
address address {
```

```

        destination address;
    }
    bundle interface-name;
    filter {
        dialer filter-name;
        input filter-name;
        output filter-name;
        group filter-group-number;
    }
    interface-mode (access | trunk);
    ipsec-sa sa-name;
    keep-address-and-control;
    llc2 {
        ack-delay-time time;
        ack-max count;
        idle-time time;
        local-window count;
        max-retry count;
        p-bit-timeout time;
        redundancy-group group-number {
            advertise-interval seconds;
            map {
                local-mac mac-address request mac-address;
            }
            preempt hold-time seconds;
            no-preempt;
            priority priority;
            track {
                dls {
                    peer ip-address priority-cost priority;
                    destination mac-address priority-cost priority;
                }
                interface interface-name priority-cost priority;
            }
        }
    }
    t1-time time;
    t2-time time;
    trej-time time;
}
mtu bytes;
multicast-only;
negotiate-address;
no-redirects;
policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
}
primary;
protocols [inet iso mpls];
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check <fail-filter filter-name>;
sampling {

```

```

    direction;
}
service {
    input {
        service-set service-set-name <service-filter filter-name>;
        post-service-filter filter-name;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
vlan-id number;
vlan-id-list (Interface in Bridge Domain) [number number-number];
unnumbered-address interface-name destination address destination-profile
    profile-name;
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    multipoint-destination address dlcid dlcid-identifier;
    multipoint-destination address {
        epd-threshold cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (disable | seconds);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
                rate burst length);
            queue-length number;
        }
        vci vpi-identifier.vci-identifier;
    }
}
primary;
preferred;
(vrrp-group | vrrp-inet6-group) group-number {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    authentication-type authentication;
    authentication-key key;
    fast-interval milliseconds;
    (preempt | no-preempt) {
        hold-time seconds;
    }
}
priority-number number;
track {
    priority-cost seconds;
    priority-hold-time interface-name {
        interface priority;
        bandwidth-threshold bits-per-second {

```

```

        priority;
    }
}
route ip-address/mask routing-instance instance-name priority-cost cost;
}
virtual-address [ addresses ];
}
}
}

```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For information about interface-specific protocol and address properties, see “Configuring T1 and NxDS0 Interfaces” on page 458.

Configuring the Protocol Family

For each logical interface, you can configure one or more of the following protocols that run on the interface:

- **any**—Protocol-independent family used for Layer 2 packet filtering. This option is not supported on J Series routers.
- **bridge**—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation. You can optionally configure this protocol family for the logical interface on which you configure VPLS.
- **ccc**—Circuit cross-connect (CCC). You can configure this protocol family for the logical interface of CCC physical interfaces. When you use this encapsulation type, you can configure the **ccc** family only.
- **inet**—IP. You must configure this protocol family for the logical interface to support IP protocol traffic, including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), Internet Control Message Protocol (ICMP), and Internet Protocol Control Protocol (IPCP).
- **inet6**—IP version 6 (IPv6). You must configure this protocol family for the logical interface to support IPv6 protocol traffic, including Routing Information Protocol for IPv6 (RIPng), Intermediate System-to-Intermediate System (IS-IS), BGP, and Virtual Router Redundancy Protocol for IPv6 (VRRP). For more information about IPv6, see “IPv6 Overview” on page 174.
- **iso**—International Organization for Standardization (ISO). You must configure this protocol family for the logical interface to support IS-IS traffic.
- **mlfr-uni-nni**—Multilink Frame Relay (MLFR) FRF.16 user-to-network network-to-network (UNI NNI). You must configure this protocol or **mlfr-end-to-end** for the logical interface to support link services and voice services bundling.

- **mlfr-end-to-end**—Multilink Frame Relay end-to-end. You must configure this protocol or multilink Point-to-Point Protocol (MLPPP) for the logical interface to support multilink bundling.
- **mlppp**—MLPPP. You must configure this protocol (or **mlfr-end-to-end**) for the logical interface to support multilink bundling.
- **mpls**—Multiprotocol Label Switching (MPLS). You must configure this protocol family for the logical interface to participate in an MPLS path.
- **tcc**—Translational cross-connect (TCC). You can configure this protocol family for the logical interface of TCC physical interfaces.
- **tnp**—Trivial Network Protocol. This protocol is used to communicate between the Routing Engine and the router's packet forwarding components. The JUNOS Software automatically configures this protocol family on the router's internal interfaces only, as discussed in “Displaying the Internal Ethernet Interface” on page 245.
- **vpls**—M Series and T Series routers support Virtual Private LAN service (VPLS). You can optionally configure this protocol family for the logical interface on which you configure VPLS. VPLS provides an Ethernet-based point-to-multipoint Layer 2 VPN to connect customer edge (CE) routers across an MPLS backbone. When you configure a VPLS encapsulation type, the **family vpls** statement is assumed by default.

MX Series routers support dynamic profiles for VPLS pseudowires, VLAN identifier translation, and automatic bridge domain configuration.

For more information about VPLS, see the *JUNOS VPNs Configuration Guide* and the *JUNOS Feature Guide*.

To configure the logical interface's protocol family, include the **family** statement, specifying the selected family. To configure more than one protocol family on a logical interface, include multiple **family** statements. Following is the minimum configuration:

```
family family {
  mtu size;
  multicast-only;
  no-redirects;
  primary;
  address address {
    destination address;
    broadcast address;
    preferred;
    primary;
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

IPv6 Overview

IP version 4 (IPv4) has been widely deployed and used to network the Internet today. With the rapid growth of the Internet, enhancements to IPv4 are needed to support the influx of new subscribers, Internet-enabled devices, and applications. IPv6 is designed to enable the global expansion of the Internet.

IPv6 builds upon the functionality of IPv4, providing improvements to addressing, configuration and maintenance, and security.

IPv6 is defined in the following documents:

- RFC 2373, *IP Version 6 Addressing Architecture*
- RFC 2460, *Internet Protocol, Version 6 (IPv6)*

IPv4-to-IPv6 Transition

Implementing IPv6 requires a transition mechanism to allow interoperability between IPv6 nodes (both routers and hosts) and IPv4 nodes. The transition mechanism is the key factor in the successful deployment of IPv6. Because millions of IPv4 nodes already exist, upgrading every node to IPv6 at the same time is not feasible.

As a result, transition from IPv4 to IPv6 happens gradually, allowing nodes to be upgraded independently and without disruption to other nodes. While a gradual upgrade occurs, compatibility between IPv6 and IPv4 nodes becomes a requirement. Otherwise, an IPv6 node would not be able to communicate with an IPv4 node.

Transition mechanisms allow IPv6 and IPv4 nodes to coexist together in the same network, and make gradual upgrading possible. The transition mechanism supported by the JUNOS Software is tunneling. Tunnels allow IPv6 packets to be encapsulated into IPv4 headers and sent across an IPv4 infrastructure. For more information about configuring tunnels to support IPv4-to-IPv6 transition, see the *JUNOS Services Interfaces Configuration Guide*.

VRRP Properties

The Virtual Router Redundancy Protocol (VRRP) provides a much faster switchover to a backup router when the default router fails. Using VRRP, a backup router can take over a failed default router within a few seconds. This is done with minimum amount of VRRP traffic and without any interactions with the hosts.

For more information on VRRP properties, see the *JUNOS High Availability Configuration Guide*.

Configuring the Interface Address

You assign an address to an interface by specifying the address when configuring the protocol family. For the `inet` family, configure the interface's IP address. For the `iso` family, configure one or more addresses for the loopback interface. For the `ccc`, `tcc`, `mpls`, `tnp`, and `vpls` families, you never configure an address.

To assign an address to an interface, include the **address** statement:

```
address address {
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    preferred;
    primary;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

In the **address** statement, specify the network address of the interface.

For each address, you can optionally configure one or more of the following:

- Broadcast address for the interface's subnet—Specify this in the **broadcast** statement; this applies only to Ethernet interfaces, such as the management interface **fxp0** or **em0**, the Fast Ethernet interface, and the Gigabit Ethernet interface.
- Address of the remote side of the connection (for point-to-point interfaces only)—Specify this in the **destination** statement.
- Assign PPP properties to the remote end—Specify this in the **destination-profile** statement. You define the profile at the [edit access group-profile *name* ppp] hierarchy level (for point-to-point interfaces only). For more information, see “Configuring IPCP Options” on page 177.
- Whether the router automatically generates the host number portion of interface addresses—The **eui-64** statement applies only to interfaces that carry IPv6 traffic, where the prefix length of the address is 64 bits or less, and the low-order 64 bits of the address are zero. This option does not apply to the loopback interface (**lo0**) because IPv6 addresses configured on the loopback interface must have a 128-bit prefix length.
- Whether this address is the preferred address—Each subnet on an interface has a preferred local address. If you configure more than one address on the same subnet, the preferred local address is chosen by default as the source address when you originate packets to destinations on the subnet. For more information about preferred addresses, see “Configuring Default, Primary, and Preferred Addresses and Interfaces” on page 192.

By default, the preferred address is the lowest numbered address on the subnet. To override the default and explicitly configure the preferred address, include the **preferred** statement when configuring the address.

- Whether this address is the primary address—Each interface has a primary local address. If an interface has more than one address, the primary local address is used by default as the source address when you originate packets out the interface where the destination gives no hint about the subnet (for example, some **ping** commands). For more information about primary addresses, see “Configuring Default, Primary, and Preferred Addresses and Interfaces” on page 192.

By default, the primary address on an interface is the lowest numbered non-127 preferred address on the interface. To override the default and explicitly configure the preferred address, include the **primary** statement when configuring the address.

- Configuring an Interface IPv4 Address on page 176
- Configuring the Interface IPv6 Address on page 176

Configuring an Interface IPv4 Address

You can configure router interfaces with a 32-bit IPv4 address and optionally with a destination prefix, sometimes called a subnet mask. An IPv4 address utilizes a 4-octet dotted decimal address syntax (for example, 192.16.1.1). An IPv4 address with destination prefix utilizes a 4-octet dotted decimal address syntax appended with a destination prefix (for example, 192.16.1.1/30).

To configure an IPv4 address on JUNOS routers, use the **edit interface** *interface-name* unit *number* family <inet> address *a.b.c.d/nn* statement at the [edit interfaces] hierarchy level.



NOTE: Juniper Networks routers support /31 destination prefixes when used in point-to-point Ethernet configurations; however, it is not supported by many other devices, such as hosts, hubs, or routers. You must determine if the peer system also supports /31 destination prefixes before configuration.

Configuring the Interface IPv6 Address

You represent IPv6 addresses in hexadecimal notation using a colon-separated list of 16-bit values.

You assign a 128-bit IPv6 address to an interface by including the **address** statement:

```
address aaaa:bbb:...:zzzz/nn;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

The double colon (::) represents all bits set to 0, as shown in the following example:

```
interfaces fe-0/0/1 {
  unit 0 {
```



```

        family inet6 {
            address fec0:1:1:1::2/64;
        }
    }
}

```



NOTE: You must manually configure the router advertisement and advertise the default prefix for autoconfiguration to work on a specific interface.

Configuring ICCP for MC-LAG

For Multi-Chassis Link Aggregation (MC-LAG), you must configure Internet Chassis Control (ICCP) to exchange information between two MC-LAG peers.

To enable ICCP, include the ICCP statement at the [edit protocol] hierarchy level:

```

[edit protocols ICCP traceoptions]
ICCP {
    traceoptions;
    local-ip-address ip address;
    session-establishment-hold-time value;
    authentication-key string;
    peer ip-address{
        local-ip-address ip address;
        session-establishment-hold-time value;
        authentication-key string;
        redundancy-group-id-list redundancy-group-id-list;
        liveness-detection;
    }
}

```

To configure ICCP for MC-LAG, identify the [local-ip-address](#) that acts as the source address. This could be a specified address or interface address.

Session-establishment-hold-time determines whether a chassis takes over as the master at the ICCP session.

The **authentication-key** is provided by TCP Message Digest 5 (md5) option for an ICCP TCP session. The **redundancy-group-id-list** specifies the redundancy groups between ICCP peers and **liveness-detection** configures Bidirectional Forwarding Detection protocol (BFD) options.

Configuring IPCP Options

For interfaces with PPP encapsulation, you can configure IPCP to negotiate IP address assignments and to pass network-related information such as Windows Name Service (WINS) and Domain Name System (DNS) servers, as defined in RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*.



NOTE: The JUNOS Software does not request name servers from the remote end; the software does, however, send name servers to the remote end if requested.

On the logical interface, the following PPP encapsulation types are supported:

- atm-mlppp-llc
- atm-ppp-llc
- atm-ppp-vc-mux
- multilink-ppp

When you enable a PPP interface, you can configure an IP address, enable the interface to negotiate an IP address assignment from the remote end, or allow the interface to be unnumbered. You can also assign a destination profile to the remote end. The destination profile includes PPP properties, such as primary and secondary DNS and NetBIOS Name Servers (NBNSS). These options are described in the following sections:

- Configuring an IP Address for an Interface on page 178
- Negotiating an IP Address Assignment from the Remote End on page 178
- Configuring an Interface to Be Unnumbered on page 179
- Assigning a Destination Profile to the Remote End on page 179

Configuring an IP Address for an Interface

You can configure an IP address for the interface by including the **address** statement in the configuration. For more information, see “Configuring the Interface Address” on page 174.

If you include the **address** statement in the configuration, you cannot include the **negotiate-address** or **unnumbered-address** statement in the configuration.

When you include the **address** statement in the interface configuration, you can assign PPP properties to the remote end, as shown in “Assigning a Destination Profile to the Remote End” on page 179.

Negotiating an IP Address Assignment from the Remote End

To enable the interface to obtain an IP address from the remote end, include the **negotiate-address** statement at the [edit interfaces *interface-name* unit *logical-unit-number* family inet] hierarchy level:

```
negotiate-address;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

If you include the `negotiate-address` statement in the configuration, you cannot include the `address` or `unnumbered-address` statement in the configuration.

Configuring an Interface to Be Unnumbered

To configure an interface to be unnumbered, include the `unnumbered-address` and `destination` statements in the configuration:

```
unnumbered-address interface-name destination address;
```

The `unnumbered-address` statement enables the local address to be derived from the specified interface. The interface name must include a logical unit number and must have a configured address (see “Configuring the Interface Address” on page 174). Specify the IP address of the remote interface with the `destination` statement.

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet]

If you include the `unnumbered-address` statement in the configuration, you cannot include the `address` or `negotiate-address` statement in the interface configuration.

When you include the `unnumbered-address` statement in the interface configuration, you can assign PPP properties to the remote end, as shown in “Assigning a Destination Profile to the Remote End” on page 179.

Assigning a Destination Profile to the Remote End

When you include the `address` or `unnumbered-address` statement in the interface configuration, you can assign PPP properties to the remote end. To do this, include the `destination-profile` statement:

```
destination-profile name;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]
- [edit interfaces *interface-name* unit *logical-unit-number* family inet unnumbered-address *interface-name*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet unnumbered-address *interface-name*]

The profile name is a PPP group profile. You define the profile by including the following statements at the [edit access group-profile *name* ppp] hierarchy level:

```
[edit access group-profile name ppp]
```

```

framed-pool pool-id;
interface-id interface-id;
primary-dns primary-dns;
primary-wins primary-win-server;
secondary-dns secondary-dns;
secondary-wins secondary-wins;

```

For more information about PPP group profiles, see the *JUNOS System Basics Configuration Guide*.

Configuring LLC2 Options

Logical link control 2 (LLC2) options can be configured for data link switching (DLSw) protocol support on J Series routers. DLSw allows you to tunnel System Network Architecture (SNA) and NetBIOS traffic over an IP network.

DLSw enables SNA clients to communicate to SNA applications on a mainframe through an IP network. After a connection is established, a DLSw circuit can be created for transporting SNA traffic.

The IP network between an SNA client and an SNA application becomes transparent with DLSw. DLSw transports any SNA traffic. DLSw has an LLC session on one SNA device and recreates it (almost identically) on the other SNA device, making the devices operate as if they were directly connected. DLSw is configured on peer IP routers and transports everything between the peers using Switch-to-Switch Protocol (SSP).

For information about configuring DLSw, see the *JUNOS Services Interfaces Configuration Guide* and the *J-series Services Router Advanced WAN Access Configuration Guide*.

For more information, see the following sections:

- Configuring LLC2 Properties on page 180
- Configuring DLSw Ethernet Redundancy Using LLC2 Properties on page 181
- Example: Configuring LLC Options on an Interface on page 183
- Example: Configuring DLSw Ethernet Redundancy on page 184

Configuring LLC2 Properties

For basic DLSw configuration, include the `llc2` statement at the [edit interfaces *interface-name* unit *logical-unit-number* family] hierarchy level. All other LLC2 statements that follow are optional and should be used only if recommended by support or a services professional to solve specific problems or for specific network designs.

To configure logical link control properties, include the `llc2` statement:

```

llc2 {
    ack-delay-time time;
    ack-max count;
    idle-time time;
    local-window count;

```

```

max-retry count;
p-bit-timeout time;
t1-time time;
t2-time time;
trej-time time;
}

```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family]

You can configure the following LLC options:

- **ack-delay-time**—The maximum time allowed for incoming Information-frames (I-frames) to remain unacknowledged. Specify the number of milliseconds from 1 through 60000. The default value is 100 milliseconds.
- **ack-max**—The maximum number of I-frames received before acknowledgment is sent. Specify the number of I-frames from 1 through 127. The default value is three I-frames.
- **idle-time**—The number of seconds that a TCP connection between DLSw peers will stay up without any circuit using the connection. Specify the number of seconds from 1 through 60000. The default value is 10 seconds.
- **local-window**—The maximum number of I-frames to send before waiting for acknowledgment. Specify the number of I-frames from 1 through 127. The default value is 7 I-frames.
- **max-retry**—The number of retries the router should attempt when waiting for a response. Specify the number of I-frames from 1 through 127. The default value is 10 I-frames.
- **p-bit-timeout**—The length of time the router waits for response to a poll bit. Specify the number of milliseconds from 1 through 60000. The default value is 3000 milliseconds.
- **t1-time**—The length of time the router waits for an acknowledgment of transmitted frames. Specify the number of milliseconds from 1 through 60000. The default value is 1000 milliseconds.
- **t2-time**—The length of time the router withholds the I-frame response. Specify the number of milliseconds from 1 through 60000. The default value is 100 milliseconds.
- **trej-time**—Q.391-specific timer for T310, in seconds. Specify the number of milliseconds from 1 through 60000. The default value is 3000 milliseconds.

Configuring DLSw Ethernet Redundancy Using LLC2 Properties

DLSw is a means of tunneling SNA and NetBIOS traffic over IP networks. To achieve fault tolerance and load sharing, you can configure Ethernet redundancy and deploy multiple DLSw routers on the same LAN segment. These redundant routers provide alternate paths to the destinations and avoid a single point of failure.

When you configure DLSw Ethernet redundancy on a LAN segment, a master router is selected from a group of DLSw neighbors. The master router establishes the circuits.

To configure DLSw Ethernet redundancy, include the **redundancy-group** statement and define redundancy group options:

```

llc2 {
  redundancy-group group-number {
    advertise-interval seconds;
    map {
      local-mac mac-address remote-mac mac-address;
      preempt hold-time seconds;
      no-preempt;
      priority priority;
      track {
        dls {
          destination mac-address priority-cost priority;
          peer ip-address priority-cost priority;
        }
        interface interface-name priority-cost priority;
      }
    }
  }
}

```

You can include these statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family]

You can configure the following redundancy options:

- **redundancy-group group-number**—The group to which this router belongs. Specify the group number, in the range from 0 through 255.
- **advertise-interval**—The advertisement interval of DLSw peers on the network. All routers in the redundancy group must use the same advertisement interval. Specify the number of seconds, from 1 through 255. The default is 1 second.
- **map**—Map a local peer MAC address to a remote peer MAC address.
 - **local-mac**—The local MAC address to be mapped to a remote destination MAC address.
 - **mac-address**—The MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nn:nn:nn:nn:nn:nn* or *nnnn .nnnn.nnnn*. For example, 0011.2233.4455 or 00:11:22:33:44:55.
 - **remote-mac**—The remote destination MAC address to be mapped to a local MAC address.
- **preempt hold-time seconds**— Configure the time to wait before a higher-priority backup router preempts the master router. Specify the number of seconds, from 0 through 3600. DLSw preemption is 0 by default.

- **no-preempt**—Prohibit the preemption of the master router.
- **priority *priority***—The router's priority for becoming the master router. The router with the highest priority within the redundancy group becomes the master. A larger value indicates a higher priority for being elected. Specify the priority from 1 through 255. The default is 100 (for backup routers).
- **track**—Enable the following tracking options for the remote peer and the destination peer:
 - **dls**—DLSw protocol.
 - **destination *mac-address* priority-cost *priority***—The local MAC address and the priority. Specify the MAC address as six hexadecimal bytes in one of the following formats: *nn:nn:nn:nn:nn:nn* or *nnnn .nnnn.nnnn*. For example, 0011.2233.4455 or 00:11:22:33:44:55. The priority cost is the value subtracted from the priority value when remote peer connectivity is lost. Specify a value from 1 through 254.
 - **peer *ip-address* priority-cost *priority***—The IP address of the remote peer. The priority cost is the value subtracted from the priority value when remote peer connectivity is lost. Specify a value from 1 through 254.
 - **interface *interface-name***—The interface name. Include the logical portion of the name, which corresponds to the logical unit number.

Example: Configuring LLC Options on an Interface


Configure LLC options on an unnumbered interface:

```
[edit]
interfaces {
  fe-0/0/0 {
    unit 0 {
      family inet;
      address 10.10.10.2/24;
    }
    family llc2 {
      ack-delay-time 3000;
      ack-max 10;
      idle-time 102;
      local-window 15;
      max-retry 20;
      p-bit-timeout 100;
      t1-time 101;
      t2-time 101;
      max-retry 5;
      trej-time 4000;
    }
  }
}
```

Example: Configuring DLSw Ethernet Redundancy

In Figure 7 on page 184, the local hosts share the same destination MAC address of 00:00:5E:00:01:01 and send DLSw traffic to the remote host with a MAC address of 00:02:00:00:00:01. Router 1 and Router 2 are configured for DLSw redundancy and map the local destination MAC address to the remote MAC address. Router 1 is also the designated master. If Router 1 becomes unavailable, Router 2, the backup router, takes over as the master router.

Figure 7: DLSw Ethernet Redundancy Topology



To configure DLSw Ethernet redundancy, do the following:

Configuration on Router 1 Configure the redundancy group, redundancy group options, and the priority cost of each redundancy group option:

```
[edit]
interfaces {
  fe-0/0/0 {
    unit 0 {
      family llc2 {
        redundancy-group 1 {
          advertise-interval 1;
          map {
            local-mac 00:00:5e:00:01:01 remote-mac 00:02:00:00:00:01;
            preempt hold-time 20;
            priority 200;
            track {
              dls {
                destination 00:02:00:00:00:01 priority-cost 50;
                peer 10.10.10.10 priority-cost 25;
              }
              interface e1-0/0/2.0 priority-cost 40;
            }
          }
        }
      }
    }
  }
}
```


Configuration on Router 2 Configure the redundancy group, redundancy group options, and the priority cost of each redundancy group option:

```
[edit]
interfaces {
  fe-0/0/1 {
    unit 0 {
      family llc2 {
        redundancy-group 1 {
          map {
            local-mac 00:00:5e:00:01:01 remote-mac 00:02:00:00:00:01;
            priority-cost 190;
            track {
              dls {
                destination 00:02:00:00:00:01 priority-cost 50;
                peer 10.10.10.10 priority-cost 25;
              }
              interface e1-0/0/2.0 priority-cost 40;
            }
          }
        }
      }
    }
  }
}
```

Configuring an Unnumbered Interface

When you need to conserve IP addresses, you can configure unnumbered interfaces. Setting up an unnumbered interface enables IP processing on the interface without assigning an explicit IP address to the interface. For IPv6, in which conserving addresses is not a major concern, you can unnumbered interfaces to share the same subnet across multiple interfaces. IPv6 unnumbered interfaces are only supported on Ethernet interfaces. The statements you use to configure an unnumbered interface depend on the type of interface you are configuring: a point-to-point interface or an Ethernet interface:

- Configuring an Unnumbered Point-to-Point Interface on page 185
- Configuring an Unnumbered Ethernet or Demux Interface on page 186

Configuring an Unnumbered Point-to-Point Interface

To configure an unnumbered point-to-point interface, configure the protocol family, but do not include the **address** statement:

```
family family;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]



NOTE: For interfaces with PPP encapsulation, you can configure an unnumbered interface by including the `unnumbered-interface` statement in the configuration. For more information, see “Configuring IPCP Options” on page 177.

When configuring unnumbered interfaces, you must ensure that a source address is configured on some interface in the router. This address is the default address. We recommend that you do this by assigning an address to the loopback interface (lo0), as described in “Configuring the Loopback Interface” on page 257. If you configure an address (other than a martian) on the lo0 interface, that address is always the default address, which is preferable because the loopback interface is independent of any physical interfaces and therefore is always accessible.

Example: Configuring an Unnumbered Point-to-Point Interface

Configure an unnumbered point-to-point interface:

```
[edit]
interfaces {
  so-6/1/0 {
    unit 0 {
      family inet;
      family iso;
    }
  }
}
```

Configuring an Unnumbered Ethernet or Demux Interface

To configure an unnumbered Ethernet or demultiplexing interface, include the `unnumbered-address` statement in the configuration:

```
unnumbered-address interface-name;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

For dynamic profiles, include the `unnumbered-address` statement at the following hierarchy levels:

- [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family*]

The `unnumbered-address` statement currently supports configuration of unnumbered demux interfaces only for the IPv4 address family. You can configure unnumbered Ethernet interfaces for both IPv4 and IPv6 address families.

The interface that you configure to be unnumbered *borrow*s an assigned IP address from another interface, and is referred to as the *borrower interface*. The interface from which the IP address is borrowed is referred to as the *donor interface*. In the `unnumbered-address` statement, *interface-name* specifies the donor interface. For an unnumbered Ethernet interface, the donor interface can be an Ethernet, ATM, SONET, or loopback interface that has a logical unit number and configured IP address and is not itself an unnumbered interface. For an unnumbered IP demultiplexing interface, the donor interface can be an Ethernet or loopback interface that has a logical unit number and configured IP address and is not itself an unnumbered interface. In addition, for either Ethernet or demux, the donor interface and the borrower interface must be members of the same routing instance and the same logical system.

When you configure an unnumbered Ethernet or demux interface, the IP address of the donor interface becomes the source address in packets generated by the unnumbered interface.

You can configure a host route that points to an unnumbered Ethernet or demux interface. For information about host routes, see the *JUNOS MPLS Applications Configuration Guide*.

For more information, see the following sections:

- Configuring a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces on page 187
- Configuring Static Routes on Unnumbered Ethernet Interfaces on page 188
- Restrictions for Configuring Unnumbered Ethernet Interfaces on page 189
- Example: Configuring an Unnumbered Ethernet Interface on page 189
- Example: Configuring the Preferred Source Address for an Unnumbered Ethernet Interface on page 190
- Example: Configuring an Unnumbered Ethernet Interface as the Next Hop for a Static Route on page 190

For additional information about dynamic-profiles, see Dynamic Profiles Overview.

Configuring a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces

When a loopback interface with multiple secondary IP addresses is configured as the donor interface for an unnumbered Ethernet or demux interface, you can optionally specify any one of the loopback interface's secondary addresses as the preferred source address for the unnumbered Ethernet or demux interface. This feature enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet or demux interfaces in your network.

To configure a secondary address on a loopback donor interface as the preferred source address for an unnumbered Ethernet or demux interface, include the `preferred-source-address` option in the `unnumbered-address` statement:

```
unnumbered-address interface-name <preferred-source-address address>;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit dynamic-profiles *profile-name* interfaces demux0 unit *logical-unit-number* family *family*]

The following considerations apply when you configure a preferred source address on an unnumbered Ethernet or demux interface:

- The **unnumbered-address** statement currently supports the configuration of a preferred source address only for the IPv4 address family for demux interfaces, and for IPv4 and IPv6 address families for Ethernet interfaces.
- If you do not specify the preferred source address, the router uses the default primary IP address of the donor interface.
- You cannot delete an address on a donor loopback interface while it is being used as the preferred source address for an unnumbered Ethernet or demux interface.
- The router uses the preferred source address, if configured for an unnumbered Ethernet or demux interface, in ARP requests and replies. ARP requests must match the preferred source address.

For a configuration example that illustrates this feature, see “Example: Configuring the Preferred Source Address for an Unnumbered Ethernet Interface” on page 190.

To display the preferred source address for an unnumbered Ethernet or demux interface, use the **show interfaces** operational mode command. For information about using this command, see the *JUNOS Interfaces Command Reference*.

Configuring Static Routes on Unnumbered Ethernet Interfaces

You can configure static routes on an unnumbered Ethernet interface. To do so, you use the **qualified-next-hop** statement to specify the unnumbered Ethernet interface as the next-hop interface for a configured static route. This feature enables you to specify independent preferences and metrics for static routes on a next-hop basis.

For a configuration example that illustrates this feature, see “Example: Configuring an Unnumbered Ethernet Interface as the Next Hop for a Static Route” on page 190.

For information about how to specify an independent preference for a static route, see the *JUNOS Routing Protocols Configuration Guide*.

Restrictions for Configuring Unnumbered Ethernet Interfaces

The following restrictions apply when you configure unnumbered Ethernet interfaces:

- The `unnumbered-address` statement currently supports the configuration of unnumbered Ethernet interfaces for IPv4 and IPv6 address families.
- You cannot assign an IP address to an Ethernet interface that is already configured as an unnumbered interface.
- The donor interface for an unnumbered Ethernet interface must have one or more configured IP addresses.
- The donor interface for an unnumbered Ethernet interfaced cannot be configured as unnumbered.
- An unnumbered Ethernet interface does not support configuration of the following `address` statement options: `arp`, `broadcast`, `primary`, `preferred`, and `vrrp-group`. For information about these options, see “Configuring the Interface Address” on page 174.
- Running IGMP and PIM are supported only on unnumbered Ethernet interfaces that directly face the host and have no downstream PIM neighbors. IGMP and PIM are not supported on unnumbered Ethernet interfaces that act as upstream interfaces in a PIM topology.
- Running OSPF and IS-IS on unnumbered Ethernet interfaces is not supported.
-

Example: Configuring an Unnumbered Ethernet Interface

In this example, `ge-1/0/0` is the unnumbered interface and `ge-0/0/0` is the donor interface from which `ge-1/0/0` “borrows” an IP address.

```

interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        address 4.4.4.1/24;
      }
    }
  }
  ge-1/0/0 {
    unit 0 {
      family inet {
        unnumbered-address ge-0/0/0.0;
      }
    }
  }
}

```

Example: Configuring the Preferred Source Address for an Unnumbered Ethernet Interface

In this example, loopback interface `lo0` is the donor interface from which unnumbered Ethernet interface `ge-4/0/0` “borrows” an IP address. The example also configures one of the loopback interface’s secondary addresses, `3.3.3.1`, as the preferred source address for the unnumbered Ethernet interface.

```

interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 2.2.2.1/32;
        address 3.3.3.1/32;
      }
    }
  }
}
interfaces {
  ge-4/0/0 {
    unit 0 {
      family inet {
        unnumbered-address lo0.0 preferred-source-address 3.3.3.1;
      }
    }
  }
}

```

Example: Configuring an Unnumbered Ethernet Interface as the Next Hop for a Static Route

In this example, `ge-0/0/0` is the unnumbered interface and a loopback interface, `lo0`, is the donor interface from which `ge-0/0/0` “borrows” an IP address. The example also configures a static route to `7.7.7.1/32` with a next hop through unnumbered interface `ge-0/0/0.0`.

```

interfaces {
  lo0 {
    unit 0 {
      family inet {
        address 5.5.5.1/32;
        address 6.6.6.1/32;
      }
    }
  }
}
interfaces {
  ge-0/0/0 {
    unit 0 {
      family inet {
        unnumbered-address lo0.0;
      }
    }
  }
}

```

```

    }
  }
  routing-options {
    static {
      route 7.7.7.1/32 {
        qualified next-hop ge-0/0/0.0;
      }
    }
  }
}

```

Setting the Protocol MTU

When you initially configure an interface, the protocol maximum transmission unit (MTU) is calculated automatically. If you subsequently change the media MTU, the protocol MTU on existing address families automatically changes.

For a list of default protocol MTU values, see “Configuring the Media MTU” on page 98.

To modify the MTU for a particular protocol family, include the `mtu` statement:

```
mtu bytes;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

If you increase the size of the protocol MTU, you must ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. For a list of encapsulation overhead values, see Table 18 on page 104. If you reduce the media MTU size, but there are already one or more address families configured and active on the interface, you must also reduce the protocol MTU size. (You configure the media MTU by including the `mtu` statement at the [edit interfaces *interface-name*] hierarchy level, as discussed in “Configuring the Media MTU” on page 98.)



NOTE: Changing the media MTU or protocol MTU causes an interface to be deleted and added again.

The maximum number of data-link connection identifiers (DLCIs) is determined by the MTU on the interface. If you have keepalives enabled, the maximum number of DLCIs is 1000, with the MTU set to 5012.

The actual frames transmitted also contain cyclic redundancy check (CRC) bits, which are not part of the MTU. For example, the default protocol MTU for a Gigabit Ethernet interface is 1500 bytes, but the largest possible frame size is actually 1504 bytes; you need to consider the extra bits in calculations of MTUs for interoperability.

Disabling the Removal of Address and Control Bytes

For Point-to-Point Protocol (PPP) CCC-encapsulated interfaces, the address and control bytes are removed by default before the packet is encapsulated into a tunnel.

You can disable the removal of address and control bytes. To do this, include the `keep-address-and-control` statement:

```
keep-address-and-control;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *ccc*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *ccc*]

Disabling the Transmission of Redirect Messages on an Interface

By default, the interface sends protocol redirect messages. To disable the sending of these messages on an interface, include the `no-redirects` statement:

```
no-redirects;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

To disable the sending of protocol redirect messages for the entire router, include the `no-redirects` statement at the [edit `system`] hierarchy level.

Configuring Default, Primary, and Preferred Addresses and Interfaces

The router has a default address and a primary interface, and interfaces have primary and preferred addresses.

The *default address* of the router is used as the source address on unnumbered interfaces. The routing protocol process tries to pick the default address as the router ID, which is used by protocols, including OSPF and internal BGP (IBGP).

The *primary interface* for the router is the interface that packets go out when no interface name is specified and when the destination address does not imply a particular outgoing interface.

An interface's *primary address* is used by default as the local address for broadcast and multicast packets sourced locally and sent out the interface. An interface's *preferred address* is the default local address used for packets sourced by the local router to destinations on the subnet.

The default address of the router is chosen using the following sequence:

1. The primary address on the loopback interface `lo0` that is not `127.0.0.1` is used.
2. The primary address on the primary interface is used.

To configure these addresses and interfaces, you can do the following:

- Configuring the Primary Interface for the router on page 193
- Configuring the Primary Address for an Interface on page 193
- Configuring the Preferred Address for an Interface on page 194

Configuring the Primary Interface for the router

The *primary interface* for the router has the following characteristics:

- It is the interface that packets go out when you type a command such as `ping 255.255.255.255`—that is, a command that does not include an interface name (there is no interface *type-0/0/0.0* qualifier) and where the destination address does not imply any particular outgoing interface.
- It is the interface on which multicast applications running locally on the router, such as Session Announcement Protocol (SAP), do group joins by default.
- It is the interface from which the default local address is derived for packets sourced out an unnumbered interface if there are no non-127 addresses configured on the loopback interface, `lo0`.

By default, the multicast-capable interface with the lowest-index address is chosen as the primary interface. If there is no such interface, the point-to-point interface with the lowest index address is chosen. Otherwise, any interface with an address could be picked. In practice, this means that, on the router, the `fxp0` or `em0` interface is picked by default.

To configure a different interface to be the primary interface, include the `primary` statement:

```
primary;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

Configuring the Primary Address for an Interface

The *primary address* on an interface is the address that is used by default as the local address for broadcast and multicast packets sourced locally and sent out the interface. For example, the local address in the packets sent by a `ping interface so-0/0/0.0 255.255.255.255` command is the primary address on interface `so-0/0/0.0`. The primary address flag also can be useful for selecting the local address used for packets sent out unnumbered interfaces when multiple non-127 addresses are configured

on the loopback interface, `lo0`. By default, the primary address on an interface is selected as the numerically lowest local address configured on the interface.

To set a different primary address, include the `primary` statement:

```
primary;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family* address *address*]

Configuring the Preferred Address for an Interface

The *preferred address* on an interface is the default local address used for packets sourced by the local router to destinations on the subnet. By default, the numerically lowest local address is chosen. For example, if the addresses `172.16.1.1/12`, `172.16.1.2/12`, and `172.16.1.3/12` are configured on the same interface, the preferred address on the subnet (by default, `172.16.1.1`) would be used as a local address when you issue a `ping 172.16.1.5` command.

To set a different preferred address for the subnet, include the `preferred` statement:

```
preferred;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family* address *address*]

Applying Policers

Policers allow you to perform simple traffic policing on specific interfaces or Layer 2 virtual private networks (VPNs) without configuring a firewall filter. To apply policers, include the `policer` statement:

```
policer {
  arp policer-template-name;
  input policer-template-name;
  output policer-template-name;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

In the `family` statement, the protocol family can be `ccc`, `inet`, `inet6`, `mpls`, `tcc`, or `vpls`.

In the **arp** statement, list the name of one policer template to be evaluated when Address Resolution Protocol (ARP) packets are received on the interface. By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the **family inet** statement. If you want more stringent or lenient policing of ARP packets, you can configure an interface-specific policer and apply it to the interface. You configure an ARP policer just as you would configure any other policer, at the **[edit firewall policer]** hierarchy level. If you apply this policer to an interface, the default ARP packet policer is overridden. If you delete this policer, the default policer takes effect again.

In the **input** statement, list the name of one policer template to be evaluated when packets are received on the interface.

In the **output** statement, list the name of one policer template to be evaluated when packets are transmitted on the interface.



NOTE: To use policing on a CCC or TCC interface, you must configure the CCC or TCC protocol family.

You can configure a different policer on each protocol family on an interface, with one input policer and one output policer for each family. When you apply policers, you can configure the family **ccc**, **inet**, **inet6**, **mpls**, **tcc**, or **vpls** only, and one ARP policer for the family **inet** protocol only. Each time a policer is referenced, a separate copy of the policer is installed on the packet forwarding components for that interface.

If you apply both policers and firewall filters to an interface, input policers are evaluated before input firewall filters, and output policers are evaluated after output firewall filters.

If you apply the policer to the interface **lo0**, it is applied to packets received or transmitted by the Routing Engine.

On T Series, M120, and M320 platforms, if the interfaces are on the same FPC, the filters or policers do not act on the sum of traffic entering and exiting the interfaces.

For more information about policers, see the *JUNOS Policy Framework Configuration Guide*.

Applying Aggregate Policers

By default, if you apply a policer to multiple protocol families on the same logical interface, the policer restricts traffic for each protocol family individually. For example, a policer with a 50 Mbps bandwidth limit applied to both IPv4 and IPv6 traffic would allow the interface to accept 50 Mbps of IPv4 traffic and 50 Mbps of IPv6 traffic. If you apply an aggregate policer, the policer would allow the interface to receive only 50 Mbps of IPv4 and IPv6 traffic combined.

To configure an aggregate policer, include the **logical-interface-policer** statement at the **[edit firewall policer *policer-template-name*]** hierarchy level:

```
[edit firewall policer policer-template-name]
```

```
logical-interface-policer;
```

For the policer to be treated as an aggregate, you must apply it to multiple protocol families on a single logical interface by including the **policer** statement:

```
policer {
  arp policer-template-name;
  input policer-template-name;
  output policer-template-name;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

In the family statement, the protocol family can be **ccc**, **inet**, **inet6**, **mpls**, **tcc**, or **vpls**.

The protocol families on which you do not apply the policer are not affected by the policer. For example, if you configure a single logical interface to accept MPLS, IPv4, and IPv6 traffic and you apply the logical interface policer **policer1** to only the IPv4 and IPv6 protocol families, MPLS traffic is not subject to the constraints of **policer1**.

If you apply **policer1** to a different logical interface, there are two instances of the policer. This means the JUNOS Software polices traffic on separate logical interfaces separately, not as an aggregate, even if the same logical-interface policer is applied to multiple logical interfaces on the same physical interface port.



NOTE: Logical interface policers are not supported for filter policers. In other words, you cannot include the **logical-interface-policer** statement at the [edit firewall filter *name* term *name* then policer] hierarchy level.

Example: Applying Aggregate Policers

Configure two logical interface policers: **aggregate_police1** and **aggregate_police2**. Apply **aggregate_police1** to IPv4 and IPv6 traffic received on logical interface **fe-0/0/0.0**. Apply **aggregate_police2** to CCC and MPLS traffic received on logical interface **fe-0/0/0.0**. This configuration causes the software to create only one instance of **aggregate_police1** and one instance of **aggregate_police2**.

Apply **aggregate_police1** to IPv4 and IPv6 traffic received on another logical interface **fe-0/0/0.1**. This configuration causes the software to create a new instance of **aggregate_police1**, one that applies to unit 0 and another that applies to unit 1.

```
[edit firewall]
policer aggregate_police1 {
  logical-interface-policer;
  if-exceeding {
    bandwidth-limit 100m;
    burst-size-limit 500k;
```

```

    }
    then {
        discard;
    }
}
policer aggregate_police2 {
    logical-interface-policer;
    if-exceeding {
        bandwidth-limit 10m;
        burst-size-limit 200k;
    }
    then {
        discard;
    }
}
[edit interfaces fe-0/0/0]
unit 0 {
    family inet {
        policer {
            input aggregate_police1;
        }
    }
    family inet6 {
        policer {
            input aggregate_police1;
        }
    }
    family ccc {
        policer {
            input aggregate_police2;
        }
    }
    family mpls {
        policer {
            input aggregate_police2;
        }
    }
}
unit 1 {
    family inet {
        policer {
            input aggregate_police1;
        }
    }
    family inet6 {
        policer {
            input aggregate_police1;
        }
    }
}
}

```

Applying Hierarchical Policers on Enhanced Intelligent Queuing PICs

M40e, M120, and M320 edge routers and T Series core routers with Enhanced Intelligent Queuing (IQE) PICs support hierarchical policers in the ingress direction

and allow you to apply a hierarchical policer for the premium and aggregate (premium plus normal) traffic levels to an interface. Hierarchical policers provide cross-functionality between the configured physical interface and the Packet Forwarding Engine.

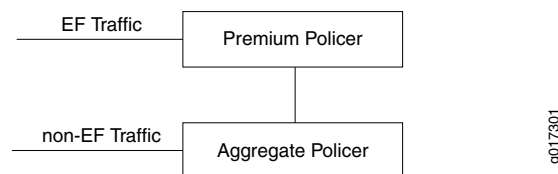
Before you begin, there are some general restrictions that apply to hierarchical policers:

- Only one type of policer can be configured for a logical or physical interface. For example, a hierarchical policer and a regular policer in the same direction for the same logical interface is not allowed.
- The chaining of the policers—that is, applying policers to both a port and the logical interfaces of that port—is not allowed.
- There is a limit of 64 policers per interface in case there is no BA classification, providing a single policer per DLCI.
- Only one kind of policer can be applied on a physical or logical interface.
- The policer should be independent of BA classification. Without BA classification, all traffic on an interface will be treated either as EF or non-EF, based on the configuration. With BA classification, an interface can support up to 64 policers. Again, the interface here may be a physical interface or logical interface (for example, DLCI).
- With BA classification, the miscellaneous traffic (the traffic *not* matching with any of the BA classification DSCP/EXP bits) will be policed as non-EF traffic. No separate policers will be installed for this traffic.

Hierarchical Policier Overview

Hierarchical policing uses two token buckets, one for aggregate (non-EF) traffic and one for premium (EF) traffic. Which traffic is EF and which is non-EF is determined by the class-of-service configuration. Logically, hierarchical policing is achieved by chaining two policers.

Figure 8: Hierarchical Policier



In the example in Figure 8 on page 198, EF traffic is policed by Premium Policier and non EF traffic is policed by Aggregate Policier. What that means is, for EF traffic the out-of-spec action will be the one that is configured for Premium Policier, but the in-spec EF traffic will still consume the tokens from the Aggregate Policier.

But EF traffic will never be submitted to the out-of-spec action of the Aggregate Policier. Also, if the out-of-spec action of the Premium Policier is not set to Discard, those out-of-spec packets will not consume the tokens from the Aggregate Policier. Aggregate Policier only polices the non-EF traffic. As you can see, the Aggregate

Policer token bucket can go negative, if all the tokens are consumed by the non-EF traffic and then you get bursts of EF traffic. But that will be for a very short time, and over a period of time it will average out. For example:

- *Premium Policer*: Bandwidth 2 Mbps, OOS Action: Discard
- *Aggregate Policer*: Bandwidth 10 Mbps, OOS Action: Discard

In the above case, EF traffic is guaranteed 2 Mbps and the non-EF traffic will get from 8 Mbps to 10 Mbps, depending on the input rate of the EF traffic.

Hierarchical Policing Characteristics

Hierarchical token buckets

- Ingress traffic is first classified into EF and non-EF traffic prior to applying a policer:
 - Classification is performed by Q-tree lookup
- Channel number selects a shared token bucket policer:
 - Dual token bucket policer is divided into two single bucket policers:
 - Policer1—EF traffic
 - Policer2—non-EF traffic
- Shared token bucket is used to police the traffic as follows:
 - Policer1 is set to EF rate (for example, 2 Mbps)
 - Policer2 is set to aggregate interface policed rate (for example, 10 Mbps).
 - EF traffic gets applied to Policer1.
 - If traffic is in-spec it is allowed to pass and decrement from both Policer1 and Policer2.
 - If traffic is out-of-spec it can be discarded or marked with a new FC or loss priority. Policer2 will not do anything with out-of-spec EF traffic.
 - Non-EF traffic gets applied only to Policer2.
 - If traffic is in-spec it is allowed to pass through and decremented Policer2.
 - If traffic is out-of-spec it is discarded or marked with a new FC or set with a new drop priority.
- Rate-limit the port speed to a desired rate at Layer 2
- Rate-limit the EF traffic
- Rate-limit the non-EF traffic
- Policing drops counted per color

Configuring Hierarchical Policers

To configure a hierarchical policer, apply the **policing-priority** statement to the proper forwarding class and configure a hierarchical policer for the aggregate and premium level. For more information about class of service, see the *JUNOS Class of Service Configuration Guide*.



NOTE: Hierarchical policers can only be configured on SONET physical interfaces hosted on an IQE PIC. Only aggregate and premium levels are supported.

CoS Configuration of Forwarding Classes for Hierarchical Policers

```
[edit class-of-service forwarding-classes]
class fc1 queue-num 0 priority high policing-priority premium;
class fc2 queue-num 1 priority low policing-priority normal;
class fc3 queue-num 2 priority low policing-priority normal;
class fc4 queue-num 3 priority low policing-priority normal;
```

For detailed information on class-of-service configuration and statements, see the *JUNOS Class of Service Configuration Guide*.

Firewall Configuration for Hierarchical Policers

```
[edit firewall hierarchical-policer foo]
aggregate {
  if-exceeding {
    bandwidth-limit 70m;
    burst-size-limit 1500;
  }
  then {
    discard;
  }
}
premium {
  if-exceeding {
    bandwidth-limit 50m;
    burst-size-limit 1500;
  }
  then {
    discard;
  }
}
```

You can apply the hierarchical policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-hierarchical-policer foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-hierarchical-policer foo;
```

Configuring a Single-Rate Two-Color Policer

You can configure a single-rate two-color policer as follows:


```
[edit firewall policer foo]
  if-exceeding {
    bandwidth-limit 50m;
    burst-size-limit 1500;
  }
  then {
    discard;
  }
}
```

You can apply the policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-policer foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-policer foo;
```

Configuring a Single-Rate Tricolor Policar

This section describes single-rate color blind and color aware policers.

Configuring a Single-Rate Color-Blind Policar

You can configure a single-rate color blind policer as follows:

```
[edit firewall three-color-policer foo]
single-rate {
  color-blind;
  committed-information-rate 50m;
  committed-burst-size 1500;
  excess-burst-size 1500;
}
```

You can apply the single-rate color blind policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-three-color foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-three-color foo;
```

Configuring a Single-Rate Color-Aware Policar

You can configure a single-rate color-aware policer as follows:

```
[edit firewall three-color-policer bar]
single-rate {
  color-aware;
  committed-information-rate 50m;
  committed-burst-size 1500;
  excess-burst-size 1500;
}
```

You can apply the single-rate color-aware policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-three-color foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-three-color bar;
```

Configuring a Two-Rate Tricolor Marker Policer

Ingress policing is implemented using a two-rate tricolor marker (trTCM). This is done with a dual token bucket (DTB) that maintains two rates, committed, and a peak. Egress static policing also uses a token bucket.

The token buckets perform the following Ingress Policing functions:

- (1K) trTCM - Dual token bucket (red, yellow, and green marking)
- Policing is based on Layer 2 packet size:
 - After +/- byte adjust offset
- Marking is color aware and color blind:
 - Color aware needs to have the color set by q-tree lookup based on:
 - ToS
 - EXP
- Programmable marking actions:
 - Color (red, yellow, green)
 - Drop based on color and congestion profile
- Policer is selected based on the arriving channel number:
 - Channel number LUT produces policer index and queue index
 - Multiple channels can share the same policer (LUT produces same policer index)
- Support ingress policing and trTCM at the following levels:
 - Queue
 - Logical interface (ifl/DLCI)
 - Physical interface (ifd)
 - Physical port (controller ifd)
 - Any combinations of logical interface, physical interface, and port
- Support percentage of interface speed and bits per second

Rate limits may be applied to selected queues on ingress and on predefined queues at egress. The token bucket operates in color aware and color blind modes (specified by RFC 2698).

Configuring a Color-Blind trTCM

```
[edit firewall three-color-policer foo]
two-rate {
  color-blind;
  committed-information-rate 50m;
  committed-burst-size 1500;
  peak-information-rate 100m;
  peak-burst-size 3k;
}
```

You can apply the three-color two-rate color-blind policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-three-color foo;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-three-color foo;
```

Configuring a Color-Aware trTCM

```
[edit firewall three-color-policer bar]
two-rate {
  color-aware;
  committed-information-rate 50m;
  committed-burst-size 1500;
  peak-information-rate 100m;
  peak-burst-size 3k;
}
```

You can apply the three-color two-rate color-aware policer as follows:

```
[edit interfaces so-0/1/0 unit 0 layer-2-policer]
input-three-color bar;
```

You also have the option to apply the policer at the physical port level as follows:

```
[edit interfaces so-0/1/0 layer-2-policer]
input-three-color bar;
```

Applying a Filter to an Interface

To apply firewall filters to an interface, include the `filter` statement:

```
filter {
  group filter-group-number;
  input filter-name;
  input-list [ filter-names ];
  output filter-name;
  output-list [ filter-names ];
}
```

To apply a single filter, include the `input` statement:

```
filter {
  input filter-name;
}
```

To apply a list of filters to evaluate packets received on an interface, include the **input-list** statement.

```
filter {
  input-list [ filter-names ];
}
```

Up to 16 filter names can be included in an input list.

To apply a list of filters to evaluate packets transmitted on an interface, include the **output-list** statement.

```
filter {
  output-list [ filter-names ];
}
```

When you apply filters using the **input-list** statement or the **output-list** statement, a new filter is created with the name `<interface-name> . <unit-direction>`. This filter is exclusively interface-specific.

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *family*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *family*]

In the **family** statement, the protocol family can be `ccc`, `inet`, `inet6`, `mpls`, or `vpls`.

In the **group** statement, specify the interface group number to associate with the filter.

In the **input** statement, list the name of one firewall filter to be evaluated when packets are received on the interface.

In the **input-list** statement, list the names of filters to evaluate when packets are received on the interface. You can include up to 16 filter names.

In the **output** statement, list the name of one firewall filter to be evaluated when packets are transmitted on the interface.



NOTE: Output filters do not work for broadcast and multicast traffic, including VPLS traffic, as shown in “Example: Applying a Filter to an Interface” on page 206.



NOTE: On an MX Series router, you cannot apply as an output filter, a firewall filter configured at the `[edit firewall filter family ccc]` hierarchy level. Firewall filters configured for the `family ccc` statement can be applied only as input filters.

In the `output-list` statement, list the names of filters to evaluate when packets are transmitted on the interface. You can include up to 16 filter names.

You can use the same filter one or more times. On M Series routers (except the M320 and M120 routers), if you apply a firewall filter or policer to multiple interfaces, the filter or policer acts on the sum of traffic entering or exiting those interfaces.

On T Series, M120, and M320 routers, interfaces are distributed among multiple packet forwarding components. Therefore, on these routers, if you apply a firewall filter or policer to multiple interfaces, the filter or policer acts on the traffic stream entering or exiting each interface, regardless of the sum of traffic on the multiple interfaces.

For more information on Understanding Ethernet Frame Statistics, see the *MX Series Layer 2 Configuration Guide*.

If you apply the filter to the interface `lo0`, it is applied to packets received or transmitted by the Routing Engine. You cannot apply MPLS filters to the management interface (`fxp0` or `em0`) or the loopback interface (`lo0`).

For more information about firewall filters, see the *JUNOS Policy Framework Configuration Guide*. For more information about MPLS filters, see the *JUNOS MPLS Applications Configuration Guide*.

See also the following sections:

- Defining Interface Groups in Firewall Filters on page 205
- Filter-Based Forwarding on the Output Interface on page 206
- Example: Applying a Filter to an Interface on page 206

Defining Interface Groups in Firewall Filters

When applying a firewall filter, you can define an interface to be part of an *interface group*. Packets received on that interface are tagged as being part of the group. You can then match these packets using the `interface-group` match statement, as described in the *JUNOS Policy Framework Configuration Guide*.

To define the interface to be part of an interface group, include the `group` statement:

```
group filter-group-number;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number family family filter]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family filter]`

Filter-Based Forwarding on the Output Interface

If port-mirrored packets are to be distributed to multiple monitoring or collection interfaces, based on patterns in packet headers, it is helpful to configure a filter-based forwarding (FBF) filter on the port-mirroring egress interface.

When an FBF filter is installed as an output filter, a packet that is forwarded to the filter has already undergone at least one route lookup. After the packet is classified at the egress interface by the FBF filter, it is redirected to another routing table for additional route lookup. To avoid packet looping inside the Packet Forwarding Engine, the route lookup in the latter routing table (designated by an FBF routing instance) must result in a different next hop from any next hop specified in a table that has already been applied to the packet.

If an input interface is configured for FBF, the source lookup is disabled for those packets headings to a different routing instance, since the routing table is not set up to handle the source lookup.

For more information about FBF configuration, see the *JUNOS Routing Protocols Configuration Guide*. For more information about port mirroring, see the *JUNOS Services Interfaces Configuration Guide*.

Example: Applying a Filter to an Interface

Input Filter for VPLS Traffic

For M Series and T Series routers only, apply an input filter to VPLS traffic. Output filters do not work for broadcast and multicast traffic, including VPLS traffic.

```
[edit interfaces]
fe-2/2/3 {
  vlan-tagging;
  encapsulation vlan-vpls;
  unit 601 {
    encapsulation vlan-vpls;
    vlan-id 601;
    family vpls {
      filter {
        input filter1; # Works for multicast destination MAC address
        output filter1; # Does not work for multicast destination MAC address
      }
    }
  }
}

[edit firewall]
family vpls {
  filter filter1 {
    term 1 {
      from {
        destination-mac-address {
          01:00:0c:cc:cc:cd/48;
        }
      }
      then {
        discard;
      }
    }
  }
}
```

```

    }
    term 2 {
        then {
            accept;
        }
    }
}
}

```

Filter-Based Forwarding at the Output Interface

The following example illustrates the configuration of filter-based forwarding at the output interface. In this example, the packet flow follows this path:

1. A packet arrives at interface **fe-1/2/0.0** with source and destination addresses **10.50.200.1** and **10.50.100.1** respectively.
2. The route lookup in routing table **inet.0** points to the egress interface **so-0/0/3.0**.
3. The output filter installed at **so-0/0/3.0** redirects the packet to routing table **fbf.inet.0**.
4. The packet matches the entry **10.50.100.0/25** in the **fbf.inet.0** table, and finally leaves the router from interface **so-2/0/0.0**.

```

[edit interfaces]
so-0/0/3 {
    unit 0 {
        family inet {
            filter {
                output fbf;
            }
            address 10.50.10.2/25;
        }
    }
}
fe-1/2/0 {
    unit 0 {
        family inet {
            address 10.50.50.2/25;
        }
    }
}
so-2/0/0 {
    unit 0 {
        family inet {
            address 10.50.20.2/25;
        }
    }
}
[edit firewall]
filter fbf {
    term 0 {
        from {
            source-address {
                10.50.200.0/25;
            }
        }
    }
}

```

```

        then routing-instance fbf;
    }
    term d {
        then count d;
    }
}
[edit routing-instances]
fbf {
    instance-type forwarding;
    routing-options {
        static {
            route 10.50.100.0/25 next-hop so-2/0/0.0;
        }
    }
}
[edit routing-options]
interface-routes {
    rib-group inet fbf-group;
}
static {
    route 10.50.100.0/25 next-hop 10.50.10.1;
}
rib-groups {
    fbf-group {
        import-rib [inet.0 fbf.inet.0];
    }
}
}

```

Configuring Unicast RPF

For interfaces that carry IPv4 or IPv6 traffic, you can reduce the impact of denial of service (DoS) attacks by configuring unicast reverse path forwarding (RPF). Unicast RPF helps determine the source of attacks and rejects packets from unexpected source addresses on interfaces where unicast RPF is enabled.



NOTE: If you want to configure unicast RPF, your router must be equipped with the Internet Processor II application-specific integrated circuit (ASIC).

If you enable unicast RPF on live traffic, some packets are dropped while the packet forwarding components are updating.

For transit packets exiting the router through the tunnel, forwarding path features, such as RPF, forwarding table filtering, source class usage, and destination class usage are not supported on the interfaces you configure as the output interface for tunnel traffic. For firewall filtering, you must allow the output tunnel packets through the firewall filter applied to input traffic on the interface that is the next-hop interface towards the tunnel destination.

The following sections describe unicast RPF in detail:

- Configuring Unicast RPF Strict Mode on page 209
- Configuring Unicast RPF Loose Mode on page 210
- Unicast RPF and Default Routes on page 210
- Unicast RPF with Routing Asymmetry on page 212
- Configuring Unicast RPF on a VPN on page 212
- Example: Configuring Unicast RPF on page 213

Configuring Unicast RPF Strict Mode

In strict mode, unicast RPF checks whether the incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix.

If the incoming packet fails the unicast RPF check, the packet is not accepted on the interface. When a packet is not accepted on an interface, unicast RPF counts the packet and sends it to an optional fail filter. If the fail filter is not configured, the default action is to silently discard the packet.

The optional fail filter allows you to apply a filter to packets that fail the unicast RPF check. You can define the fail filter to perform any filter operation, including accepting, rejecting, logging, sampling, or policing.

When unicast RPF is enabled on an interface, Bootstrap Protocol (BOOTP) packets and Dynamic Host Configuration Protocol (DHCP) packets are not accepted on the interface. To allow the interface to accept BOOTP packets and DHCP packets, you must apply a fail filter that accepts all packets with a source address of 0.0.0.0 and a destination address of 255.255.255.255. For a configuration example, see “Example: Configuring Unicast RPF” on page 213.

For more information about unicast RPF, see the *JUNOS Routing Protocols Configuration Guide*. For more information about defining fail filters, see the *JUNOS Policy Framework Configuration Guide*.

To configure unicast RPF, include the `rpf-check` statement:

```
rpf-check <fail-filter filter-name>;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6)]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6)]

Using unicast RPF can have several consequences when implemented with traffic filters:

- RPF fail filters are evaluated after input filters and before output filters.
- If you configure a filter counter for packets dropped by an input filter, and you want to know the total number of packets dropped, you must also configure a filter counter for packets dropped by the RPF check.
- To count packets that fail the RPF check and are accepted by the RPF fail filter, you must configure a filter counter.
- If an input filter forwards packets anywhere other than the `inet.0` or `inet6.0` routing tables, the unicast RPF check is not performed.
- If an input filter forwards packets anywhere other than the routing instance the input interface is configured for, the unicast RPF check is not performed.

Configuring Unicast RPF Loose Mode

By default, unicast RPF uses strict mode. Unicast RPF loose mode is similar to unicast RPF strict mode and has the same configuration restrictions. The only check in loose mode is whether the packet has a source address with a corresponding prefix in the routing table; loose mode does not check whether the interface expects to receive a packet with a specific source address prefix. If a corresponding prefix is not found, unicast RPF loose mode does not accept the packet. As in strict mode, loose mode counts the failed packet and optionally forwards it to a fail filter, which either accepts, rejects, logs, samples, or polices the packet.

To configure unicast RPF loose mode, include the `mode`:

```
mode loose;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) rpf-check <fail-filter *filter-name*>]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) rpf-check <fail-filter *filter-name*>]

Unicast RPF and Default Routes

When the active route cannot be chosen from the routes in a routing table, the router chooses a default route. A default route is equivalent to an IP address of 0.0.0.0/0. If you configure a default route, and you configure unicast RPF on an interface that the default route uses, unicast RPF behaves differently than it does otherwise. For information about configuring default routes, see the *JUNOS Routing Protocols Configuration Guide*.

To determine whether the default route uses an interface, enter the `show route` command:

```
user@host> show route address
```

address is the next-hop address of the configured default route. The default route uses the interfaces shown in the output of the **show route** command.

The following sections describe how unicast RPF behaves when a default route uses an interface and when a default route does not use an interface:

- Unicast RPF Behavior with a Default Route on page 211
- Unicast RPF Behavior Without a Default Route on page 211

Unicast RPF Behavior with a Default Route

If you configure a default route that uses an interface configured with unicast RPF, unicast RPF behaves as follows:

- Loose mode—All packets are automatically accepted. For this reason, we recommend that you not configure unicast RPF loose mode on interfaces that the default route uses.
- Strict mode—The packet is accepted when either of the following is true:
 - The source address of the packet matches any of the routes (either default or learned) that can be originated from the interface. Note that routes can have multiple destinations associated with them; therefore, if one of the destinations matches the incoming interface of the packet, the packet is accepted.
 - The source address of the packet does not match any of the routes.

The packet is not accepted when either of the following is true:

- The source address of the packet does not match a prefix in the routing table.
 - The interface does not expect to receive a packet with this source address prefix.

Unicast RPF Behavior Without a Default Route


If you do not configure a default route, or if the default route does not use an interface configured with unicast RPF, unicast RPF behaves as described in “Configuring Unicast RPF Strict Mode” on page 209 and “Configuring Unicast RPF Loose Mode” on page 210. To summarize, unicast RPF without a default route behaves as follows:

- Strict mode—The packet is not accepted when either of the following is true:
 - The packet has a source address that does not match a prefix in the routing table.
 - The interface does not expect to receive a packet with this source address prefix.
- Loose mode—The packet is not accepted when the packet has a source address that does not match a prefix in the routing table.

Unicast RPF with Routing Asymmetry

In general, we recommend that you not enable unicast RPF on interfaces that are internal to the network because internal interfaces are likely to have *routing asymmetry*. Routing asymmetry means that a packet's outgoing and return paths are different. Routers in the core of the network are more likely to have asymmetric reverse paths than routers at the customer or provider edge. Figure 9 on page 212 shows unicast RPF in an environment with routing asymmetry.

Figure 9: Unicast RPF with Routing Asymmetry



In Figure 9 on page 212, if you enable unicast RPF on interface `so-0/0/0`, traffic destined for Router A is not rejected. If you enable unicast RPF on interface `so-1/0/1`, traffic from Router A is rejected.

If you need to enable unicast RPF in an asymmetric routing environment, you can use fail filters to allow the router to accept incoming packets that are known to be arriving by specific paths. For an example of a fail filter that accepts packets with a specific source and destination address, see “Example: Configuring Unicast RPF” on page 213.

Configuring Unicast RPF on a VPN

You can configure unicast RPF on a VPN interface by enabling unicast RPF on the interface and including the `interface` statement at the `[edit routing-instances routing-instance-name]` hierarchy level.

You can configure unicast RPF only on the interfaces you specify in the routing instance. This means the following:

- For Layer 3 VPNs, unicast RPF is supported on the CE router interface.
- Unicast RPF is not supported on core-facing interfaces.
- For virtual-router routing instances, unicast RPF is supported on all interfaces you specify in the routing instance.
- If an input filter forwards packets anywhere other than the routing instance the input interface is configured for, the unicast RPF check is not performed.

For unicast RPF configuration guidelines, see “Configuring Unicast RPF” on page 208. For more information about VPNs and virtual-router routing instances, see the *JUNOS VPNs Configuration Guide*. For more information about FBF, see the *JUNOS Routing Protocols Configuration Guide*.

Example: Configuring Unicast RPF on a VPN

Configure unicast RPF on a Layer 3 VPN interface:

```
[edit interfaces]
so-0/0/0 {
  unit 0 {
    family inet {
      rpf-check;
    }
  }
}
[edit routing-instance]
VPN-A {
  interface so-0/0/0.0;
}
```

Example: Configuring Unicast RPF

Configure unicast RPF strict mode, and apply a fail filter that allows the interface to accept BOOTP packets and DHCP packets. The filter accepts all packets with a source address of 0.0.0.0 and a destination address of 255.255.255.255.

```
[edit firewall]
filter rpf-special-case-dhcp-bootp {
  term allow-dhcp-bootp {
    from {
      source-address {
        0.0.0.0/32;
      }
      address {
        255.255.255.255/32;
      }
    }
    then {
      count rpf-dhcp-bootp-traffic;
      accept;
    }
  }
  term default {
    then {
      log;
      reject;
    }
  }
}
[edit]
interfaces {
  so-0/0/0 {
    unit 0 {
      family inet {
        rpf-check fail-filter rpf-special-case-dhcp-bootp;
      }
    }
  }
}
```

```
}
}
```

Enabling Source Class and Destination Class Usage

For interfaces that carry IPv4, IPv6, or MPLS traffic, you can maintain packet counts based on the entry and exit points for traffic passing through your network. Entry and exit points are identified by source and destination prefixes grouped into disjoint sets defined as *source classes* and *destination classes*. You can define classes based on a variety of parameters, such as routing neighbors, autonomous systems, and route filters.

Source class usage (SCU) counts packets sent to customers by performing lookup on the IP source address. SCU makes it possible to track traffic originating from specific prefixes on the provider core and destined for specific prefixes on the customer edge. You must enable SCU accounting on both the inbound and outbound physical interfaces, and the route for the source of the packet must be in located in the forwarding table.



NOTE: SCU and DCU accounting do not work with directly connected interface routes. Source class usage does not count packets coming from sources with direct routes in the forwarding table because of software architecture limitations.

Destination class usage (DCU) counts packets from customers by performing lookup of the IP destination address. DCU makes it possible to track traffic originating from the customer edge and destined for specific prefixes on the provider core router.



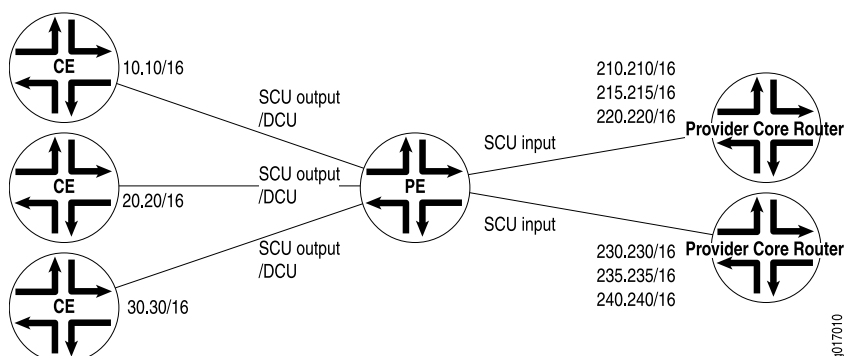
NOTE: SCU and DCU accounting are supported on the J Series router only for IPv4 and IPv6 traffic.



NOTE: We recommend that you stop the network traffic on an interface before you modify the DCU or SCU configuration for that interface. Modifying the DCU or SCU configuration without stopping the traffic might corrupt the DCU or SCU statistics. Before you restart the traffic after modifying the configuration, enter the **clear interfaces statistics** command.

Figure 10 on page 215 illustrates an Internet service provider (ISP) network. In this topology, you can use DCU to count packets customers send to specific prefixes. For example, you can have three counters, one per customer, that count the packets destined for prefix 210.210/16 and 220.220/16.

You can use SCU to count packets the provider sends from specific prefixes. For example, you can count the packets sent from prefix 210.210/16 and 215.215/16 and transmitted on a specific output interface.

Figure 10: Prefix Accounting with Source and Destination Classes

You can configure up to 126 source classes and 126 destination classes. For each interface on which you enable destination class usage and source class usage, the JUNOS Software maintains an interface-specific counter for each corresponding class up to the 126 class limit.



NOTE: To configure source class and destination class usage, your router must be equipped with the Internet Processor II ASIC.

For transit packets exiting the router through the tunnel, forwarding path features, such as RPF, forwarding table filtering, source class usage, and destination class usage are not supported on the interfaces you configure as the output interface for tunnel traffic. For firewall filtering, you must allow the output tunnel packets through the firewall filter applied to input traffic on the interface that is the next-hop interface towards the tunnel destination.



NOTE:

Performing DCU accounting when an output service is enabled produces inconsistent behavior in the following configuration:

- Both SCU input and DCU are configured on the packet input interface.
- SCU output is configured on the packet output interface.
- Interface services is enabled on the output interface.

For an incoming packet with source and destination prefixes matching the SCU and DCU classes respectively configured in the router, both SCU and DCU counters will be incremented. This behavior is not harmful or negative. However, it is inconsistent with non-serviced packets, in that only the SCU count will be incremented (because the SCU class ID will override the DCU class ID in this case).

To enable packet counting on an interface, include the **accounting** statement:

```
accounting {
```

```

    destination-class-usage;
    source-class-usage {
        direction;
    }
}

```

direction can be one of the following:

- **input**—Configure at least one expected ingress point.
- **output**—Configure at least one expected egress point.
- **input output**—On a single interface, configure at least one expected ingress point and one expected egress point.

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6 | mpls)]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6 | mpls)]

For SCU to work, you must configure at least one input interface and at least one output interface. An incoming packet is counted only once, and SCU takes priority over DCU. This means that when a packet arrives on an interface on which you include the **source-class-usage input** and **destination-class-usage** statements in the configuration, and when the source and destination both match accounting prefixes, the JUNOS Software associates the packet with the source class only. To ensure the outgoing packet is counted, include the **source-class-usage output** statements in the configuration of the outgoing interface.

On T Series, M120, and M320 routers, the source class and destination classes are not carried across the router fabric. The implications of this are as follows:

- On T Series, M120, and M320 routers, SCU and DCU accounting is performed before the packet enters the fabric.
- On T Series, M120, and M320 routers, DCU is performed before output filters are evaluated. On other M Series routers, DCU is performed after output filters are evaluated.
- If an output filter drops traffic on T Series, M120, and M320 routers, the dropped packets are included in DCU statistics. If an output filter drops traffic on other M Series routers, the dropped packets are excluded from DCU statistics.

Once you enable accounting on an interface, the JUNOS Software maintains packet counters for that interface, with separate counters for **inet**, **inet6**, and **mpls** protocol families. You must then configure the source class and destination class attributes in policy action statements, which must be included in forwarding-table export policies.

For a complete discussion about source and destination class accounting profiles, see the *JUNOS Network Management Configuration Guide*. For more information about MPLS, see the *JUNOS MPLS Applications Configuration Guide*.

Examples: Enabling Source Class and Destination Class Usage

Configure DCU and SCU output on one interface:

```
[edit]
interfaces {
  so-6/1/0 {
    unit 0 {
      family inet {
        accounting {
          destination-class-usage;
          source-class-usage {
            output;
          }
        }
      }
    }
  }
}
```

Complete SCU Configuration

Source routers A and B use loopback addresses as the prefixes to be monitored. Most of the configuration tasks and actual monitoring occur on transit Router SCU.

The loopback address on Router A contains the origin of the prefix that is to be assigned to source class A on Router SCU. However, no SCU processing happens on this router. Therefore, configure Router A for basic Open Shortest Path First (OSPF) routing and include your loopback interface and interface **so-0/0/2** in the OSPF process.

Router A

```
[edit]
interfaces {
  so-0/0/2 {
    unit 0 {
      family inet {
        address 10.255.50.2/24;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.192.10/32;
      }
    }
  }
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface so-0/0/2.0;
      interface lo0.0;
    }
  }
}
```

Router SCU Router SCU handles the bulk of the activity in this example. On Router SCU, enable source class usage on the inbound and outbound interfaces at the `[edit interfaces interface-name unit unit-number family inet accounting]` hierarchy level. Make sure you specify the expected traffic: input, output, or, in this case, both.

Next, configure a route filter policy statement that matches the prefixes of the loopback addresses from routers A and B. Include statements in the policy that classify packets from Router A in one group named **scu-class-a** and packets from Router B in a second class named **scu-class-b**. Notice the efficient use of a single policy containing multiple terms.

Last, apply the policy to the forwarding table.

```
[edit]
interfaces {
  so-0/0/1 {
    unit 0 {
      family inet {
        accounting {
          source-class-usage {
            input;
            output;
          }
        }
        address 10.255.50.1/24;
      }
    }
  }
  so-0/0/3 {
    unit 0 {
      family inet {
        accounting {
          source-class-usage {
            input;
            output;
          }
        }
        address 10.255.10.3/24;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.6.111/32;
      }
    }
  }
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface so-0/0/1.0;
      interface so-0/0/3.0;
    }
  }
}
```

```

    }
  }
  routing-options {
    forwarding-table {
      export scu-policy;
    }
  }
  policy-options {
    policy-statement scu-policy {
      term 0 {
        from {
          route-filter 10.255.192.0/24 orlonger;
        }
        then source-class scu-class-a;
      }
      term 1 {
        from {
          route-filter 10.255.165.0/24 orlonger;
        }
        then source-class scu-class-b;
      }
    }
  }
}

```

Router B Just as Router A provides a source prefix, Router B's loopback address matches the prefix assigned to **scu-class-b** on Router SCU. Again, no SCU processing happens on this router, so configure Router B for basic OSPF routing and include your loopback interface and interface **so-0/0/4** in the OSPF process.

```

interfaces {
  so-0/0/4 {
    unit 0 {
      family inet {
        address 10.255.10.4/24;
      }
    }
  }
  lo0 {
    unit 0 {
      family inet {
        address 10.255.165.226/32;
      }
    }
  }
}
protocols {
  ospf {
    area 0.0.0.0 {
      interface so-0/0/4.0;
      interface lo0.0;
    }
  }
}

```

Enabling Packet Counting for Layer 3 VPNs

You can use SCU and DCU to count packets on Layer 3 VPNs. To enable packet counting for Layer 3 VPN implementations at the egress point of the MPLS tunnel, you must configure a virtual loopback tunnel interface (vt) on the PE router, map the virtual routing and forwarding (VRF) instance type to the virtual loopback tunnel interface, and send the traffic received from the VPN out the source class output interface, as shown in the following example:

1. Configure a virtual loopback tunnel interface on a provider edge router equipped with a tunnel PIC:

```
[edit interfaces]
vt-0/3/0 {
  unit 0 {
    family inet {
      accounting {
        source-class-usage {
          input;
        }
      }
    }
  }
}
```

2. Map the VRF instance type to the virtual loopback tunnel interface.

For SCU and DCU to work, you must not include the `vrf-table-label` statement at the `[edit routing-instances instance-name]` hierarchy level.

```
[edit routing-instances]
VPN-A {
  instance-type vrf;
  interface at-2/1/1.0;
  interface vt-0/3/0.0;
  route-distinguisher 10.255.14.225:100;
  vrf-import import-policy-A;
  vrf-export export-policy-A;
  protocols {
    bgp {
      group to-r4 {
        local-address 10.27.253.1;
        peer-as 400;
        neighbor 10.27.253.2;
      }
    }
  }
}
```

3. Send traffic received from the VPN out the source class output interface:

```
[edit interfaces]
at-2/1/0 {
  unit 0 {
    family inet {
      accounting {
        source-class-usage {
          output;
        }
      }
    }
  }
}
```

```
}  
}  
}  
}
```

For more information about VPNs, see the *JUNOS VPNs Configuration Guide*. For more information about virtual loopback tunnel interfaces, see the *JUNOS Services Interfaces Configuration Guide*.

Chapter 6

Configuring Circuit and Translational Cross-Connects

This chapter describes circuit and translational cross-connects:

- Circuit and Translational Cross-Connects Overview on page 223
- Defining the Encapsulation for Switching Cross-Connects on page 225
- Defining the Connection for Switching Cross-Connects on page 228
- Configuring MPLS for Switching Cross-Connects on page 229
- Configuring IS-IS or MPLS Traffic for TCC Interfaces on page 229
- Configuring ATM-to-Ethernet Interworking on page 229
- Examples: Configuring Switching Cross-Connects on page 233

Circuit and Translational Cross-Connects Overview

Circuit cross-connect (CCC) and translational cross-connect (TCC) allow you to configure transparent connections between two circuits, where a circuit can be a Frame Relay data-link connection identifier (DLCI), an Asynchronous Transfer Mode (ATM) virtual circuit (VC), a Point-to-Point Protocol (PPP) interface, a Cisco High-level Data Link Control (HDLC) interface, or a Multiprotocol Label Switching (MPLS) label-switched path (LSP).

Using CCC or TCC, packets from the source circuit are delivered to the destination circuit with, at most, the Layer 2 address being changed. No other processing, such as header checksums, time-to-live (TTL) decrementing, or protocol processing, is done.

To connect interfaces of the same type, use CCC. To connect unlike interfaces, use TCC.

CCC and TCC circuits fall into three categories: logical interfaces, which include ATM VCs and Frame Relay DLCIs; physical interfaces, which include PPP and Cisco HDLC; and paths, which include LSPs. The three circuit categories provide three types of cross-connect:

- Layer 2 switching (interface-to-interface)—Cross-connects between logical interfaces provide what is essentially Layer 2 switching.

- **MPLS tunneling (interface-to-LSP)**—Cross-connects between interfaces and LSPs allow you to connect two distant interface circuits by creating MPLS tunnels that use LSPs as the conduit.
- **LSP stitching (LSP-to-LSP)**—Cross-connects between LSPs provide a way to “stitch” together two label-switched paths, including paths that fall in two different traffic engineering database (TED) areas.

The cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first interface.

For most CCC connections that connect interfaces, the interfaces must be of the same type; that is, ATM to ATM, Frame Relay to Frame Relay, PPP to PPP, or Cisco HDLC to Cisco HDLC.


ATM-to-Ethernet interworking cross-connect circuits connect logical interfaces configured on an ATM2 and Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet IQ2 and IQ2-E physical interfaces.

For all TCC connections that connect interfaces, the interfaces can be of unlike types. Mainly, TCC is used for Layer 2.5 virtual private networks (VPNs), but it can also be used as a simple “unlike circuit” switch.

Switching cross-connects join logical interfaces to form what is essentially Layer 2 switching.

Figure 11 on page 224 illustrates a Layer 2 switching circuit cross-connect. In this topology, Router A and Router C have Frame Relay connections to Router B, which is a Juniper Networks router. CCC allows you to configure Router B to act as a Frame Relay (Layer 2) switch. To do this, configure a circuit from Router A to Router C that passes through Router B, effectively configuring Router B as a Frame Relay switch with respect to these routers. This configuration allows Router B to transparently switch packets (frames) between Router A and Router C without regard to the packets’ contents or the Layer 3 protocols. The only processing that Router B performs is to translate DLCI 600 to 750.

Figure 11: Layer 2 Switching Circuit Cross-Connect



If the Router A-to-Router B and Router B-to-Router C circuits are PPP, for example, the Link Control Protocol and Network Control Protocol exchanges occur between Router A and Router C. These messages are handled transparently by Router B, allowing Router A and Router C to use various PPP options (such as header or address compression and authentication) that Router B might not support. Similarly, Router A and Router C exchange keepalives, providing circuit-to-circuit connectivity status.

You can configure Layer 2 switching cross-connects on PPP, Cisco HDLC, Frame Relay, Ethernet CCC, Ethernet VLAN, and ATM circuits. With CCC, only like interfaces

can be connected in a single cross-connect. With TCC, unlike interfaces can be connected in a single cross-connect. In Layer 2 switching cross-connects, the exchanges take place between point-to-point links.

This chapter discusses the Layer 2 switching cross-connect configuration tasks. For information about MPLS tunneling and LSP stitching, see the *JUNOS MPLS Applications Configuration Guide*.

For information about Layer 2 and Layer 2.5 VPNs, see the *JUNOS VPNs Configuration Guide*.

To configure switching cross-connects, you must configure the following on the router that is acting as the switch (Router B in Figure 11 on page 224):

Defining the Encapsulation for Switching Cross-Connects

To configure Layer 2 or Layer 2.5 switching cross-connects, configure the CCC or TCC encapsulation on the router that is acting as the switch (Router B in Figure 11 on page 224).



NOTE: When you use CCC encapsulation, you can configure the `ccc` family only. Likewise, when you use TCC encapsulation, you can configure the `tcc` family only.

This section contains the following topics:

- Configuring PPP or Cisco HDLC Circuits on page 225
- Configuring ATM Circuits on page 225
- Configuring Frame Relay Circuits on page 226
- Configuring Ethernet CCC Circuits on page 227
- Configuring Ethernet VLAN Circuits on page 228

Configuring PPP or Cisco HDLC Circuits

For PPP or Cisco HDLC circuits, specify the encapsulation by including the `encapsulation` statement at the `[edit interfaces interface-name]` hierarchy level. This statement configures the entire physical device. For these circuits to work, you must configure a logical interface unit 0.

```
[edit interfaces interface-name]
encapsulation (ppp-ccc | cisco-hdlc-ccc | ppp-tcc | cisco-hdlc-tcc);
unit 0;
```

Configuring ATM Circuits

For ATM circuits, include the `vpi` statement `[edit interfaces interface-name atm-options]` hierarchy level:

```
[edit interfaces at-fpc/pic/port]
atm-options {
```

```

    vpi vpi-identifier;
}

```

On the logical interface, include the following statements:

```

point-to-point;
encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-tcc-vc-mux | atm-tcc-snap);
vci vpi-identifier.vci-identifier;

```

You can include the logical interface statements at the following hierarchy levels:

- [edit interfaces *at-fpc/pic/port* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *at-fpc/pic/port* unit *logical-unit-number*]

For each VC, configure whether it is a circuit or a regular logical interface. The default interface type is point-to-point.

Configuring Frame Relay Circuits

For Frame Relay circuits, include the **encapsulation** statement at the [edit interfaces *interface-name*] hierarchy level:

```

[edit interfaces interface-name]
encapsulation type;

```

On the logical interface, include the following statements:

```

point-to-point;
encapsulation type;
dlci dlci-identifier;

```

You can include the logical interface statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The encapsulation type can be one of the following:

- **Flexible Frame Relay (flexible-frame-relay)**—Intelligent queuing (IQ) interfaces can use flexible Frame Relay encapsulation. You use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.
- **Frame Relay CCC version (frame-relay-ccc)**—For E1, E3, SONET/SDH, T1, and T3 interfaces, this encapsulation type is the same as standard Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to CCC. The logical interface must also have **frame-relay-ccc** encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.

- Frame Relay TCC version (**frame-relay-tcc**)—Similar to Frame Relay CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- Extended CCC version (**extended-frame-relay-ccc**)—This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC. The logical interface must have **frame-relay-ccc** encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.
- Extended TCC version (**extended-frame-relay-tcc**)—Similar to extended Frame Relay CCC, this encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC, which is used for circuits with different media on either side of the connection.
- Port CCC version (**frame-relay-port-ccc**)—Defined in the IETF document *Frame Relay Encapsulation over Pseudo-Wires* (expired December 2002). This encapsulation type allows you to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. The connection between the two CE routers can be either user-to-network interface (UNI) or network-to-network interface (NNI); this is completely transparent to the PE routers. The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the **ccc** family only.

For each DLCI, configure whether it is a circuit or a regular logical interface. The DLCI for regular interfaces must be from 1 through 511. For CCC and TCC interfaces, it must be from 512 through 1022. This restriction does not apply to IQ interfaces. The default interface type is point to point.

Configuring Ethernet CCC Circuits

You can configure Ethernet CCC encapsulation on Fast Ethernet, Gigabit Ethernet, and aggregated Ethernet interfaces.



NOTE: CCC over aggregated Ethernet requires an M Series Enhanced Flexible PIC Concentrator (FPC).

For Ethernet CCC circuits, specify the encapsulation by including the **encapsulation** statement at the [edit interfaces *interface-name*] hierarchy level. This statement configures the entire physical device.

```
[edit interfaces interface-name]
encapsulation ethernet-ccc;
unit logical-unit-number {
    ...
}
[edit interfaces aex]
encapsulation ethernet-ccc;
unit logical-unit-number {
    ...
}
```

Configuring Ethernet VLAN Circuits

You can configure Ethernet virtual local area network (VLAN) circuits on Fast Ethernet, Gigabit Ethernet, and aggregated Ethernet interfaces. For Ethernet VLAN circuits, specify the encapsulation by including the `encapsulation` statement at the `[edit interfaces interface-name]` hierarchy level. This statement configures the entire physical device. You must also enable VLAN tagging. To do this, include the following statements:

```
[edit interfaces interface-name]
vlan-tagging;
encapsulation (extended-vlan-ccc | vlan-ccc);
[edit interfaces aex]
vlan-tagging;
encapsulation vlan-ccc;
```

On the logical interface, include the following statements:

```
encapsulation vlan-ccc;
vlan-id number;
```

You can include the logical interface statements at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]`

Ethernet interfaces in VLAN mode can have multiple logical interfaces. For encapsulation type `vlan-ccc`, VLAN IDs 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 1023 are reserved for CCC VLANs. For encapsulation type `extended-vlan-ccc`, VLAN IDs 1 through 4094 are valid. VLAN ID 0 is reserved for tagging the priority of frames.

Defining the Connection for Switching Cross-Connects

To configure Layer 2 switching cross-connects, define the connection between the two circuits. You configure this on the router that is acting as the switch (Router B in Figure 11 on page 224). The connection joins the interface that comes from the circuit's source to the interface that leads to the circuit's destination. When you specify the interface names, include the logical portion of the name, which corresponds to the logical unit number. The cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first interface.

```
[edit protocols]
connections {
  remote-interface-switch connection-name {
    interface interface-name.unit-number;
  }
  lsp-switch connection-name {
    transmit-lsp lsp-number;
```

```

        receive-lsp lsp-number;
    }
}

```

Configuring MPLS for Switching Cross-Connects

For Layer 2 switching cross-connects to work, you must configure MPLS. The following is a minimal MPLS configuration:

```

[edit protocols]
mpls {
    interface (interface-name | all);
}

```

For more information, see the *JUNOS MPLS Applications Configuration Guide*.

Configuring IS-IS or MPLS Traffic for TCC Interfaces

Layer 2.5 VPNs on T Series, M120, and M320 routers support IPv4, IS-IS, and MPLS traffic types. By default, IPv4 traffic runs on T Series, M120, and M320 routers and over TCC interfaces. To configure IS-IS (ISO traffic) or MPLS traffic on Layer 2.5 VPNs, you must configure the same traffic type on both ends of the Layer 2.5 VPN.

To specify which traffic can run over a TCC interface, include the **protocols** statement with the appropriate value (**inet**, **mpls**, and **iso**) at the [edit interfaces *interface-name* unit *logical-unit-number* family *tcc*] hierarchy level:

```

[edit interfaces interface-name unit logical-unit-number family tcc]
protocols [ inet iso mpls ];

```



NOTE: Layer 2.5 VPNs running on M Series Multiservice Edge Routers support only IPv4 traffic. IPv6 is not supported on Layer 2.5 VPNs.

When enabling ISO over a Layer 2.5 VPN that is configured on a CE Ethernet interface, you must also include the **point-to-point** statement at the [edit protocols isis interface *interface-name*] hierarchy level:

```

[edit protocols isis interface interface-name]
point-to-point;

```

For more information about Layer 2 VPNs, see the *JUNOS VPNs Configuration Guide*.

Configuring ATM-to-Ethernet Interworking

The ATM-to-Ethernet interworking feature is useful where ATM2 interfaces are used to terminate ATM DSLAM traffic. The ATM traffic can be forwarded with encapsulation type **ccc** (circuit cross-connect) to a local or remote Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet IQ2 and IQ2-E interface or label-switched path (LSP). The ATM VPI and VCI are converted to stacked VLAN inner and outer VLAN tags.

These ATM-to-Ethernet interworking circuits can be mapped to individual logical interfaces configured on an ATM2 IQ interface and Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet IQ2 and IQ2-E physical interface.

The ATM-to-Ethernet interworking cross-connect essentially provides Layer 2 switching, and statistics are reported at the logical interface level.

During conversion from ATM to Ethernet, the least significant 12 bits of the ATM cell VCI are copied to the Ethernet frame inner VLAN tag. Cells received on an ATM logical interface configured with encapsulation type `vlan-vci-ccc` and falling within the configured VCI range are reassembled into packets and forwarded to a designated Ethernet logical interface that is configured with encapsulation type `vlan-vci-ccc`.

During conversion from Ethernet to ATM, the Ethernet frame inner VLAN tags that fall within the configured range, are copied to the least significant 12 bits of the ATM cell VCI. The ATM logical interface uses its configured VPI when segmenting the Ethernet packets into cells.

ATM-to-Ethernet interworking is supported on M120, M320, and T Series routers.

The following sections discuss ATM-to-Ethernet interworking:

- Enabling ATM-to-Ethernet Interworking on page 230
- Configuring the ATM-to-Ethernet Interworking Ethernet Interface on page 230
- Configuring the ATM-to-Ethernet Interworking Ethernet Encapsulation on page 231
- Configuring the ATM-to-Ethernet Interworking Outer VLAN Identifier on page 231
- Configuring the ATM-to-Ethernet Interworking Inner VLAN Identifier Range on page 231
- Configuring the ATM-to-Ethernet Interworking Physical Interface VPI on page 232
- Configuring the ATM-to-Ethernet Interworking ATM Logical Interface on page 232
- Configuring the ATM-to-Ethernet Interworking Protocol Family on page 232
- Configuring the ATM-to-Ethernet Interworking Logical Interface VPI on page 233
- Configuring the ATM-to-Ethernet Interworking Logical Interface VCI on page 233

Enabling ATM-to-Ethernet Interworking

To enable the ATM-to-Ethernet interworking cross-connect function, include the `vlan-vci-tagging` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
vlan-vci-tagging;
```

Configuring the ATM-to-Ethernet Interworking Ethernet Interface

Configure the Ethernet or aggregated Ethernet physical interface by including the `encapsulation` statement with the `vlan-vci-ccc` option at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]
```

```
encapsulation vlan-vci-ccc;
```

When the encapsulation type `vlan-vci-ccc` is configured on the physical interface, all logical interfaces configured on the Ethernet interface must also have the encapsulation type set to `vlan-vci-ccc`.

Configuring the ATM-to-Ethernet Interworking Ethernet Encapsulation

Configure the Ethernet logical interface by including the `encapsulation` statement with the `vlan-vci-ccc` option at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
encapsulation vlan-vci-ccc;
```

The chassis configuration cannot contain the `atm-l2circuit-mode` statement if any logical interfaces are configured with the `vlan-vci-ccc` encapsulation option.

Configuring the ATM-to-Ethernet Interworking Outer VLAN Identifier

Configure the Ethernet logical interface outer VLAN ID by including the `vlan-id` statement specifying the outer VLAN ID at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
vlan-id outer-vlan-identifier;
```

It is the administrator's responsibility to ensure that the outer VLAN tag and VPI match and the inner VLAN tags fall within the VCI range of the VPI.

The allowable VPI range is from 0 to 255. So the outer VLAN tags must not be configured for values above 255.

Configuring the ATM-to-Ethernet Interworking Inner VLAN Identifier Range

Configure the Ethernet logical interface inner VLAN ID range by including the `inner-vlan-id-range` statement and specifying the starting VLAN ID and ending VLAN ID at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
inner-vlan-id-range start start-id end end-id;
```

VLAN IDs 0 and 4095 are reserved by IEEE 801.1q and must not be used for the inner or outer VLAN ID.

VCIs 0 through 31 are reserved for ATM management purposes by convention. Therefore inner VLAN IDs 1 through 31 should not be used.

VLAN ID 1 might be used by Ethernet switches for certain bridge management services, so using VLAN ID 1 for the inner or outer VLAN ID is discouraged.

Configuring the ATM-to-Ethernet Interworking Physical Interface VPI

Configure the ATM physical interface VPI by including the `vpi` statement at the [edit interfaces *interface-name* atm-options] hierarchy level:

```
[edit interfaces interface-name atm-options]
vpi virtual-path-identifier;
```

VPI 0 is reserved, and must not be used.

ATM F4/F5 OAM is not supported for VPIs used in ATM-to-Ethernet interworking cross-connects. Any F4/F5 OAM cells received are discarded.

Only one logical interface may be declared per virtual path specified in the `atm-options` statement hierarchy.

It is not necessary to dedicate all the VPIs of an ATM2 interface for ATM-to-Ethernet interworking cross-connects.

Configuring the ATM-to-Ethernet Interworking ATM Logical Interface

Configure the ATM logical interface by including the `encapsulation` statement and specifying the encapsulation type `vlan-vci-ccc` at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
encapsulation vlan-vci-ccc;
```

An ATM logical interface configured with the encapsulation type `vlan-vci-ccc` only supports the `epd-threshold`, `shaping`, `traps | no-traps`, `disable`, and `description` statements. No other configuration statements are supported. ATM interface CoS features are not supported by logical interfaces configured with the encapsulation type `vlan-vci-ccc`.

The ATM2 OC48 PIC does not support the encapsulation type `vlan-vci-ccc`.

The encapsulation type `vlan-vci-ccc` only supports the `ccc` protocol family. Attempts to configure any other interface protocol family are rejected.

Configuring the ATM-to-Ethernet Interworking Protocol Family

Configure the ATM logical interface protocol family by including the `family` statement and specifying the `ccc` option at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
family ccc;
```


Configuring the ATM-to-Ethernet Interworking Logical Interface VPI

Configure the ATM logical interface virtual path identifier by including the `vpi` statement at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]  
vpi virtual-path-identifier;
```

VPI 0 is reserved, and must not be used.

It is the administrator's responsibility to ensure the outer VLAN tag and VPI match and the inner VLAN tags fall within the VCI range of the VPI.

Once a VPI is used in an ATM-to-Ethernet interworking cross-connect, it cannot be used with any other logical interface, even if the `vpi.vci` value falls outside the VCI range for the cross-connect.

Configuring the ATM-to-Ethernet Interworking Logical Interface VCI

Configure the ATM logical interface virtual channel identifier range by including the `vci-range` statement and specifying the starting VCI and ending VCI at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]  
vci-range start start-vci end end-vci;
```

Do not use VCIs 0 through 31, which are reserved for ATM management purposes by convention.


Examples: Configuring Switching Cross-Connects

This section includes the following examples:

- Example: Configuring a CCC over Frame Relay Encapsulated Interface on page 233
- Example: Configuring a TCC on page 234
- Example: Configuring CCC over Aggregated Ethernet on page 236
- Example: Configuring a Remote LSP CCC over Aggregated Ethernet on page 237
- Example: Configuring ATM-to-Ethernet Interworking on page 239

Example: Configuring a CCC over Frame Relay Encapsulated Interface

Configure a full-duplex Layer 2 switching circuit cross-connect between Router A and Router C, using a Juniper Networks router, Router B, as the virtual switch. See the topology in Figure 12 on page 234.


Figure 12: Example Topology of a Switching Circuit Cross-Connect with Frame Relay CCC Encapsulation

```
[edit]
interfaces {
  so-1/0/0 {
    encapsulation frame-relay-ccc;
    unit 1 {
      point-to-point;
      eui-64 frame-relay-ccc;
      dlci 600;
    }
  }
  so-2/0/0 {
    encapsulation frame-relay-ccc;
    unit 2 {
      point-to-point;
      encapsulation frame-relay-ccc;
      dlci 750;
    }
  }
}
protocols {
  connections {
    interface-switch router-a-router-c {
      interface so-1/0/0.1;
      interface so-2/0/0.2;
    }
  }
  mpls {
    interface all;
  }
}
```

Example: Configuring a TCC

Configure a full-duplex switching translational cross-connect with PPP TCC encapsulation between Router A and Router C, using a Juniper Networks router, Router B, as the virtual switch. See the topology in Figure 13 on page 235.

In this topology, Router B has a PPP connection to Router A and an ATM connection to Router C.

Figure 13: Layer 2.5 Switching Translational Cross-Connect

On Router A

```
[edit]
interfaces {
  so-0/1/0 {
    description "to Router B so-1/0/0";
    encapsulation ppp;
    unit 0 {
      family inet {
        address 10.1.1.1/30;
      }
    }
  }
}
```

On Router B

```
[edit]
interfaces {
  so-1/0/0 {
    description "to Router A so-0/1/0";
    encapsulation ppp-tcc;
    unit 0 {
    }
  }
  at-1/1/0 {
    description "to Router C at-0/3/0";
    atm-options {
      vpi 0 maximum-vc 2000;
    }
    unit 32 {
      vci 32;
      encapsulation atm-tcc-vc-mux;
    }
  }
}
[edit]
protocols {
  mpls {
    interface so-1/0/0.0;
    interface at-1/1/0.32;
  }
  connections {
    interface-switch PPP-to-ATM {
      interface so-1/0/0.0;
      interface at-1/1/0.32;
    }
  }
}
```

On Router C

```
[edit]
interfaces {
  at-0/3/0 {
```

```


description "to Router B at-1/1/0";
atm-options {
    vpi 0 maximum-vc 2000;
}
unit 32 {
    vci 32;
    encapsulation atm-vc-mux;
    family inet {
        address 10.1.1.2/30;
    }
}
}
}

```

Example: Configuring CCC over Aggregated Ethernet

See the topology in Figure 14 on page 236. In this topology, CE Routers A and C have aggregated Ethernet connections to PE Router B. With CCC, you specify that the circuit from Router A is connected to the circuit from Router C. Router B functions as a cross-connect switch between the two circuits. For a back-to-back connection, all VLAN IDs must be the same on Router A through Router C. You configure Router A and Router C as standard aggregated Ethernet interfaces. For more information about aggregated Ethernet, see “Configuring Aggregated Ethernet Interfaces” on page 623.

Figure 14: Interface-to-Interface Circuit Cross-Connect over Aggregated Ethernet Interfaces



On Router A

```

[edit interfaces]
ae0 {
    vlan-tagging;
    aggregated-ether-options {
        minimum-links 1;
        link-speed 1g;
    }
    unit 0 {
        vlan-id 600;
        family inet {
            address 192.168.1.1/30;
        }
    }
}

```

On Router B

```

[edit interfaces]
ae0 {
    encapsulation vlan-ccc;
    vlan-tagging;
    aggregated-ether-options {
        minimum-links 1;
        link-speed 1g;
    }
}

```

```

    }
    unit 0 { # CCC switch
        encapsulation vlan-ccc;
        vlan-id 600;
        family ccc;
    }
    ae1 {
        encapsulation vlan-ccc;
        vlan-tagging;
        aggregated-ether-options {
            minimum-links 1;
            link-speed 100m;
        }
        unit 0 {
            encapsulation vlan-ccc;
            vlan-id 600;
            family ccc;
        }
    }
    [edit protocols]
    mpls {
        interface all;
    }
    connections {
        interface-switch layer2-cross-connect {
            interface ae0.0;
            interface ae1.0;
        }
    }
}

```

On Router C

```

[edit interfaces]
ae1 {
    vlan-tagging;
    aggregated-ether-options {
        minimum-links 1;
        link-speed 1g;
    }
    unit 0 {
        vlan-id 600;
        family inet {
            address 192.168.1.2/30;
        }
    }
}


```

Example: Configuring a Remote LSP CCC over Aggregated Ethernet

See the topology in Figure 15 on page 238. In this topology, CE Router G has an aggregated Ethernet connection to PE Router F. CE Router D has an aggregated Ethernet connection to PE Router E. Router E and Router F have an MPLS LSP between them. With remote CCC, you specify that the circuit from Router D is connected to the circuit from Router G. The circuit from Router D is connected to the LSP on Router E; the circuit from Router G is connected to the LSP on Router F. In other words, **ae0.0** and **ae1.0** are connected using **lsp1-2** and **lsp2-1**. You configure Router D and

Router G as standard aggregated Ethernet interfaces. For more information about aggregated Ethernet, see “Configuring Aggregated Ethernet Interfaces” on page 623.

Figure 15: Remote Interface-LSP-Interface Circuit Cross-Connect over Aggregated Ethernet Interfaces



On Router D

```
[edit interface]
ae0 {
  aggregated-ether-options {
    minimum-links 1;
    link-speed 1g;
    lacp {
      active;
      periodic fast;
    }
  }
  unit 0 {
    family inet {
      address 192.168.2.1/30;
    }
  }
}
```

On Router E

```
[edit interfaces]
ae0 {
  encapsulation ethernet-ccc;
  aggregated-ether-options {
    minimum-links 1;
    link-speed 100m;
    lacp {
      active;
      periodic fast;
    }
  }
  unit 0 {
    encapsulation vlan-ccc; # default
    family ccc; # default
  }
}

[edit protocols]
mpls {
  interface all;
}

connections {
  remote-interface-switch remote-sw-1 {
    interface ae1.0;
    receive-lsp lsp2_1;
    transmit-lsp lsp1_2;
  }
}
```

```


On Router F [edit interfaces]
ae1 {
  encapsulation ethernet-ccc;
  aggregated-ether-options {
    minimum-links 1;
    link-speed 100m;
    lacp {
      active;
      periodic fast;
    }
  }
}
unit 0 {
  encapsulation vlan-ccc; # default
  family ccc; # default
}
}
[edit protocols]
mpls {
  interface all;
}
connections {
  remote-interface-switch remote-sw-2 {
    interface ae0.0;
    receive-lsp lsp1_2;
    transmit-lsp lsp2_1;
  }
}
}

On Router G [edit interface]
ae1 {
  aggregated-ether-options {
    minimum-links 1;
    link-speed 1g;
    lacp {
      active;
      periodic fast;
    }
  }
}
unit 0 {
  family inet {
    address 192.168.2.2/30;
  }
}
}

```

Example: Configuring ATM-to-Ethernet Interworking

The following example shows the configuration of the ATM and Ethernet interfaces for an ATM-to-Ethernet interworking cross connect. In the example ATM DSLAM traffic is terminated on an ATM2 interface. The ATM traffic is forwarded using encapsulation type `vlan-vci-ccc` to a local Ethernet IQ2 and IQ2-E interface. See the topology in Figure 16 on page 240.

Figure 16: ATM-to-Ethernet Interworking

In this example, the ATM traffic comes from the DSLAM to the router on ATM interface `at-4/0/0` and is forwarded out on Ethernet interface `ge-2/2/1`.

```
[edit interfaces]
ge-2/2/1 {
  vlan-vci-tagging;
  encapsulation vlan-vci-ccc;
  unit 0 {
    encapsulation vlan-vci-ccc;
    vlan-id 100;
    inner-vlan-id-range start 100 end 500;
  }
}
at-4/0/0 {
  atm-options {
    vpi 100;
  }
  unit 0 {
    encapsulation vlan-vci-ccc;
    family ccc;
    vpi 100;
    vci-range start 100 end 500;
  }
}
```


Chapter 7

Tracing Interface Operations

You can trace the operations of individual router interfaces and those of the interface process (dcd). For a general discussion of tracing and of the precedence of multiple tracing operations, see the *JUNOS System Basics Configuration Guide*.

For information about the operations of Virtual Router Resolution Protocol (VRRP)-enabled interfaces, see the *JUNOS High Availability Configuration Guide*.

This chapter discusses the following interface trace operation configuration tasks:

- Tracing Operations of an Individual Router Interface on page 241
- Tracing Operations of the Interface Process on page 241

Tracing Operations of an Individual Router Interface

To trace the operations of individual router interfaces, include the **traceoptions** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]  
traceoptions {  
    flag flag;  
}
```

You can specify the following interface tracing flags:

- **all**—Trace all interface operations.
- **event**—Trace all interface events.
- **ipc**—Trace all interface interprocess communication (IPC) messages.
- **media**—Trace all interface media changes.

The interfaces **traceoptions** statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system **syslog** files.

Tracing Operations of the Interface Process

To trace the operations of the router's interface process, dcd, include the **traceoptions** statement at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
```

```
traceoptions {  
  file <filename> <files number> <match regular-expression> <size size>  
    <world-readable | no-world-readable>;  
  flag <flag> <disable>;  
  no-remote-trace;  
}
```

By default, interface process operations are placed in the file named `dcd` and three 1-MB files of tracing information are maintained.

You can specify the following flags in the `interfaces traceoptions` statement:

- **change-events**—Log changes that produce configuration events.
- **config-states**—Log the configuration state machine changes.
- **kernel**—Log configuration IPC messages to kernel.
- **kernel-detail**—Log details of configuration messages to kernel.

For general information about tracing, see the tracing and logging information in the *JUNOS System Basics Configuration Guide*.

Part 3

Configuring Special Router Interfaces

This section describes configuration of multiple unique interface types used for special purposes in the router.

- [Displaying the Internal Ethernet Interface on page 245](#)
- [Configuring Discard Interfaces on page 249](#)
- [Configuring IP Demultiplexing Interfaces on page 251](#)
- [Configuring the Loopback Interface on page 257](#)

Chapter 8

Displaying the Internal Ethernet Interface

This section contains the following topics:

- Displaying the Internal Ethernet Interface for M Series, MX Series, and Most T Series Routers on page 245
- Displaying Internal Ethernet Interfaces for a Routing Matrix with a TX Matrix Plus Router on page 246

Displaying the Internal Ethernet Interface for M Series, MX Series, and Most T Series Routers

The router internal Ethernet interface connects the Routing Engine with the router's packet forwarding components. The JUNOS Software automatically configures internal Ethernet interfaces. For M Series and MX Series routers and T Series routers not configured in a routing matrix with a TX Matrix Plus router, the internal Ethernet interface is `fxp1`.



NOTE: Do not modify or remove the configuration for the internal Ethernet interface that the JUNOS Software automatically configures. If you do, the router will stop functioning.

The following example shows the command output for the `show configuration` command for an M Series or MX Series router or a T Series router not configured in a routing matrix with a TX Matrix Plus router. The example shows only the portion of the `interfaces` stanza that shows the configuration of the internal Ethernet interface.

```
user@host> show configuration
...
interfaces {
  ...
  fxp1 {
    unit 0 {
      family tnp {
        address 1;
      }
    }
  }
}
```

Displaying Internal Ethernet Interfaces for a Routing Matrix with a TX Matrix Plus Router

The router internal Ethernet interface connects the Routing Engine with the router's packet forwarding components. The JUNOS Software automatically configures internal Ethernet interfaces. For TX Matrix Plus routers, the internal Ethernet interfaces are `ixgbe0` and `ixgbe1`. For T1600 routers configured in a routing matrix, the internal Ethernet interfaces are `bcm0` and `em1`. For more information about internal Ethernet interfaces, see "Permanent Interfaces" on page 32.



NOTE: Do not modify or remove the configuration for the internal Ethernet interface that the JUNOS Software automatically configures. If you do, the router will stop functioning.

The following example is a sequence of `show interfaces` commands issued in a JUNOS command-line interface (CLI) session with a TX Matrix Plus router in a routing matrix. In the example, the TX Matrix Plus router, which is also called the switch-fabric chassis (SFC), is known by the IP host name `host-sfc-0` and contains redundant Routing Engines. The commands display information about the management Ethernet interface and both internal Ethernet interfaces configured on the Routing Engine to which you are currently logged in:

```
user@host-sfc-0> show interfaces em0 terse
Interface      Admin Link Proto  Local          Remote
em0            up    up
em0.0          up    up    inet   192.168.35.95/24

user@host-sfc-0> show interfaces ixgbe0 terse
Interface      Admin Link Proto  Local          Remote
ixgbe0         up    up
ixgbe0.0       up    up    inet   10.34.0.4/8
               up    up    inet   162.0.0.4/2
               up    up    inet6  fe80::200:ff:fe22:4/64
               up    up    inet6  fec0::a:22:0:4/64
               tnp    0x22000004

user@host-sfc-0> show interfaces ixgbe1 terse
Interface      Admin Link Proto  Local          Remote
ixgbe1         up    up
ixgbe1.0       up    up    inet   10.34.0.4/8
               up    up    inet   162.0.0.4/2
               up    up    inet6  fe80::200:1ff:fe22:4/64
               up    up    inet6  fec0::a:22:0:4/64
               tnp    0x22000004
```

The following example is a sequence of `show interfaces` commands issued in a CLI session with a T1600 router in a routing matrix. In the example, the T1600 router, which is also called the line-card chassis (LCC), is known by the IP host name `host-sfc-0-lcc-2` and contains redundant Routing Engines.

This T1600 router is connected to the routing matrix through a connection in the TXP-SIB-F13 in slot 2 of the SCC. The commands display information about the management Ethernet interface and both internal Ethernet interfaces configured on the Routing Engine to which you are currently logged in:



NOTE: In a routing matrix, the **show interfaces** command displays information about the current router only. If you are logged in to the TX Matrix Plus router, the **show interfaces** command output does not include information about any of the attached T1600 routers. To display interface information about a specific T1600 router in the routing matrix, you must first log in to that router.

The previous example shows a CLI session with the TX Matrix Plus router. To display interface information about the T1600 router known as **host-sfc-0-lcc-2**, first use the **request routing-engine login** command to log in to that LCC.

```
user@host-sfc-0> request routing-engine login lcc 2
--- JUNOS 9.6I built 2009-06-22 18:13:04 UTC
% cli
warning: This chassis is a Line Card Chassis (LCC) in a multichassis system.
warning: Use of interactive commands should be limited to debugging.
warning: Normal CLI access is provided by the Switch Fabric Chassis (SFC).
warning: Please logout and log into the SFC to use CLI.
```

```
user@host-sfc-0-lcc-2> show interfaces em0 terse
Interface      Admin Link Proto  Local      Remote
em0            up    up
em0.0          up    up    inet    192.168.35.117/24
```

```
user@host-sfc-0-lcc-2> show interfaces bcm0 terse
Interface      Admin Link Proto  Local      Remote
bcm0           up    up
bcm0.0         up    up    inet    10.1.0.5/8
                                   129.0.0.5/2
                                   inet6   fe80::201:ff:fe01:5/64
                                   fec0::a:1:0:5/64
                                   tnp     0x1000005
```

```
user@host-sfc-0-lcc-2> show interfaces em1 terse
Interface      Admin Link Proto  Local      Remote
em1            up    up
em1.0          up    up    inet    10.1.0.5/8
                                   129.0.0.5/2
                                   inet6   fe80::201:1ff:fe01:5/64
                                   fec0::a:1:0:5/64
                                   tnp     0x1000005
```


Chapter 9

Configuring Discard Interfaces

On the routing platform, you can configure one physical discard interface, **dsc**. The discard interface allows you to identify the ingress point of a denial-of-service (DoS) attack. When your network is under attack, the target host IP address is identified, and the local policy forwards attacking packets to the discard interface. When traffic is routed out of the discard interface, the traffic is silently discarded.

You can configure the **inet** family protocol on the discard interface, which allows you to apply an output filter to the interface. If you apply an output filter to the interface, the action specified by the filter is executed before the traffic is discarded.

Once you configure a discard interface, you must then configure a local policy to forward attacking traffic to the discard interface. For a complete discussion about using the discard interface to protect your network against DoS attacks, see the *JUNOS Policy Framework Configuration Guide*.

To configure a discard interface, include the following statements at the **[edit interfaces]** hierarchy level:

```
[edit interfaces]
dsc {
  unit 0 {
    family inet {
      filter {
        output filter-name;
      }
      address address {
        destination address;
      }
    }
  }
}
```

Keep the following guidelines in mind when configuring the discard interface:

- Only the logical interface unit 0 is supported.
- The **filter** and **address** statements are optional.
- Although you can configure an input filter and a filter group, these configuration statements have no effect because traffic is not transmitted from the discard interface.

- The `show interface` command is not relevant for the discard interface.
- The discard interface does not support class of service (CoS).

Chapter 10

Configuring IP Demultiplexing Interfaces

This chapter contains the following topics:

- Configuring an IP Demultiplexing Interface on page 251
- Configuring an IP Demux Underlying Interface on page 252
- Specifying the Demux Underlying Interface on page 253
- Configuring IP Demux Prefixes on page 253
- Configuring MAC Address Validation on Static Demux Interfaces on page 254
- Example: Configuring a Demux Interface on page 255

Configuring an IP Demultiplexing Interface

IP demultiplexing (demux) interfaces are logical interfaces that share a common, underlying logical interface. The demux interfaces use the IP source or destination address to demultiplex received packets when the subscriber is not uniquely identified by a Layer 2 circuit.

To configure a demux interface, include the following statements at the [edit interfaces] hierarchy level:

```
[edit interfaces]
interface-name {
  unit logical-unit-number {
    ... logical-interface-configuration ...
  }
}
demux0 {
  unit logical-unit-number {
    demux-options {
      underlying-interface interface-name;
    }
    family family {
      demux-destination {
        destination-prefix;
      }
      demux-source {
        source-prefix;
      }
      mac-validate (loose | strict)
      unnumbered-address interface-name <preferred-source-address address>;
    }
  }
}
```

```
}
}
```

Keep the following guidelines in mind when configuring the demux interface:

- Demux interfaces are supported on M120 or MX Series routers only.
- You can configure only one **demux0** interface per chassis, but you can define logical demux interfaces on top of it (for example, **demux0.1**, **demux0.2**, and so on).
- You must associate demux interfaces with an underlying logical interface.



NOTE: IP demux interfaces currently support only Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet underlying interfaces.

- The demux underlying interface must reside on the same logical system as the demux interfaces that you configure over it.
- IP demux interfaces currently support only the Internet Protocol version 4 (IPv4) suite **inet** family type.
- You can configure more than one demux prefix for a given demux unit. However, you cannot configure the exact same demux prefix on two different demux units with the same underlying interface.
- You can configure overlapping demux prefixes on two different demux units with the same underlying prefix. However, under this configuration, best match rules apply (in other words, the most specific prefix wins).
- If the address in a received packet does not match any demux prefix, the packet is logically received on the underlying interface. For this reason, the underlying interface is often referred to as the “primary” interface.

Configuring an IP Demux Underlying Interface

An IP demux interface uses an underlying logical interface to receive packets.



NOTE: IP demux interfaces support only Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.

To determine which IP demux interface to use, the destination or source prefix is matched against the destination or source address of packets that the underlying interface receives. The underlying interface family type must match the demux interface prefix type.

To configure a logical interface as an IP demux underlying interface, configure the logical demultiplexing destination or source family type. To configure, include the **demux-destination** (underlying interface) statement or the **demux-source** (underlying interface) statement:

```
interfaces {
```

```

interface-name {
    unit logical-unit-number {
        (demux-destination family | demux-source family);
    }
}

```

You can include these statements at the following hierarchy levels:

- [edit]
- [edit logical-system *logical-system-name*]
- [edit logical-system *logical-system-name* routing-instances *routing-instance-name*]

Specifying the Demux Underlying Interface

You must specify an underlying interface for the demux interfaces to use. The underlying interface must reside on the same logical system as the demux interface.

To specify the logical underlying interface, include the `underlying-interface` statement:

```

interfaces {
    demux0 {
        unit logical-unit-number {
            demux-options {
                underlying-interface interface-name;
            }
        }
    }
}

```

You can include these statements at the following hierarchy levels:

- [edit]
- [edit logical-systems *logical-system-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

Configuring IP Demux Prefixes

You configure demux prefixes for use by the underlying interface. The demux prefixes can represent individual hosts or networks. For a given demux interface unit, you can configure either demux source or demux destination prefixes but not both. You can choose not to configure a demux source or demux destination prefix. This type of configuration results in a transmit-only interface.

To configure IP demux prefixes, include the `demux-destination` (demux interface) statement or the `demux-source` (demux interface) statement:

```

interfaces {
    demux0 {
        unit logical-unit-number {
            family family;

```

```

        demux-destination {
            destination-prefix;
        }
    }
    family family;
    demux-source {
        source-prefix;
    }
}
}
}

```

You can include these statements at the following hierarchy levels:

- [edit]
- [edit logical-systems *logical-system-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

Configuring MAC Address Validation on Static Demux Interfaces

MAC address validation enables the router to validate that received packets contain a trusted IP source and an Ethernet MAC source address.

MAC address validation is supported on static demux interfaces on MX Series routers only.

There are two types of MAC address validation that you can configure:

- Loose—Forwards packets when both the IP source address and the MAC source address match one of the trusted address tuples.

Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not support the MAC address of the tuple

Continues to forward packets when the source address of the incoming packet does not match any of the trusted IP addresses.

- Strict—Forwards packets when both the IP source address and the MAC source address match one of the trusted address tuples.

Drops packets when the MAC address does not match the tuple's MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.

To configure MAC address validation on static Ethernet interfaces, include the `mac-validate (loose | strict)` statement at the [edit interfaces demux0 unit *logical-unit-number* family *family*] hierarchy level.

Example: Configuring a Demux Interface

Configure two VLANs, each with two IP demux interfaces. One VLAN demultiplexes based on the source address; the other VLAN demultiplexes based on the destination address.



NOTE: This example is not intended to depict any realistic deployment; it is intended to demonstrate many possible CLI variations.

```
[edit]
interfaces {
  fe-0/0/0 {
    vlan-tagging;
    unit 100 {
      vlan-id 100;
      demux-source inet; # Enable demux of inet prefixes
      family inet {
        address 10.1.1.1/24;
        filter {
          input vlan1-primary-in-filter;
          output vlan1-primary-out-filter;
        }
        mac-validate loose;
      }
    }
    unit 200 {
      vlan-id 200;
      demux-destination inet; # Enable demux of inet using destination addresses
      family inet {
        address 20.1.1.1/24;
      }
    }
    unit 300 {
      vlan-id 300;
      demux-source inet; # Enable demux of inet using source addresses
      family inet {
        address 20.1.2.1/24;
      }
    }
  }
}
demux0 {
  unit 101 {
    description vlan1-sub1;
    demux-options {
      underlying-interface fe-0/0/0.100;
    }
    family inet {
      demux-source 10.1.1.0/24;
      filter {
        input vlan1-sub1-in-filter;
        output vlan1-sub1-out-filter;
      }
    }
  }
}
```

```

    }
    mac-validate loose;
  }
}
unit 102 {
  description vlan1-sub2;
  demux-options {
    underlying-interface fe-0/0/0.100;
  }
  family inet {
    demux-source {
      10.1.0.0/16;
      10.2.1.0/24;
    }
    filter {
      input vlan1-sub2-in-filter;
      output vlan1-sub2-out-filter;
    }
    mac-validate loose;
  }
}
unit 202 {
  description vlan2-sub2;
  demux-options {
    underlying-interface fe-0/0/0.200;
  }
  family inet {
    demux-destination 100.1.2.0/24;
  }
}
unit 302 {
  description vlan2-sub2;
  demux-options {
    underlying-interface fe-0/0/0.300;
  }
  family inet {
    demux-source 100.1.2.0/24;
  }
}
}
}

```


Chapter 11

Configuring the Loopback Interface

On the router, you can configure one physical loopback interface, `lo0`, and one or more addresses on the interface.

- Configuring the Loopback Interface on page 257

Configuring the Loopback Interface

To configure the physical loopback interface, include the following statements at the `[edit interfaces]` hierarchy level:

```
[edit interfaces]
lo0 {
  unit 0 {
    family inet {
      address loopback-address;
      address <loopback-address2>;
      ...
    }
    family inet6 {
      address loopback-address;
    }
  }
}
```

When specifying the loopback address, do not include a destination prefix. Also, in most cases, do not specify a loopback address on any unit other than unit 0.



NOTE: For Layer 3 virtual private networks (VPNs), you can configure multiple logical units for the loopback interface. This allows you to configure a logical loopback interface for each virtual routing and forwarding (VRF) routing instance. For more information, see the *JUNOS VPNs Configuration Guide*.

For some applications, such as SSL for JUNOScript, the address for the interface `lo0.0` must be `127.0.0.1`.

You can configure loopback interfaces using a subnetwork address for both `inet` and `inet6` address families. Many protocols require a subnetwork address as their source address. Configuring a subnetwork loopback address as a donor interface enables these protocols to run on unnumbered interfaces.

If you configure the loopback interface, it is automatically used for unnumbered interfaces. If you do not configure the loopback interface, the router chooses the first interface to come online as the default. If you configure more than one address on the loopback interface, we recommend that you configure one to be the primary address to ensure that it is selected for use with unnumbered interfaces. By default, the primary address is used as the source address when packets originate from the interface.

For more information about unnumbered interfaces, see “Configuring an Unnumbered Interface” on page 185. For more information about primary addresses, see “Configuring the Interface Address” on page 174.

Example: Configuring the Loopback Interface

Configure two addresses on the loopback interface with host routes:

```
[edit]
user@host# edit interfaces lo0 unit 0 family inet
[edit interfaces lo0 unit 0 family inet]
user@host# set address 172.16.0.1
[edit interfaces lo0 unit 0 family inet]
user@host# set address 10.0.0.1
[edit interfaces lo0 unit 0 family inet]
user@host# top
[edit]
user@host# show
interfaces {
  lo0 {
    unit 0 {
      family inet {
        10.0.0.1;
        127.0.0.1;
        172.16.0.1;
      }
    }
  }
}
```

Configure two addresses on the loopback interface with subnetwork routes:

```
[edit]
user@host# edit interfaces lo0 unit 0 family inet
[edit interfaces lo0 unit 0 family inet]
user@host# set address 192.16.0.1/24
[edit interfaces lo0 unit 0 family inet]
user@host# set address 10.2.0.1/16
[edit interfaces lo0 unit 0 family inet]
user@host# top
[edit]
user@host# show
interfaces {
  lo0 {
    unit 0 {
      family inet {
        10.2.0.1/16;

```

```

        127.0.0.1/32;
        192.16.0.1/24;
    }
}
}

```

Configure an IP and an IPv6 address on the loopback interface with subnetwork routes:

```

[edit]
user@host# edit interfaces lo0 unit 0 family inet
[edit interfaces lo0 unit 0 family inet]
user@host# set address 192.16.0.1/24
[edit interfaces lo0 unit 0 family inet]
user@host# up
[edit interfaces lo0 unit 0 family]
user@host# edit interfaces lo0 unit 0 family inet6
[edit interfaces lo0 unit 0 family inet6]
user@host# set address 3ffe::1:200:f8ff:fe75:50df/64
[edit interfaces lo0 unit 0 family inet6]
user@host# top
[edit]
user@host# show
interfaces {
  lo0 {
    unit 0 {
      family inet {
        127.0.0.1/32;
        192.16.0.1/24;
      }
      family inet6 {
        3ffe::1:200:f8ff:fe75:50df/64;
      }
    }
  }
}

```


Part 4

Configuring Serial Interfaces

This section describes configuration of serial interfaces.

- Configuring Serial Interfaces on page 263

Chapter 12

Configuring Serial Interfaces

This chapter discusses configuration of the following serial interface properties:

- Serial Interfaces Overview on page 263
- Physical Interface Configuration Statements for Serial Interfaces on page 264
- Configuring the Serial Line Protocol on page 265
- Configuring the Serial Clocking Mode on page 269
- Configuring the Serial Idle Cycle Flag on page 271
- Configuring the Serial Signal Handling on page 271
- Configuring the Serial DTR Circuit on page 274
- Configuring Serial Signal Polarities on page 274
- Configuring Serial Loopback Capability on page 275
- Configuring Serial Line Encoding on page 277

Serial Interfaces Overview

Devices that communicate over a serial interface are divided into two classes: data terminal equipment (DTE) and data circuit-terminating equipment (DCE). Juniper Networks Serial Physical Interface Cards (PICs) have two ports per PIC and support full-duplex data transmission. These PICs support DTE mode only. On the Serial PIC, you can configure three types of serial interfaces:

- EIA-530—An Electronics Industries Alliance (EIA) standard for the interconnection of DTE and DCE using serial binary data interchange with control information exchanged on separate control circuits.
- V.35—An ITU-T standard describing a synchronous, physical layer protocol used for communications between a network access device and a packet network. V.35 is most commonly used in the United States and in Europe.
- X.21—An ITU-T standard for serial communications over synchronous digital lines. The X.21 protocol is used primarily in Europe and Japan.

The following standards apply to serial interfaces:

- TIA/EIA Standard 530, *High-Speed 25-Position Interface for Data Terminal Equipment and Data Circuit-Terminating Equipment*, defines the signals on the cable and specifies the connector at the end of the cable.
- TIA/EIA Standard 232, *Interface between Data Terminal Equipment and Data Circuit-Terminating Equipment Employing Serial Binary Data Interchange*, describes the physical interface and protocol for serial data communication.
- ITU-T Recommendation V.35, *Data Transmission at 48 kbit/s Using 60-108 kHz Group Band Circuits*. Note that the Juniper Networks Serial PIC supports V.35 interfaces with speeds higher than 48 kilobits per second.
- ITU-T Recommendation X.21, *Interface Between Data Terminal Equipment and Data Circuit-Terminating Equipment for Synchronous Operation on Public Data Networks*.

There are no serial interface-specific logical properties. For information about general logical properties that you can configure, see “Configuring Logical Interface Properties” on page 143. On J Series routers, link fragmentation and interleaving (LFI) and Multilink Point-to-Point Protocol (MLPPP) support has been extended to serial interfaces. This support on serial interfaces is the same as the existing LFI and MLPPP support on T1 and E1 interfaces.

Physical Interface Configuration Statements for Serial Interfaces

To configure serial physical interface properties, include the `serial-options` statement for the J Series router at the `[edit interfaces se-pim/O/port]` hierarchy level or at the `[edit interfaces se-fpc/pic/port]` hierarchy level for M Series and T Series routers:

```
[edit interfaces se-fpc/pic/port]
serial-options {
  clock-rate rate;
  clocking-mode (dce | internal | loop);
  control-polarity (negative | positive);
  cts-polarity (negative | positive);
  dcd-polarity (negative | positive);
  dce-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
  }
  dsr-polarity (negative | positive);
  dte-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
```



```

        indication (ignore | normal | require);
        rts (assert | de-assert | normal);
        tm (ignore | normal | require);
    }
    dtr-circuit (negative | positive);
    dtr-polarity (negative | positive);
    encoding (nrz | nrzi);
    idle-cycle-flag flag;
    indication-polarity (negative | positive);
    line-protocol protocol;
    loopback mode;
    rts-polarity (negative | positive);
    tm-polarity (negative | positive);
    transmit-clock invert;
}

```

Configuring the Serial Line Protocol

By default, serial interfaces use the EIA-530 line protocol. You can configure each port on the PIC independently to use one of the following line protocols:

- EIA-530
- V.35
- X.21

To configure the serial line protocol, include the **line-protocol** statement, specifying the **eia530**, **v.35**, or **x.21** option:

```
line-protocol protocol;
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *se-pim*/*0/port* serial-options]
- [edit interfaces *se-fpc/pic/port* serial-options]

For more information about serial interfaces, see the following sections:

- Serial Interface Default Settings on page 265
- Invalid Serial Interface Statements on page 267

Serial Interface Default Settings

The following sections show the default settings for serial interfaces:

- EIA-530 Interface Default Settings on page 266
- V.35 Interface Default Settings on page 266
- X.21 Interface Default Settings on page 267

EIA-530 Interface Default Settings

If you do not include the `line-protocol` statement or if you explicitly configure the default EIA-530 line protocol, the default settings are as follows:

```
dce-options | dte-options {
  cts normal;
  dcd normal;
  dsr normal;
  dtr normal;
  rts normal;
  tm normal;
}
clock-rate 16.384mhz;
clocking-mode loop;
cts-polarity positive;
dcd-polarity positive;
dsr-polarity positive;
dtr-circuit balanced;
dtr-polarity positive;
encoding nrz;
rts-polarity positive;
tm-polarity positive;
```



NOTE: On M Series routers, you can set the DCE clocking mode for EIA-530 interfaces and commit. An error message is not displayed and the CLI is not blocked.

You can include the `line-protocol` statement at the following hierarchy levels:

- [edit interfaces *se-pim*/0/*port* serial-options]
- [edit interfaces *se-fpc*/*pic*/*port* serial-options]

V.35 Interface Default Settings

If you include the `line-protocol v.35` statement, the default settings are as follows:

```
dce-options | dte-options {
  cts normal;
  dcd normal;
  dsr normal;
  dtr normal;
  rts normal;
}
clock-rate 16.384mhz;
clocking-mode loop;
cts-polarity positive;
dcd-polarity positive;
dsr-polarity positive;
dtr-circuit balanced;
dtr-polarity positive;
encoding nrz;
```

```
rts-polarity positive;
```

You can include the `line-protocol` statement at the following hierarchy levels:

- [edit interfaces *se-pim/0/port* serial-options]
- [edit interfaces *se-fpc/pic/port* serial-options]

X.21 Interface Default Settings

If you include the `line-protocol x.21` statement, the default settings are as follows:

```
dce-options | dte-options {
  control-signal normal;
  indication normal;
}
clock-rate 16.384mhz;
clocking-mode loop;
control-polarity positive;
encoding nrz;
indication-polarity positive;
```

You can include the `line-protocol` statement at the following hierarchy levels:

- [edit interfaces *se-pim/0/port* serial-options]
- [edit interfaces *se-fpc/pic/port* serial-options]

Invalid Serial Interface Statements

The following sections show the invalid configuration statements for each type of serial interface. If you include the following statements in the configuration, an error message indicates the location of the error and the configuration is not activated.

- Invalid EIA-530 Interface Statements on page 267
- Invalid V.35 interface Statements on page 268
- Invalid X.21 Interface Statements on page 268

Invalid EIA-530 Interface Statements

If you do not include the `line-protocol` statement or if you explicitly configure the default EIA-530 line protocol, the following statements are invalid:

```
dce-options | dte-options {
  control-signal (assert | de-assert | normal);
  indication (ignore | normal | require);
}
control-polarity (negative | positive);
indication-polarity (negative | positive);
```

You can include the `line-protocol` statement at the following hierarchy levels:

- [edit interfaces *se-pim/0/port* serial-options]

- [edit interfaces *se-fpc/pic/port* serial-options]

Invalid V.35 Interface Statements

If you include the line-protocol v.35 statement, the following statements are invalid:

```
dce-options | dte-options {
  control-signal (assert | de-assert | normal);
  indication (ignore | normal | require);
  tm (ignore | normal | require);
}
control-polarity (negative | positive);
indication-polarity (negative | positive);
loopback (dce-local | dce-remote);
tm-polarity (negative | positive);
```

You can include the line-protocol statement at the following hierarchy levels:

- [edit interfaces *se-pim/0/port* serial-options]
- [edit interfaces *se-fpc/pic/port* serial-options]

Invalid X.21 Interface Statements

If you include the line-protocol x.21 statement, the following statements are invalid:

```
dce-options | dte-options {
  cts (ignore | normal | require);
  dcd (ignore | normal | require);
  dsr (ignore | normal | require);
  dtr (assert | de-assert | normal);
  rts (assert | de-assert | normal);
  tm (ignore | normal | require);
}
clocking-mode (dce | internal);
cts-polarity (negative | positive);
dce-polarity (negative | positive);
dsr-polarity (negative | positive);
dtr-circuit (balanced | unbalanced);
dtr-polarity (negative | positive);
loopback (dce-local | dce-remote);
rts-polarity (negative | positive);
tm-polarity (negative | positive);
```

You can include the line-protocol statement at the following hierarchy levels:

- [edit interfaces *se-pim/0/port* serial-options]
- [edit interfaces *se-fpc/pic/port* serial-options]

Configuring the Serial Clocking Mode

By default, serial interfaces use loop clocking mode. For EIA-530 and V.35 interfaces, you can configure each port on the PIC independently to use loop, DCE, or internal clocking mode. For X.21 interfaces, only loop clocking mode is supported.

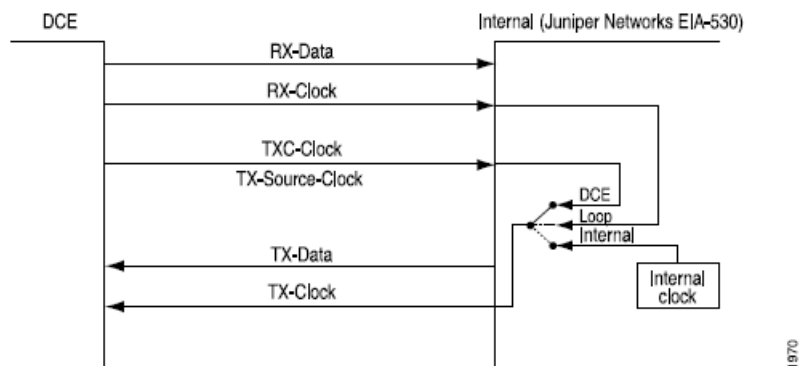
The three clocking modes work as follows:

- Loop clocking mode—Uses the DCE's RX clock to clock data from the DCE to the DTE.
- DCE clocking mode—Uses the TXC clock, which is generated by the DCE specifically to be used by the DTE as the DTE's transmit clock.
- Internal clocking mode—Also known as line timing, uses an internally generated clock. You can configure the speed of this clock by including the `clock-rate` statement at the `[edit interfaces se-pim/0/port serial-options]` or `[edit interfaces se-fpc/pic/port dte-options]` hierarchy levels. For more information about the DTE clock rate, see “Configuring the DTE Clock Rate” on page 270.

Note that DCE clocking mode and loop clocking mode use external clocks generated by the DCE.

Figure 17 on page 269 shows the clock sources of loop, DCE, and internal clocking modes.

Figure 17: Serial Interface Clocking Mode



To configure the clocking mode of a serial interface, include the `clocking-mode` statement:

```
clocking-mode (dce | internal | loop);
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces se-pim/0/port serial-options]`
- `[edit interfaces se-fpc/pic/port serial-options]`

For more information about clocking on serial interfaces, see the following sections:

- Inverting the Serial Interface Transmit Clock on page 270
- Configuring the DTE Clock Rate on page 270

Inverting the Serial Interface Transmit Clock

When an externally timed clocking mode (DCE or loop) is used, long cables might introduce a phase shift of the DTE-transmitted clock and data. At high speeds, this phase shift might cause errors. Inverting the transmit clock corrects the phase shift, thereby reducing error rates.

By default, the transmit clock is not inverted. To invert the transmit clock, include the `transmit-clock invert` statement:

```
transmit-clock invert;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *se-pim*/0/*port* serial-options]
- [edit interfaces *se-fpc*/*pic*/*port* serial-options]

Configuring the DTE Clock Rate

By default, the serial interface has a clock rate of 16.384 MHz. For EIA-530 and V.35 interfaces with internal clocking mode configured, you can configure the clock rate. For more information about internal clocking mode, see “Configuring the Serial Clocking Mode” on page 269.

To configure the clock rate, include the `clock-rate` statement:

```
clock-rate rate;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *se-pim*/0/*port* serial-options]
- [edit interfaces *se-fpc*/*pic*/*port* serial-options]

You can configure the following interface speeds:

- 2.048 MHz
- 2.341 MHz
- 2.731 MHz
- 3.277 MHz
- 4.096 MHz
- 5.461 MHz
- 8.192 MHz
- 16.384 MHz

Although the serial interface is intended for use at the default rate of 16.384 MHz, you might need to use a slower rate if any of the following conditions prevail:

- The interconnecting cable is too long for effective operation.
- The interconnecting cable is exposed to an extraneous noise source that might cause an unwanted voltage in excess of + 1 volt measured differentially between the signal conductor and circuit common at the load end of the cable, with a 50-ohm resistor substituted for the generator.
- You need to minimize interference with other signals.
- You need to invert signals.

For detailed information about the relationship between signaling rate and interface cable distance, see the following standards:

- EIA-422-A, *Electrical Characteristics of Balanced Voltage Digital Interface Circuits*
- EIA-423-A, *Electrical Characteristics of Unbalanced Voltage Digital Interface Circuits*

Configuring the Serial Idle Cycle Flag

By default, a serial interface on J Series routers transmits the value 0x7E in the idle cycles. To have the interface transmit the value 0xFF (all ones) instead, include the `idle-cycle-flag` statement at the `[edit interfaces interface-name serial-options]` hierarchy level, specifying the `ones` option:

```
[edit interfaces interface-name serial-options]
idle-cycle-flag ones;
```

To explicitly configure the default value of 0x7E, include the `idle-cycle-flag` statement with the `flags` option:

```
[edit interfaces interface-name serial-options]
idle-cycle-flag flags;
```

Configuring the Serial Signal Handling

By default, normal signal handling is enabled for all signals. For each signal, the `normal` option applies to the normal signal handling for that signal, as defined by the following standards:

- TIA/EIA Standard 530
- ITU-T Recommendation V.35
- ITU-T Recommendation X.21

Table 24 on page 272 shows the serial interface modes that support each signal type.

Table 24: Signal Handling by Serial Interface Type

Signal	Serial Interfaces
From-DCE signals	
Clear to send (CTS)	EIA-530 and V.35
Data carrier detect (DCD)	EIA-530 and V.35
Data set ready (DSR)	EIA-530 and V.35
Indication	X.21 only
Test mode (TM)	EIA-530 only
To-DCE signals	
Control signal	X.21 only
Data transfer ready (DTR)	EIA-530 and V.35
Request to send (RTS)	EIA-530 and V.35

You configure serial interface signal characteristics by including the **dce-options** or **dte-options** statement:

```
dce-options | dte-options {
  control-signal (assert | de-assert | normal);
  cts (ignore | normal | require);
  dcd (ignore | normal | require);
  dsr (ignore | normal | require);
  dtr signal-handling-option;
  ignore-all;
  indication (ignore | normal | require);
  rts (assert | de-assert | normal);
  tm (ignore | normal | require);
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *se-pim/0/port* serial-options]
- [edit interfaces *se-fpc/pic/port* serial-options]

For EIA-530 and V.35 interfaces, configure to-DCE signals by including the **dtr** and **rts** statements, specifying the **assert**, **de-assert**, or **normal** option:

```
dtr (assert | de-assert | normal);
rts (assert | de-assert | normal);
```

For X.21 interfaces, configure to-DCE signals by including the **control-signal** statement, specifying the **assert**, **de-assert**, or **normal** option:

```
control-signal (assert | de-assert | normal);
```


Assertion is when the positive side of a given signal is at potential high-level output voltage (Voh), while the negative side of the same signal is at potential low-level output voltage (Vol). *Deassertion* is when the positive side of a given signal is at potential Vol, while the negative side of the same signal is at potential Voh.

For the DTR signal, you can configure normal signal handling using the signal for automatic resynchronization by including the **dtr** statement, and specifying the **auto-synchronize** option:

```
dtr {
  auto-synchronize {
    duration milliseconds;
    interval seconds;
  }
}
```

The pulse duration of resynchronization can be from 1 through 1000 milliseconds. The offset interval for resynchronization can be from 1 through 31 seconds.

For EIA-530 and V.35 interfaces, configure from-DCE signals by including the **cts**, **dcd**, and **dsr** statements, specifying the **ignore**, **normal**, or **require** option:

```
cts (ignore | normal | require);
dcd (ignore | normal | require);
dsr (ignore | normal | require);
```

For X.21 interfaces, configure from-DCE signals by including the **indication** statement, specifying the **ignore**, **normal**, or **require** option:

```
indication (ignore | normal | require);
```

For EIA-530 interfaces only, you can configure from-DCE test-mode (TM) signaling by including the **tm** statement, specifying the **ignore**, **normal**, or **require** option:

```
tm (ignore | normal | require);
```

To specify that the from-DCE signal must be asserted, include the **require** option in the configuration. To specify that the from-DCE signal must be ignored, include the **ignore** option in the configuration.



NOTE: For V.35 and X.21 interfaces, you cannot include the **tm** statement in the configuration.

For X.21 interfaces, you cannot include the **cts**, **dcd**, **dsr**, **dtr**, and **rts** statements in the configuration.

For EIA-530 and V.35 interfaces, you cannot include the **control-signal** and **indication** statements in the configuration.

For a complete list of serial options statements that are not supported by each serial interface mode, see “Invalid Serial Interface Statements” on page 267.

To return to the default normal signal handling, delete the **require**, **ignore**, **assert**, **de-assert**, or **auto-synchronize** statement from the configuration, as shown in the following example:

```
[edit]
user@host# delete interfaces se-fpc/pic/port dte-options control-leads cts require
```

To explicitly configure normal signal handling, include the **control-signal** statement with the **normal** option:

```
control-signal normal;
```

You can configure the serial interface to ignore all control leads by including the **ignore-all** statement:

```
ignore-all;
```

You can include the **ignore-all** statement in the configuration only if you do not explicitly enable other signal handling options at the **[edit interfaces se-pim/0/port serial-options dce-options]** or **[edit interfaces se-fpc/pic/port serial-options dte-options]** hierarchy levels.

You can include the **control-signal**, **cts**, **dcd**, **dsr**, **dtr**, **indication**, **rts**, and **tm** statements at the following hierarchy levels:

- **[edit interfaces se-pim/0/port serial-options dte-options]**
- **[edit interfaces se-fpc/pic/port serial-options dte-options]**

Configuring the Serial DTR Circuit

A balanced circuit has two currents that are equal in magnitude and opposite in phase. An unbalanced circuit has one current and a ground; if a pair of terminals is unbalanced, one side is connected to electrical ground and the other carries the signal. By default, the DTR circuit is balanced.

For EIA-530 and V.35 interfaces, configure the DTR circuit by including the **dtr-circuit** statement:

```
dtr-circuit (balanced | unbalanced);
```

You can include the **dtr-circuit** statement at the following hierarchy levels:

- **[edit interfaces se-pim/0/port serial-options]**
- **[edit interfaces se-fpc/pic/port serial-options]**

Configuring Serial Signal Polarities

Serial interfaces use a differential protocol signaling technique. Of the two serial signals associated with a circuit, the one referred to as the A signal is denoted with a plus sign, and the one referred to as the B signal is denoted with a minus sign; for

example, DTR + and DTR-. If DTR is low, then DTR + is negative with respect to DTR-. If DTR is high, then DTR + is positive with respect to DTR-.

By default, all signal polarities are positive. You can reverse this polarity on a Juniper Networks serial interface. You might need to do this if signals are miswired as a result of reversed polarities.

For EIA-530 and V.35 interfaces, configure signal polarities by including the `cts-polarity`, `dcd-polarity`, `dsr-polarity`, `dtr-polarity`, `rts-polarity`, and `tm-polarity` statements:

```
cts-polarity (negative | positive);
dcd-polarity (negative | positive);
dsr-polarity (negative | positive);
dtr-polarity (negative | positive);
rts-polarity (negative | positive);
tm-polarity (negative | positive);
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *se-pim/0/port* serial-options]
- [edit interfaces *se-fpc/pic/port* serial-options]

For X.21 interfaces, configure signal polarities by including the `control-polarity` and `indication-polarity` statements:

```
control-polarity (negative | positive);
indication-polarity (negative | positive);
```

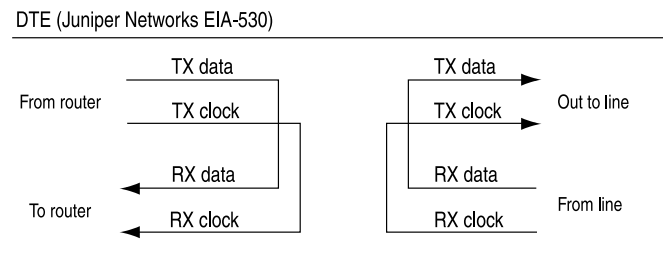
You can include these statements at the following hierarchy levels:

- [edit interfaces *se-pim/0/port* serial-options]
- [edit interfaces *se-fpc/pic/port* serial-options]

Configuring Serial Loopback Capability

From the router, remote line interface unit (LIU) loopback loops the TX (transmit) data and TX clock back to the router as RX (receive) data and RX clock. From the line, LIU loopback loops the RX data and RX clock back out the line as TX data and TX clock, as shown in Figure 18 on page 275.

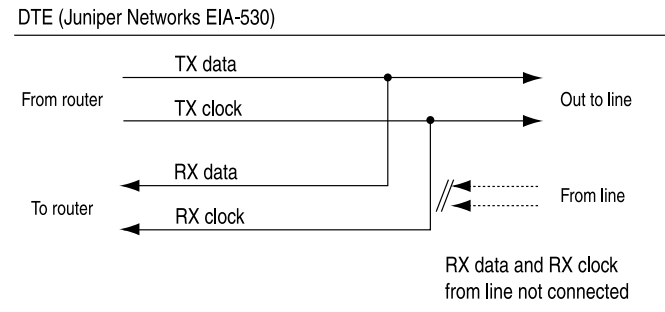
Figure 18: Serial Interface LIU Loopback



1972

DCE local and DCE remote control the EIA-530 interface-specific signals for enabling local and remote loopback on the link partner DCE. Local loopback is shown in Figure 19 on page 276.

Figure 19: Serial Interface Local Loopback



1971

For EIA-530 interfaces, you can configure DCE local, DCE remote, local, and remote (LIU) loopback capability.

For V.35, you can configure remote LIU and local loopback capability. DCE local and DCE remote loopbacks are not supported on V.35 and X.21 interfaces. Local and remote loopbacks are not supported on X.21 interfaces.

To configure the loopback capability on a serial interface, include the **loopback** statement, specifying the **dce-local**, **dce-remote**, **local**, or **remote** option:

```
loopback mode;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *se-pim/0/port* serial-options]
- [edit interfaces *se-fpc/pic/port* serial-options]

To disable the loopback capability, remove the **loopback** statement from the configuration:

```
[edit]
user@host# delete interfaces se-fpc/pic/port serial-options loopback
```

You can determine whether there is an internal or external problem by checking the error counters in the output of the **show interface se-fpc/pic/port extensive** command:

```
user@host> show interfaces se-fpc/pic/port extensive
```

Example: Configuring Serial Loopback Capability

To determine the source of a problem, loop packets on the local router, the local DCE, the remote DCE, and the remote line interface unit (LIU). To do this, include the **no-keepalives** and **encapsulation cisco-hdlc** statements at the [edit interfaces *se-fpc/pic/port*] hierarchy level, and the **loopback local** option at the [edit interfaces *se-pim/0/port* serial-options] or [edit interfaces *se-fpc/pic/port* serial-options] hierarchy

level. With this configuration, the link stays up, so you can loop ping packets to a remote router. The **loopback local** statement causes the interface to loop within the PIC just before the data reaches the transceiver.

```
[edit interfaces]
se-1/0/0 {
  no-keepalives;
  encapsulation cisco-hdlc;
  serial-options {
    loopback local;
  }
  unit 0 {
    family inet {
      address 10.100.100.1/24;
    }
  }
}
```

Configuring Serial Line Encoding

By default, serial interfaces use non-return to zero (NRZ) line encoding. You can configure non-return to zero inverted (NRZI) line encoding if necessary.

To have the interface use NRZI line encoding, include the **encoding** statement, specifying the **nrzi** option:

```
encoding nrzi;
```

To explicitly configure the default NRZ line encoding, include the **encoding** statement, specifying the **nrz** option:

```
encoding nrz;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *se-pim*/0/*port* serial-options]
- [edit interfaces *se-fpc/pic*/0/*port* serial-options]

When setting the line encoding parameter, you must set the same value for paired ports. Ports 0 and 1 must share the same value.

Part 5

Configuring ATM Interfaces

This part describes the configuration of the following ATM interfaces:

- Configuring ATM Interfaces on page 281
- Configuring ATM-over-ADSL Interfaces on page 355
- Configuring ATM-over-SHDSL Interfaces on page 361

Chapter 13

Configuring ATM Interfaces

This section contains the following topics:

- ATM Interfaces Overview on page 282
- ATM1 Physical and Logical Configuration Statement Hierarchies on page 283
- ATM2 IQ Physical and Logical Configuration Statement Hierarchies on page 285
- Supported Features on ATM1 and ATM2 IQ Interfaces on page 287
- Configuring Communication with Directly Attached ATM Switches and Routers on page 291
- Enabling ILMI for Cell Relay on page 292
- Enabling Passive Monitoring on ATM Interfaces on page 293
- Removing MPLS Labels from Incoming Packets on page 294
- Configuring the ATM PIC Type on page 295
- Configuring ATM Cell-Relay Promiscuous Mode on page 296
- Configuring the Maximum Number of ATM1 VCs on a VP on page 300
- Configuring Layer 2 Circuit Transport Mode on page 300
- Configuring Layer 2 Circuit Cell-Relay Promiscuous Mode on page 308
- Configuring Layer 2 Circuit Trunk Mode Scheduling on page 309
- Configuring CoS Queues in Layer 2 Circuit Trunk Mode on page 311
- Configuring the Layer 2 Circuit Cell-Relay Cell Maximum on page 313
- Configuring the OAM F4 Cell Flows on page 315
- Defining Virtual Path Tunnels on page 316
- Configuring a Point-to-Point ATM1 or ATM2 IQ Connection on page 316
- Configuring a Point-to-Multipoint ATM1 or ATM2 IQ Connection on page 317
- Configuring a Multicast-Capable ATM1 or ATM2 IQ Connection on page 318
- Configuring Inverse ATM1 or ATM2 ARP on page 318
- Defining the ATM Traffic-Shaping Profile on page 319
- Configuring the ATM1 Queue Length on page 325
- Configuring the ATM2 IQ EPD Threshold on page 326
- Configuring Two EPD Thresholds per Queue on page 328
- Configuring the ATM2 IQ Transmission Weight on page 329

- Defining the ATM OAM F5 Loopback Cell Period on page 329
- Configuring the ATM OAM F5 Loopback Cell Threshold on page 329
- Configuring ATM Interface Encapsulation on page 330
- Configuring an ATM1 Cell-Relay Circuit on page 332
- Configuring PPP over ATM2 Encapsulation on page 334
- Configuring E3 and T3 Parameters on ATM Interfaces on page 337
- Configuring SONET/SDH Parameters on ATM Interfaces on page 338
- Configuring ATM2 IQ VC Tunnel CoS Components on page 339
- Example: Configuring ATM1 Interfaces on page 350
- Example: Configuring ATM2 IQ Interfaces on page 352

ATM Interfaces Overview

Asynchronous Transfer Mode (ATM) is a network protocol designed to facilitate the simultaneous handling of various types of traffic streams (voice, data, and video) at very high speeds over the same physical connection. By always using 53-byte cells, ATM simplifies the design of hardware, enabling it to quickly determine the destination address of each cell. This allows simple switching of network traffic at much higher speeds than are easily accomplished using protocols with variable sizes of transfer units, such as Frame Relay and Transmission Control Protocol/Internet Protocol (TCP/IP).

Although ATM was designed to operate without the requirement of any other networking protocol, other protocols are frequently segmented and encapsulated across multiple, smaller ATM cells. This makes ATM a transport mechanism for pre-existing technologies such as Frame Relay and the TCP/IP family of protocols.

ATM relies on the concepts of virtual paths and virtual circuits. A virtual path, represented by a specific virtual path identifier (VPI), establishes a route between two devices in a network. Each VPI can contain multiple virtual circuits, each represented by a virtual circuit identifier (VCI).

VPIs and VCIs are local to the router, which means that only the two devices connected by the VCI or VPI need know the details of the connection. In a typical ATM network, user data might traverse multiple connections, using many different VPI and VCI connections. Each end device, just like each device in the network, needs to know only the VCI and VPI information for the path to the next device.



NOTE: The ATM three-bit payload type identifier (PTI) field is not supported.

With ATM2 intelligent queuing (IQ) interfaces, you can configure virtual path (VP) shaping and Operation, Administration, and Management (OAM) F4 cell flows.

ATM1 Physical and Logical Configuration Statement Hierarchies

To configure ATM1 physical interface properties, include the `atm-options`, `e3-options`, `t3-options`, and `sonet-options` statements at the [edit interfaces *at-fpc/pic/port*] hierarchy level:

ATM1 Physical Configuration Hierarchy

```
[edit interfaces at-fpc/pic/port]
atm-options {
  ilmi;
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  pic-type atm1;
  promiscuous-mode {
    vpi vpi-identifier;
  }
  vpi vpi-identifier {
    maximum-vcs maximum-vcs;
  }
}
e3-options {
  atm-encapsulation (direct | plcp);
  buildout feet;
  framing (g.751 | g.832);
  loopback (local | remote);
  (payload-scrambler | no-payload-scrambler);
}
encapsulation (atm-ccc-cell-relay | atm-pvc | ethernet-over-atm);
sonet-options {
  aps {
    advertise-interval milliseconds;
    authentication-key key;
    force;
    hold-time milliseconds;
    lockout;
    neighbor address;
    paired-group group-name;
    protect-circuit group-name;
    request;
    revert-time seconds;
    working-circuit group-name;
  }
  bytes {
    e1-quiet value;
    f1 value;
    f2 value;
    s1 value;
    z3 value;
    z4 value;
  }
  loopback (local | remote);
  (payload-scrambler | no-payload-scrambler);
}
```

```

rfc-2615;
trigger {
    defect ignore {
        hold-time up milliseconds down milliseconds;
    }
}
(z0-increment | no-z0-increment);
}
t3-options {
    atm-encapsulation (direct | plcp);
    buildout feet;
    (cbit-parity | no-cbit-parity);
    loopback (local | payload | remote);
    (payload-scrambler | no-payload-scrambler);
}

```

To configure ATM1 logical interface properties, include the following statements:

ATM1 Logical Configuration Hierarchy

```

allow-any-vci;
multicast-vci vpi-identifier.vci-identifier;
oam-liveness {
    up-count cells;
    down-count cells;
}
oam-period (disable | seconds);
shaping {
    (cbr rate | vbr peak rate sustained rate burst length);
    queue-length number;
}
vci vpi-identifier.vci-identifier;
vpi vpi-identifier;
family inet {
    address address {
        multipoint-destination address {
            inverse-arp;
            oam-liveness {
                up-count cells;
                down-count cells;
            }
            oam-period (disable | seconds);
            shaping {
                (cbr rate | vbr peak rate sustained rate burst length);
                queue-length number;
            }
        }
        vci vpi-identifier.vci-identifier;
    }
}
}

```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

ATM2 IQ Physical and Logical Configuration Statement Hierarchies

To configure ATM2 IQ physical interface properties, include the `atm-options` and `sonet-options` statements at the [edit interfaces *at-fpc/pic/port*] hierarchy level:

ATM2 IQ Physical Configuration Hierarchy

```
[edit interfaces at-fpc/pic/port]
atm-options {
  cell-bundle-size cells;
  ilmi;
  linear-red-profiles profile-name{
    high-plp-max-threshold percent;
    low-plp-max-threshold percent;
    queue-depth cells high-plp-threshold percent low-plp-threshold percent;
  }
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  pic-type atm2;
  plp-to-clp;
  promiscuous-mode {
    vpi vpi-identifier;
  }
  scheduler-maps map-name {
    forwarding-class class-name {
      epd-threshold cells plp1 cells;
      linear-red-profile profile-name;
      priority (high | low);
      transmit-weight (cells number | percent number);
    }
    vc-cos-mode (alternate | strict);
  }
  vpi vpi-identifier {
    oam-liveness {
      up-count;
      down-count;
    }
    oam-period (disable | seconds);
    shaping {
      (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
        rate burst length);
    }
  }
}
sonet-options {
  aps {
    advertise-interval milliseconds;
    authentication-key key;
    force;
    hold-time milliseconds;
    lockout;
    neighbor address;
    paired-group group-name;
  }
}
```

```

    protect-circuit group-name;
    request;
    revert-time seconds;
    working-circuit group-name;
}
bytes {
    e1-quiet value;
    f1 value;
    f2 value;
    s1 value;
    z3 value;
    z4 value;
}
loopback (local | remote);
(payload-scrambler | no-payload-scrambler);
rfc-2615;
trigger {
    defect ignore {
        hold-time up milliseconds down milliseconds;
    }
}
(z0-increment | no-z0-increment);
}

```

To configure ATM2 IQ logical interface properties, include the following statements:

ATM2 IQ Logical Configuration Hierarchy

```

allow-any-vci;
atm-scheduler-map (map-name | default);
cell-bundle-size cells;
epd-threshold cells;
multicast-vci vpi-identifier.vci-identifier;
oam-liveness {
    up-count cells;
    down-count cells;
}
oam-period (disable | seconds);
plp-to-clp;
shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
    burst length);
}
transmit-weight number;
trunk-id number;
vci vpi-identifier.vci-identifier;
vpi vpi-identifier;
family inet address address {
    multipoint-destination address;
    epd-threshold cells;
    inverse-arp;
    oam-liveness {
        up-count cells;
        down-count cells;
    }
}
oam-period (disable | seconds);
shaping {

```

```

        (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
         rate burst length);
    }
    transmit-weight number;
    vci vpi-identifier.vci-identifier;
}

```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Supported Features on ATM1 and ATM2 IQ Interfaces

Table 25 on page 287 lists the supported features on ATM1 and ATM2 IQ interfaces.

Table 25: ATM1 and ATM2 IQ Supported Features

Item	ATM1	ATM2 IQ	Comments
Encapsulation and Transport Modes			
ATM Adaptation Layer 5 (AAL5) circuit cross-connect (CCC)	Supported	Supported	For ATM1 and ATM2 IQ Physical Interface Cards (PICs), you can configure any combination of AAL5 CCC, nonpromiscuous cell relay, and AAL5 permanent virtual connections (PVCs) on the same PIC at the same time. See “Configuring ATM Interface Encapsulation” on page 330.
Cell-relay accumulation mode: The incoming cells (1 to 8) are packaged into a single packet and forwarded to the label-switched path (LSP).	Supported	Not supported	Cell-relay accumulation mode is per PIC, not per port. If you configure accumulation mode, the entire ATM1 PIC uses the configured mode. See “Configuring ATM Interface Encapsulation” on page 330.
Cell-relay promiscuous port mode: All cells from 0 through 65,535 of all VPIs (0 through 255) are sent to or received from an LSP.	Supported	Supported	For promiscuous mode, you must configure the port with atm-ccc-cell-relay encapsulation. For ATM2 IQ multiport PICs, you can configure one or more ports in port promiscuous mode, and the other ports with any ATM encapsulation.
Cell-relay promiscuous VPI mode: All cells in the VCI range 0 through 65,535 of a single VPI are sent to or received from an LSP.	Supported	Supported	For ATM2 IQ PICs, you can configure one or more logical interfaces in VPI promiscuous mode, and the other logical interfaces with any ATM encapsulation. For ATM1 PICs, if you configure one port in port mode, all ports on the PIC operate in port mode. Likewise if you configure one logical interface in VPI mode, all logical interfaces on the PIC operate in VPI mode. See “Configuring ATM Cell-Relay Promiscuous Mode” on page 296.

Table 25: ATM1 and ATM2 IQ Supported Features (*continued*)

Item	ATM1	ATM2 IQ	Comments
Cell-relay VP shaping	Supported	Supported	For ATM2 PICs, you can configure ATM CC cell relay promiscuous mode. VP promiscuous mode allows incoming traffic on all VCIs under the VPI to be bundled and directed to an LSP. Port promiscuous mode allows all traffic coming in on the entire VPI/VCI range to be forwarded to an LSP. In both modes, traffic shaping is not permitted. The ATM2 PIC supports traffic shaping in VP promiscuous mode and cell relay VC mode.
Cell-relay VCI mode: All cells in a VCI are sent to or received from an LSP.	Supported	Supported	For ATM1 PICs, nonpromiscuous cell-relay VCI, VPI, and port modes are supported on the same PIC with ATM AAL5 PVCs or ATM AAL5 CCC.
Cell-relay VPI mode: All cells in the VCI range (0 through <i>maximum-vc</i> s) of a single VPI are sent to or received from an LSP.	Supported	Not supported	For ATM2 IQ PICs, nonpromiscuous cell-relay VCI mode is supported on the same PIC with ATM AAL5 PVCs or ATM AAL5 CCC. See “Configuring ATM Interface Encapsulation” on page 330.
Cell-relay port mode: All cells in the VCI range (0 through <i>maximum-vc</i> s) of all VPIs (0 through 255) are sent to or received from an LSP.	Supported	Not supported	@@@amp@@@mdash;
Ethernet over ATM encapsulation: Allows ATM interfaces to connect to devices that support only bridged-mode protocol data units (PDUs).	Supported	Supported	See “Configuring ATM Interface Encapsulation” on page 330.
Layer 2 circuit cell-relay, Layer 2 circuit AAL5, and Layer 2 circuit trunk transport modes: Allow you to send ATM cells or AAL5 PDUs between ATM2 IQ interfaces across a Layer 2 circuit-enabled network. Layer 2 circuits are designed to transport Layer 2 frames between provider edge (PE) routers across a Label Distribution Protocol (LDP)-signaled Multiprotocol Label Switching (MPLS) backbone.	Not supported	Supported	Transport mode is per PIC, not per port. If you configure Layer 2 circuit cell-relay, Layer 2 circuit AAL5, or Layer 2 circuit trunk transport mode, the entire ATM2 IQ PIC uses the configured transport mode. Layer 2 circuit cell-relay mode supports both VP- and port-promiscuous modes. See “Configuring Layer 2 Circuit Transport Mode” on page 300.
Layer 2 VPN cell relay and Layer 2 VPN AAL5: Allow you to carry ATM cells or AAL5 PDUs over an MPLS backbone.	Supported	Supported	See the <i>JUNOS VPNs Configuration Guide</i> .

Table 25: ATM1 and ATM2 IQ Supported Features (continued)

Item	ATM1	ATM2 IQ	Comments
Point-to-Point Protocol (PPP) over ATM encapsulation: Associates a PPP link with an ATM AAL5 PVC.	Not supported	Supported	<p>For ATM2 IQ interfaces, the JUNOS Software supports three PPP over ATM encapsulation types:</p> <ul style="list-style-type: none"> ■ <code>atm-ppp-llc</code>—PPP over AAL5 logical link control (LLC). ■ <code>atm-ppp-vc-mux</code>—PPP over AAL5 multiplex. ■ <code>atm-mlppp-llc</code>—Multilink PPP over AAL5 LLC. Requires a Link Services or Voice Services PIC. <p>See “Configuring PPP over ATM2 Encapsulation” on page 334.</p>
Other ATM Attributes			
EPD (early packet discard) threshold: Limits the queue size in ATM cells of a particular VC or forwarding class configured over a VC when using VC tunnel class of service (CoS). When the first ATM cell of a new packet is received, the VC’s queue depth is checked against the EPD threshold. If the VC’s queue depth exceeds the EPD threshold, the first and all subsequent ATM cells in the packet are discarded.	Not supported	Supported	<p>If you are using VC tunnel CoS, the EPD threshold configured at the logical unit level has no effect. You should configure each forwarding class for congestion management using either an individual EPD threshold (in other words, tail drop) or weighted random early detection (WRED) profile.</p> <p>See “Configuring the ATM2 IQ EPD Threshold” on page 326 and “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.</p>
OAM F4 cell flows: Identify and report virtual path connection (VPC) defects and failures.	Not supported	Supported	See “Configuring the OAM F4 Cell Flows” on page 315.
OAM F5 loopback cell responses	Supported	Supported	<p>For ATM1 interfaces, when an OAM F5 loopback request is received, the response cell is sent by the PIC. The request and response cells are not counted in the VC, logical interface, or physical interface statistics.</p> <p>For ATM2 IQ interfaces, when an OAM F5 loopback request is received, the response is sent by the routing engine. The OAM, VC, logical interface, and physical interface statistics are incremented.</p> <p>See “Defining the ATM OAM F5 Loopback Cell Period” on page 329 and “Configuring the ATM OAM F5 Loopback Cell Threshold” on page 329.</p>
Passive monitoring mode	Supported	Supported	See “Enabling Passive Monitoring on ATM Interfaces” on page 293.
PIC type	Supported	Supported	<p>For ATM1 interfaces, you can include the <code>pic-type atm1</code> statement.</p> <p>For ATM2 IQ interfaces, you can include the <code>pic-type atm2</code> statement.</p> <p>See “Configuring the ATM PIC Type” on page 295.</p>

Table 25: ATM1 and ATM2 IQ Supported Features (continued)

Item	ATM1	ATM2 IQ	Comments
Ping	Supported	Supported	<p>For ATM1 and ATM2 IQ interfaces, when you issue the ATM ping command, you must include a logical unit number in the interface name, as shown in the following example:</p> <pre>ping atm interface at-1/0/0.5 vci 0.123 count 3</pre> <p>The logical unit number is 5 on physical interface at-1/0/0.</p> <p>See the <i>JUNOS Interfaces Command Reference</i>.</p>
Queue length: Limits the queue size in packets of a particular VC.	Supported	Not supported	See “Configuring the ATM1 Queue Length” on page 325.
Real-time variable bit rate (VBR): Supports VBR data traffic with average and peak traffic parameters.	Not supported	Supported	<p>Compared to non-real-time VBR, real-time VBR data is serviced at a higher priority. Real-time VBR is suitable for carrying packetized video and audio.</p> <p>See “Configuring ATM2 IQ Real-Time VBR” on page 321.</p>
Shaping rates: Peak and sustained rates of traffic.	Supported	Supported	<p>For ATM1 OC3 interfaces, the rate can be from 33 kilobits per second (Kbps) through 135.6 megabits per second (Mbps); for ATM1 OC12 interfaces, the rate can be from 33 Kbps through 276 Mbps.</p> <p>For ATM2 IQ OC3 interfaces, the rate can be from 33 Kbps through 135,600,000 bits per second (bps). For ATM2 IQ OC12 interfaces, the rate can be from 33 Kbps through 271,273,396 bps (up to 50 percent of the line rate).</p> <p>For ATM2 IQ OC48 interfaces, the rate can be from 33 Kbps through 2,170,107,168 bits per second (bps).</p> <p>For ATM2 IQ DS3 and E3 interfaces, the rate can be from 33 Kbps to the maximum rate. The maximum rate varies depending on the ATM encapsulation and framing you configure:</p> <ul style="list-style-type: none"> ■ For DS3 interfaces with direct ATM encapsulation, the maximum rate is 40,038,968 bps. ■ For DS3 interfaces with Physical Layer Convergence Protocol (PLCP) ATM encapsulation, the maximum rate is 36,864,000 bps. ■ For E3 interfaces with g.751 framing and direct ATM encapsulation, the maximum rate is 30,801,509 bps. ■ For E3 interfaces with g.751 framing PLCP ATM encapsulation, the maximum rate is 27,648,000 bps. ■ For E3 interfaces with g.832 framing, the maximum rate is 30,720,000 bps. <p>See “Defining the ATM Traffic-Shaping Profile” on page 319.</p>

Table 25: ATM1 and ATM2 IQ Supported Features (continued)

Item	ATM1	ATM2 IQ	Comments
VC tunnel CoS: Allows VCs to be opened as VC tunnels.	Not supported	Supported	<p>On M Series routers (except the M320 and M120 routers), a VC tunnel can support four CoS queues. On the M320, M120, and T Series routers, a VC tunnel can support eight CoS queues. Within the VC tunnel, the class-based weighted fair queuing algorithm is used to schedule packet transmission from each queue. You can configure the queue admission policies, such as EPD or WRED, to control the queue size during congestion.</p> <p>See “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.</p>
VCI management	Supported	Supported	<p>For ATM1 interfaces, you must specify the maximum number of VCIs by including the <code>maximum-vc</code> statement in the configuration. This restricts VCIs to the range 0 through <code>maximum-vc</code>. See “Configuring the Maximum Number of ATM1 VCs on a VP” on page 300.</p> <p>For ATM2 interfaces, you must not include the <code>maximum-vc</code> statement in the configuration. All ATM2 IQ interfaces support VCI numbers from 0 through 65,535. The total number of VCIs that you can open on an ATM2 IQ port depends on two factors:</p> <ul style="list-style-type: none"> ■ Number of tunnels ■ Sparseness of VCI numbers (the more sparse, the fewer VCIs supported) <p>For ATM1 and ATM2 IQ interfaces with promiscuous mode, the allowable maximum number of VCIs is 65,535.</p>
VCI statistics	Supported	Supported	<p>For ATM1 interfaces, multipoint VCI statistics are collected from indirect sources.</p> <p>For ATM2 IQ interfaces, multipoint VCI statistics are collected directly from the PIC.</p> <p>For ATM1 and ATM2 IQ interfaces, point-to-point VCI statistics are the same as logical interface statistics.</p>

Configuring Communication with Directly Attached ATM Switches and Routers

For ATM1 and ATM2 IQ interfaces, you can configure communication with directly attached ATM switches and routers to enable querying of the IP addresses and switch port numbers. You query the switch or router by entering the following `show` command:

```
user@host> show ilmi interface interface-name
```

The router uses VC 0.16 to communicate with the ATM switch or router.

To configure communication between the router and its directly attached ATM switches and routers, include the `ilmi` statement at the [edit interfaces *interface-name* atm-options] hierarchy level:

```
[edit interfaces interface-name atm-options]
ilmi;
```

Example: Configuring Communication with Directly Attached ATM Switches and Routers

Enable an interface to communicate directly with an ATM switch or router:

```
[edit interfaces]
at-0/1/0 {
  atm-options {
    vpi 0;
    ilmi;
  }
  unit 0 {
    vci 0.120;
    family inet {
      address 10.33.33.1/30;
    }
  }
}
```

Enabling ILMI for Cell Relay

The JUNOS Software supports standard AAL5 and three Layer 2 circuit transport modes: Layer 2 circuit AAL5, Layer 2 circuit cell-relay, and Layer 2 circuit trunk transport mode.

Integrated local management interface (ILMI) is supported on standard AAL5 interfaces, regardless of encapsulation. To enable ILMI on interfaces with cell-relay encapsulation, you must configure an ATM2 IQ PIC to use Layer 2 circuit trunk transport mode. ILMI is not supported with cell-relay encapsulation when the ATM2 IQ PIC is configured with Layer 2 AAL5 or Layer 2 circuit cell-relay transport mode, as shown in as shown in Table 26 on page 292.

Layer 2 circuit cell-relay trunk mode is not supported on ATM OC48 PICs.

Table 26: ILMI Support by Encapsulation Type

Encapsulation Type	ILMI Support
Standard AAL5, with any encapsulation type	Yes
Layer 2 circuit AAL5 mode	No
Layer 2 circuit cell-relay mode	No
Layer 2 circuit trunk mode	Yes

For more information about Layer 2 circuit transport modes, see “Configuring Layer 2 Circuit Transport Mode” on page 300.

To configure ILMI on an interface with cell-relay encapsulation, include the following statements:

```
[edit chassis fpc slot-number pic pic-number]
atm-l2circuit-mode trunk trunk;
[edit interfaces at-fpc/pic/port]
encapsulation atm-ccc-cell-relay;
atm-options {
    ilmi;
    pic-type atm2;
}
unit logical-unit-number {
    trunk-id number;
}
```

For more information about ILMI, see “Configuring Communication with Directly Attached ATM Switches and Routers” on page 291.

Example: Enabling ILMI for Cell Relay

On an ATM2 IQ PIC with Layer 2 circuit trunk transport mode, enable ILMI on an interface with cell-relay encapsulation:

```
[edit chassis]
fpc 0 {
    pic 1 {
        atm-l2circuit-mode trunk uni;
    }
}
[edit interfaces]
at-0/0/0 {
    encapsulation atm-ccc-cell-relay;
    atm-options {
        pic-type atm2;
        ilmi;
    }
}
```

Enabling Passive Monitoring on ATM Interfaces

The Monitoring Services I and Monitoring Services II PICs are designed to enable IP services. If you have a Monitoring Services PIC and an ATM PIC installed in an M160, M40e, or T Series router, you can monitor IP version 4 (IPv4) traffic from another router.

On ATM interfaces, you enable packet flow monitoring by including the `passive-monitor-mode` statement at the `[edit interfaces at-fpc/pic/port]` hierarchy level:

```
[edit interfaces at-fpc/pic/port]
passive-monitor-mode;
```

If you include the **passive-monitor-mode** statement in the configuration, the ATM interface is always up, and the interface does not receive or transmit incoming control packets, such as OAM cell and ILMI.

On monitoring services interfaces, you enable packet flow monitoring by including the **family** statement at the `[edit interfaces mo-fpc/pic/port unit logical-unit-number]` hierarchy level, specifying the **inet** option:

```
[edit interfaces mo-fpc/pic/port unit logical-unit-number]
family inet;
```

For conformity with cflowd record structure, you must include the **receive-options-packets** and **receive-ttl-exceeded** statements at the `[edit interfaces mo-fpc/pic/port unit logical-unit-number family inet]` hierarchy level:

```
[edit interfaces mo-fpc/pic/port unit logical-unit-number family inet]
receive-options-packets;
receive-ttl-exceeded;
```

For the monitoring services interface, you can configure multiservice physical interface properties. For more information, see “Configuring Multiservice Physical Interface Properties” on page 138 and the *JUNOS Services Interfaces Configuration Guide*.

Removing MPLS Labels from Incoming Packets

The JUNOS Software can forward only IPv4 packets to a Monitoring Services PIC. IPv4 packets with MPLS labels cannot be forwarded to a Monitoring Services PIC. By default, if packets with MPLS labels are forwarded to the Monitoring Services PIC, they are discarded. To monitor packets with MPLS labels, you must remove the MPLS labels as the packets arrive on the interface.

You can remove up to two MPLS labels from an incoming packet by including the **pop-all-labels** statement at the `[edit interfaces interface-name atm-options mpls]` hierarchy level:

```
[edit interfaces interface-name atm-options mpls]
pop-all-labels {
    required-depth number;
}
```

By default, the **pop-all-labels** statement takes effect for incoming packets with one or two labels. You can specify the number of MPLS labels an incoming packet must have for the **pop-all-labels** statement to take effect by including the **required-depth** statement at the `[edit interfaces interface-name atm-options mpls pop-all-labels]` hierarchy level:

```
[edit interfaces interface-name atm-options mpls pop-all-labels]
required-depth number;
```

The required depth can be 1, 2, or [1 2]. If you include the **required-depth 1** statement, the **pop-all-labels** statement takes effect for incoming packets with one label only. If you include the **required-depth 2** statement, the **pop-all-labels** statement takes effect for incoming packets with two labels only. If you include the **required-depth**

[1 2] statement, the **pop-all-labels** statement takes effect for incoming packets with one or two labels. A required depth of [1 2] is equivalent to the default behavior of the **pop-all-labels** statement.

When you remove MPLS labels from incoming packets, note the following:

- The **pop-all-labels** statement has no effect on IP packets with three or more MPLS labels.
- When you enable MPLS label removal, you must configure all ports on a PIC with the same label popping mode and required depth.
- You use the **pop-all-labels** statement to enable passive monitoring applications, not active monitoring.
- You cannot apply MPLS filters or accounting to the MPLS labels because the labels are removed as soon as the packet arrives on the interface.
- The following ATM encapsulation types are not supported on interfaces with MPLS label removal:
 - atm-ccc-cell-relay
 - atm-ccc-vc-mux
 - atm-mlppp-llc
 - atm-tcc-snap
 - atm-tcc-vc-mux
 - ether-over-atm-llc
 - ether-vpls-over-atm-llc

Configuring the ATM PIC Type

For ATM1 and ATM2 IQ interfaces, the JUNOS Software does not determine from the interface name *at- fpc/pic/port* whether your router has an ATM1 or ATM2 IQ PIC installed. You can configure the PIC type as ATM1 or ATM2 IQ by including the **pic-type** statement at the [edit interfaces *interface-name* atm-options] hierarchy level:

```
[edit interfaces interface-name atm-options]
pic-type (atm1 | atm2);
```

The following guidelines apply to configuring the ATM PIC type:

- If you include the **pic-type** statement in the configuration, and you include other statements at the [edit interfaces *interface-name* atm-options] hierarchy level that do not match the configured PIC type, the configuration does not commit. For example, you cannot commit a configuration that includes the **pic-type atm2** statement and the **maximum-vcs** statement.
- If you do not include the **pic-type** statement and you do include the **maximum-vcs** statement in the configuration, the JUNOS Software assumes you are configuring an ATM1 interface, and sets the PIC type option accordingly. If you do not include the **maximum-vcs** statement in the configuration, the JUNOS Software assumes

you are configuring an ATM2 IQ interface, and sets the PIC type option accordingly.

- If you include the **promiscuous-mode** statement in the configuration of an ATM2 interface, you must also include the **pic-type atm2** statement.

Example: Configuring the ATM PIC Type

Configure the PIC type on an ATM1 and an ATM2 interface.

On an ATM1 Interface

```
[edit interfaces]
at-1/0/0 {
  atm-options {
    pic-type atm1;
    vpi 0 maximum-vcs 256;
    vpi 1 maximum-vcs 512;
  }
  ...
}
```

On an ATM2 IQ Interface

```
[edit interfaces]
at-1/1/0 {
  atm-options {
    pic-type atm2;
    vpi 0;
    vpi 2 {
      oam-period 6;
    }
  }
  ...
}
```

Configuring ATM Cell-Relay Promiscuous Mode

For ATM1 and ATM2 IQ interfaces with **atm-ccc-cell-relay** encapsulation, you can map all incoming cells from either an interface port or a virtual path (VP) to a single LSP without restricting the VCI number. Promiscuous mode allows you to map traffic from all 65,535 VCIs to a single LSP, or from all 256 VPIs to a single LSP.

To map incoming traffic from a port or VC to an LSP, include the **promiscuous-mode** statement at the `[edit interfaces interface-name atm-options]` hierarchy level:

```
[edit interfaces interface-name]
atm-options {
  promiscuous-mode {
    vpi vpi-identifier;
  }
}
```

You can include multiple **vpi** statements in the configuration.

To enable all VCIs in a VPI to open in ATM CCC cell-relay mode, you must also map the logical interface to a VPI by including the `vpi` statement in the logical interface configuration:

```
vpi vpi-identifier;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Also, note the following:

- For promiscuous mode, you must configure the port with `atm-ccc-cell-relay encapsulation`.
- For ATM1 and ATM2 IQ PICs, changing modes between promiscuous and nonpromiscuous causes all physical interfaces to be deleted and re-added.
- For ATM1 and ATM2 IQ PICs, when you configure promiscuous mode, you cannot configure VCIs.
- For ATM1 PICs, if you configure one port in port mode, all ports on the PIC operate in port mode. Likewise if you configure one logical interface in VPI mode, all logical interfaces on the PIC must operate in VPI mode.
- For ATM2 IQ PICs, you can configure one or more logical interfaces in VPI promiscuous mode, and the other logical interfaces with any ATM encapsulation.
- For ATM2 IQ PICs, when you configure promiscuous mode, you must also include the `pic-type atm2` statement. For more information, see “Configuring the ATM PIC Type” on page 295.
- For ATM2 IQ multiport PICs, you can configure one or more ports in port promiscuous mode, and the other ports with any ATM encapsulation.
- For interfaces that are configured for cell-relay promiscuous virtual path identifier (VPI) mode, the `show interfaces` command output does not show OAM F4 cell statistics.

Examples: Configuring ATM Cell-Relay Promiscuous Mode

This section includes the following examples:

Configuring Port-Promiscuous Mode

```
[edit interfaces]
at-0/2/1 {
  encapsulation atm-ccc-cell-relay; # at the physical interface level only
  atm-options {
    pic-type atm2;
    promiscuous-mode;
  }
  unit 0 {
    allow-any-vci;
  }
}
```

**Configuring
VP-Promiscuous Mode**

```
[edit interfaces]
at-0/2/0 {
  atm-options {
    pic-type atm2;
    promiscuous-mode {
      vpi 0;
      vpi 1;
    }
    vpi 2;
    vpi 3;
  }
  unit 0 {
    encapsulation atm-ccc-cell-relay; # at the logical interface level only
    vpi 0;
  }
  unit 1 {
    encapsulation atm-ccc-cell-relay;
    vpi 1;
  }
  unit 2 {
    encapsulation atm-snap;
    vci 2.100;
  }
  unit 3 {
    encapsulation atm-vc-mux;
    vci 3.100;
  }
}
```

To map incoming traffic from a port to an LSP, include the **allow-any-vci** statement at the [edit interfaces *interface-name* unit 0] hierarchy level. When you include the **allow-any-vci** statement, you cannot configure other logical interfaces in the same physical interface. Next, you must map unit 0 to an LSP using the CCC connection.

**Mapping Incoming
Traffic from a Port to an
LSP**

```
[edit interfaces at-1/2/0]
encapsulation atm-ccc-cell-relay;
atm-options {
  promiscuous-mode;
}
unit 0 {
  allow-any-vci;
}
```

**Mapping Unit 0 to an
LSP**

```
protocols {
  connections {
    remote-interface-switch router-a-router-c {
      interface at-1/2/0.0;
    }
    lsp-switch router-a-router-c {
      transmit-lsp lsp1
      receive-lsp lsp2;
    }
  }
}
```

To map a VPI to an LSP, you must define the allowed VPIs. You can configure one or more logical interfaces, each mapped to a different VPI. You can then route traffic from each of these interfaces to different LSPs.

Mapping a VPI to an LSP

```
[edit interfaces at-1/1/0]
encapsulation atm-ccc-cell-relay;
atm-options {
  pic-type atm1;
  promiscuous-mode {
    vpi 10;
    vpi 20;
  }
}
unit 0 {
  encapsulation atm-ccc-cell-relay;
  vpi 10;
}
unit 1 {
  encapsulation atm-ccc-cell-relay;
  vpi 20;
}
[edit interfaces at-3/1/0]
encapsulation atm-ccc-cell-relay;
atm-options {
  pic-type atm2;
  promiscuous-mode {
    vpi 10;
    vpi 20;
  }
}
unit 0 {
  encapsulation atm-ccc-cell-relay;
  vpi 10;
}
unit 1 {
  encapsulation atm-ccc-cell-relay;
  vpi 20;
}
[edit protocols]
mpls {
  connections {
    interface-switch router-a-router-c {
      interface at-1/1/0.0;
      interface at-3/1/0.0;
    }
    interface-switch router-a-router-d {
      interface at-1/1/0.1;
      interface at-3/1/0.1;
    }
  }
}
```

Configuring the Maximum Number of ATM1 VCs on a VP

For ATM1 interfaces, you must configure the maximum number of virtual circuits (VCs) allowed on a virtual path (VP) so that sufficient memory on the ATM1 PIC can be allocated for each VC.

To configure the highest-numbered VCs on a VP, include the `maximum-vcs` and `vpi` statements at the `[edit interfaces interface-name atm-options]` hierarchy level:

```
[edit interfaces interface-name atm-options]
vpi vpi-identifier {
    maximum-vcs maximum-vcs;
}
```

The VP identifier can be a value from 0 through 255. For most interfaces, you can define a maximum of 4090 VCs per interface, and some interfaces have higher limits. Promiscuous mode removes these limits. For more information, see “Configuring ATM Cell-Relay Promiscuous Mode” on page 296.

All VPIs that you configure in the `atm-options` statement are stored in a single table. If you modify the VPIs—for example, by editing them in configuration mode or by issuing a `load override` command—all VCs on the interface are closed and then reopened, resulting in a temporary loss of connectivity for all the VCs on the interface.

You can also include some of the statements in the `sonet-options` statement to set SONET/SDH parameters on ATM interfaces, as described in “Configuring SONET/SDH Parameters on ATM Interfaces” on page 338.

Configuring Layer 2 Circuit Transport Mode

On ATM2 IQ interfaces only, you can configure Layer 2 circuit cell-relay, Layer 2 circuit AAL5, or Layer 2 circuit trunk transport mode.

Layer 2 circuit cell-relay and Layer 2 circuit AAL5 are defined in Internet draft `draft-martini-l2circuit-encap-mpls-07.txt`, *Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks* (expires December 2004).

Layer 2 circuit cell-relay and Layer 2 circuit AAL5 transport modes allow you to send ATM cells between ATM2 IQ interfaces across a Layer 2 circuit-enabled network. Layer 2 circuits are designed to transport Layer 2 frames between PE routers across an LDP-signaled MPLS backbone. You use Layer 2 circuit AAL5 transport mode to send AAL5 segmentation and reassembly protocol data units (SAR-PDUs) over the Layer 2 circuit.

A trunk is a collection of ATM VPs. Layer 2 circuit trunk transport mode allows you to send ATM cells over MPLS trunking.

By default, ATM2 IQ PICs are in standard AAL5 transport mode. Standard AAL5 allows multiple applications to tunnel the protocol data units of their Layer 2 protocols over an ATM virtual circuit. Encapsulation of these Layer 2 protocol data units allows a number of these emulated virtual circuits to be carried in a single tunnel. Protocol

data units are segmented at one end of the tunnel and reassembled at the other end. The ingress router reassembles the protocol data units received from the incoming VC and transports each PDU as a single packet.

In contrast, Layer 2 circuit cell-relay and Layer 2 circuit AAL5 transport modes accept a stream of ATM cells, convert these to an encapsulated Layer 2 format, then tunnel them over an MPLS or IP backbone, where a similarly configured router segments these packets back into a stream of ATM cells, to be forwarded to the virtual circuit configured for the far-end router.

In Layer 2 circuit cell-relay transport mode, ATM cells are bundled together and transported in packet form to the far-end router, where they are segmented back into individual ATM cells and forwarded to the ATM virtual circuit configured for the far-end router.

The uses for the four transport modes are defined as follows:

- To tunnel IP packets over an ATM backbone, use the default standard AAL5 transport mode.
- To tunnel a stream of AAL5-encoded ATM SAR-PDUs over an MPLS or IP backbone, use Layer 2 circuit AAL5 transport mode.
- To tunnel a stream of ATM cells over an MPLS or IP backbone, use Layer 2 circuit cell-relay transport mode.
- To transport ATM cells over an MPLS core network that is implemented between other vendors' switches or routers, use Layer 2 circuit trunk transport mode.



NOTE: You can transport AAL5-encoded traffic with Layer 2 circuit cell-relay transport mode, because Layer 2 circuit cell-relay transport mode ignores the encoding of the cell data presented to the ingress interface.

When you configure AAL5 mode Layer 2 circuits, the control word carries cell loss priority (CLP) information by default.

The Layer 2 circuit trunk transport mode is not supported on the ATM2 IQ OC48c/STM16 PIC.

To configure Layer 2 circuit AAL5, Layer 2 circuit cell-relay, or Layer 2 circuit trunk mode, you must perform the following tasks:

1. Identify the interface as an ATM2 IQ interface by including the **pic-type atm2** statement at the **[edit interfaces at-fpc/pic/port atm-options]** hierarchy level:

```
[edit interfaces at-fpc/pic/port atm-options]
pic-type atm2;
```

2. Include the **atm-l2circuit-mode** statement at the **[edit chassis fpc slot-number pic pic-number]** hierarchy level, specifying **aal5**, **cell**, or **trunk**:

```
[edit chassis fpc slot-number pic pic-number]
atm-l2circuit-mode (aal5 | cell | trunk trunk );
```

By default, the trunk mode uses user-to-network interface (UNI) mode. The trunk option can be UNI or network-to-network interface (NNI). For more information about UNI and NNI, see the *JUNOS VPNs Configuration Guide* and the *JUNOS Feature Guide*.

Transport mode is per PIC, not per port. If you do not include the **atm-l2circuit-mode** statement in the configuration, the ATM2 IQ PIC uses standard AAL5 transport mode. If you configure Layer 2 circuit cell-relay, Layer 2 circuit AAL5 transport mode, or Layer 2 circuit trunk mode, the entire ATM2 PIC uses the configured transport mode.

3. For Layer 2 circuit trunk mode only, you must also configure a trunk identification number by including the **trunk-id** statement:

```
trunk-id number;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The trunk identification number can be from 0 through 31; each trunk on an interface must have a unique trunk ID. When you associate a trunk ID number with a logical interface, you are in effect specifying the interfaces that are allowed to send ATM traffic over an LSP. For UNI mode, the trunk ID range is from 0 through 7. For NNI mode, the trunk ID range is from 0 through 31. Trunk IDs on connecting trunks do not need to be the same.

For information about proportional bandwidth sharing in trunk mode, see “Configuring Layer 2 Circuit Trunk Mode Scheduling” on page 309.

4. For Layer 2 circuit AAL5 mode, configure logical interface encapsulation by including the **encapsulation** statement, specifying the **atm-ccc-vc-mux** encapsulation type:

```
encapsulation atm-ccc-vc-mux;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

5. For Layer 2 circuit cell-relay and Layer 2 circuit trunk modes, configure physical interface encapsulation by including the **encapsulation** statement at the [edit interfaces *interface-name*] hierarchy level, specifying the **atm-ccc-cell-relay** encapsulation type:

```
[edit interfaces interface-name]  
encapsulation atm-ccc-cell-relay;
```

For more information about Layer 2 circuits, see the *JUNOS VPNs Configuration Guide* and the *JUNOS Routing Protocols Configuration Guide*. For a comprehensive example, see the *JUNOS Feature Guide*.

Examples: Configuring IQ Layer 2 Circuit Transport Mode

Configure Layer 2 circuit AAL5 transport mode and cell-relay transport mode.


Configuring Layer 2 Circuit AAL5 Transport Mode	<pre>[edit chassis] fpc 0 { pic 1 { atm-l2circuit-mode aal5; } } [edit interfaces] at-0/1/0 { atm-options { pic-type atm2; vpi 0; } unit 0 { encapsulation atm-ccc-vc-mux; point-to-point; vci 0.32; } }</pre>
Configuring Layer 2 Circuit Cell-Relay Transport Mode	<pre>[edit chassis] fpc 0 { pic 1 { atm-l2circuit-mode cell; } } [edit interfaces] at-0/1/0 { encapsulation atm-ccc-cell-relay; atm-options { pic-type atm2; vpi 0; } unit 0 { encapsulation atm-ccc-cell-relay; point-to-point; vci 0.32; } }</pre>

Configuring Layer 2 Circuit Trunk Transport Mode

In Figure 20 on page 304, Router A is a local PE router. Router B is a remote PE router. Both Juniper Networks routers have Layer 2 circuit cell-relay capability. You configure an ATM physical interface on Router A in Layer 2 circuit trunk mode and specify trunks that are allowed to send traffic over the LSP. As a cell is received on this interface, it is classified using the CoS bits in the cell header, and encapsulated as a labeled packet. It is then queued on one of the outgoing queues according to its classification and sent over the LSP to Router B. At Router B, the packet label is removed and the raw cell is put on one of the queues of the ATM interface and forwarded to the second ATM switch. To carry the CoS information and CLP of the cell over the network, the CoS and CLP bits are copied into the EXP bits of the MPLS label. This CoS information is used to select the output queues. Using EPD profiles, the CLP is used to determine whether the cell should be dropped.

For more information about ATM CoS capability, see “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.

Figure 20: Layer 2 Circuit Trunk Topology



```
On Router A    [edit chassis]
                fpc 0 {
                  pic 1 {
                    atm-l2circuit-mode trunk uni;
                  }
                }
                [edit interfaces]
                at-0/0/0 {
                  encapsulation atm-ccc-cell-relay;
                  atm-options {
                    pic-type atm2;
                    ilmi;
                  }
                  unit 0 {
                    trunk-id 0;
                    epd-threshold 10240;
                  }
                  unit 1 {
                    trunk-id 1;
                    epd-threshold 10240;
                  }
                  unit 2 {
                    trunk-id 2;
                    epd-threshold 10240;
                  }
                  unit 3 {
                    trunk-id 3;
                    epd-threshold 10240;
                  }
                  unit 4 {
```



```

        trunk-id 4;
        epd-threshold 10240;
    }
    unit 5 {
        trunk-id 5;
        epd-threshold 10240;
    }
    unit 6 {
        trunk-id 6;
        epd-threshold 10240;
    }
    unit 7 {
        trunk-id 7;
        epd-threshold 10240;
    }
}
so-3/0/0 {
    mtu 9192;
    unit 0 {
        family inet {
            address 10.0.1.1/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 172.16.0.1/32;
            address 10.255.245.1/32;
        }
    }
}
[edit protocols]
rsvp {
    interface all;
}
mpls {
    interface all;
}
ldp {
    interface all;
}
ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
l2circuit {
    neighbor 10.255.245.2 {
        interface at-0/1/0.0 {

```

```

        virtual-circuit-id 100;
    }
    interface at-0/1/0.1 {
        virtual-circuit-id 101;
    }
    interface at-0/1/0.2 {
        virtual-circuit-id 102;
    }
    interface at-0/1/0.3 {
        virtual-circuit-id 103;
    }
    interface at-0/1/0.4 {
        virtual-circuit-id 104;
    }
    interface at-0/1/0.5 {
        virtual-circuit-id 105;
    }
    interface at-0/1/0.6 {
        virtual-circuit-id 106;
    }
    interface at-0/1/0.7 {
        virtual-circuit-id 107;
    }
    }
}

```

On Router B

```

[edit chassis]
fpc 0 {
    pic 1 {
        atm-l2circuit-mode trunk uni;
    }
}
[edit interfaces]
at-0/0/1 {
    encapsulation atm-ccc-cell-relay;
    atm-options {
        pic-type atm2;
    }
    unit 0 {
        trunk-id 0;
        epd-threshold 10240;
    }
    unit 1 {
        trunk-id 1;
        epd-threshold 10240;
    }
    unit 2 {
        trunk-id 2;
        epd-threshold 10240;
    }
    unit 3 {
        trunk-id 3;
        epd-threshold 10240;
    }
    unit 4 {

```

```

        trunk-id 4;
        epd-threshold 10240;
    }
    unit 5 {
        trunk-id 5;
        epd-threshold 10240;
    }
    unit 6 {
        trunk-id 6;
        epd-threshold 10240;
    }
    unit 7 {
        trunk-id 7;
        epd-threshold 10240;
    }
}
so-3/0/1 {
    mtu 9192;
    unit 0 {
        family inet {
            address 10.0.1.2/24;
        }
        family mpls;
    }
}
lo0 {
    unit 0 {
        family inet {
            address 172.16.0.1/32;
            address 10.255.245.2/32;
        }
    }
}
[edit protocols]
rsvp {
    interface all;
}
mpls {
    interface all;
}
ldp {
    interface all;
}
ospf {
    traffic-engineering;
    reference-bandwidth 4g;
    area 0.0.0.0 {
        interface all;
        interface fxp0.0 {
            disable;
        }
    }
}
l2circuit {
    neighbor 10.255.245.1 {
        interface at-0/1/0.0 {

```

```

        virtual-circuit-id 100;
    }
    interface at-0/1/0.1 {
        virtual-circuit-id 101;
    }
    interface at-0/1/0.2 {
        virtual-circuit-id 102;
    }
    interface at-0/1/0.3 {
        virtual-circuit-id 103;
    }
    interface at-0/1/0.4 {
        virtual-circuit-id 104;
    }
    interface at-0/1/0.5 {
        virtual-circuit-id 105;
    }
    interface at-0/1/0.6 {
        virtual-circuit-id 106;
    }
    interface at-0/1/0.7 {
        virtual-circuit-id 107;
    }
}
}

```

Configuring Layer 2 Circuit Cell-Relay Promiscuous Mode

By default, all incoming cells are mapped from a single VC to an external LSP. For ATM interfaces with Layer 2 circuit cell-relay transport mode and `atm-ccc-cell-relay` encapsulation, you can configure promiscuous mode. Promiscuous mode allows you to map all incoming cells from either an interface port or a VP to a single LSP without restricting the VCI number. You can map traffic from all 65,535 VCIs to a single LSP, or from all 256 VPIs to a single LSP. For promiscuous-mode configuration guidelines, see “Configuring ATM Cell-Relay Promiscuous Mode” on page 296.

Example: Configuring Layer 2 Circuit Cell-Relay Promiscuous Mode

Configure Layer 2 circuit cell-relay VP- and port-promiscuous mode:

```

VP-Promiscuous Mode  [edit interfaces]
                        at-0/1/0 {
                          encapsulation atm-ccc-cell-relay;
                          atm-options {
                            pic-type atm2;
                            cell-bundle-size 4;
                            promiscuous-mode {
                              vpi 0;
                            }
                          }
                        }
                        unit 0 {
                          encapsulation atm-ccc-cell-relay;
                          point-to-point;
                        }

```

```

        vci 0.32;
    }
}

```

Port-Promiscuous Mode

```

[edit interfaces]
at-0/1/0 {
    encapsulation atm-ccc-cell-relay;
    atm-options {
        pic-type atm2;
        promiscuous-mode;
    }
    unit 0 {
        allow-any-vci;
    }
}

```

Configuring Layer 2 Circuit Trunk Mode Scheduling

For ATM2 IQ interfaces configured to use Layer 2 circuit trunk mode, you can share a scheduler among 32 trunks on an ATM port. A weighted round robin scheduling algorithm ensures each trunk receives a proportional share of the bandwidth when all trunks are active, and redistributes bandwidth that would have otherwise been reserved by an inactive trunk, thus minimizing the latency on each trunk. For general information about Layer 2 circuit trunk mode, see “Configuring Layer 2 Circuit Transport Mode” on page 300. For general information about ATM CoS scheduling, see “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.

Each trunk is associated with a trunk bandwidth. The trunk bandwidth is the maximum bandwidth used each time a trunk is serviced. We recommend configuring trunk bandwidths so that the ratio between the minimum and maximum bandwidths does not exceed 1:500.

To minimize latency, the JUNOS Software does not shape the trunks. As cells are received, they are immediately transmitted.

To configure trunk bandwidth, include the **trunk-bandwidth** statement:

```
trunk-bandwidth rate;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The trunk bandwidth can be from 1,000,000 through 542,526,792 bps. You can specify the rate in bits per second or cells per second (cps). You can specify a bits-per-second value either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can specify a cells-per-second value by entering a decimal number followed by the abbreviation c; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps.

The JUNOS Software rounds off the configured value. Therefore, we recommend that you configure a minimum trunk bandwidth of 1m. From 1m, configure values in increments of 500k.

Example: Configuring Layer 2 Circuit Trunk Mode Scheduling

Configure two logical interfaces to use Layer 2 circuit trunk mode, ATM CoS scheduling, and proportional bandwidth sharing:

```
[edit interface]
at-1/1/0 {
  encapsulation atm-ccc-cell-relay;
  atm-options {
    pic-type atm2;
    ilmi;
    scheduler-maps {
      trunk-map {
        vc-cos-mode strict;
        forwarding-class cbr-class {
          priority high;
          transmit-weight percent 40;
          epd-threshold 100;
        }
        forwarding-class rtvbr-class {
          priority low;
          transmit-weight percent 30;
          epd-threshold 100;
        }
        forwarding-class nrtvbr-class {
          priority low;
          transmit-weight percent 20;
          epd-threshold 100;
        }
        forwarding-class ubr-class {
          priority low;
          transmit-weight percent 10;
          epd-threshold 100;
        }
      }
    }
  }
}
unit 0 {
  encapsulation atm-ccc-cell-relay;
  trunk-id 1;
  trunk-bandwidth 10m;
  atm-scheduler-map trunk-map;
  family ccc {
    filter {
      output atm-trunk-01;
    }
  }
}
unit 1 {
  encapsulation atm-ccc-cell-relay;
  trunk-id 3;
```

```

        trunk-bandwidth 30m;
        atm-scheduler-map trunk-map;
    }
}

```

Configuring CoS Queues in Layer 2 Circuit Trunk Mode

On ATM2 IQ interfaces, you can configure ATM CoS scheduling for AAL5 mode and Layer 2 circuit trunk mode. For general information about ATM CoS, see “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.

When you configure CoS scheduling in Layer 2 circuit trunk mode, the trunk is defined on the logical interface, and four CoS queues are opened in the trunk. For each CoS queue, you specify a priority and a transmit weight. CoS queues are serviced using a weighted round robin (WRR) algorithm. One queue is serviced with strictly high priority and the remaining queues are serviced with the WRR.

For Layer 2 circuit trunk mode, only strict mode is supported. Alternate mode is not supported.

To configure CoS queues in Layer 2 circuit trunk mode, perform the following tasks:

1. Include the `encapsulation atm-ccc-cell-relay` statement at the [edit interfaces *at-fpc/pic/port*] hierarchy level:

```

[edit interfaces at-fpc/pic/port]
encapsulation atm-ccc-cell-relay;

```

2. Include the `scheduler-maps` statement at the [edit interfaces *at-fpc/pic/port* *atm-options*] hierarchy level:

```

[edit interfaces at-fpc/pic/port atm-options]
scheduler-maps map-name {
    forwarding-class (class-name | assured-forwarding | best-effort |
        expedited-forwarding | network-control);
    vc-cos-mode strict;
}

```

3. Include the `atm-scheduler-map`, `trunk-bandwidth`, and `trunk-id` statements at the [edit interfaces *at-fpc/pic/port* unit *logical-unit-number*] hierarchy level:

```

[edit interfaces at-fpc/pic/port unit logical-unit-number]
atm-scheduler-map (map-name | default);
trunk-bandwidth rate;
trunk-id number;

```

For information about ATM scheduler maps, see “Configuring an ATM Scheduler Map” on page 341.

For information about trunk identification numbers, see “Configuring Layer 2 Circuit Transport Mode” on page 300. For information about trunk bandwidths, see “Configuring Layer 2 Circuit Trunk Mode Scheduling” on page 309.

Strict mode CoS queue priority works as follows:

- **Scheduling**—One queue has strictly high priority and is always serviced before the remaining queues are serviced by a weighted round robin. This means the packets in a **high** priority queue are sent first until the queue is empty. Then **low** priority queues send packets until their weight quota becomes zero or negative.
- **Latency**—Each trunk is associated with a trunk bandwidth. The trunk bandwidth is the maximum bandwidth used each time a trunk is serviced. In the scheduling process, each trunk is serviced in a WRR. The maximum latency for any trunk to begin transmitting is equal to the sum of the weights of all previously queued trunks. Trunks without data do not affect output scheduling. As long as all the trunks have data, the exact weight proportions are maintained. If a trunk runs out of data during its turn, it is no longer included in the WRR. When the trunk gets more data, the trunk is placed at the end of the queue. For more information, see “Configuring Layer 2 Circuit Trunk Mode Scheduling” on page 309.

Within a single trunk, the maximum latency of a **high** priority queue is the time it takes to transmit one ATM cell. The latency of a **low** priority queue is the sum of **high** priority queue burst time and the transmission time of the remaining **low** priority queues’ weight.

- **Bandwidth distribution**—Trunks are serviced in a WRR based on the trunk bandwidth.

Within a single trunk, the **high** priority queue consumes the bandwidth first regardless of its weight. The remaining bandwidth is distributed to the **low** priority queues in proportion to their weights.

Consider the following example:

- You configure a trunk with weights of 10 percent, 20 percent, 30 percent, and 40 percent for queues 0, 1, 2, and 3, respectively.
- You configure queue 0 to be a high priority queue.
- Queue 0 does not have cells to transmit.

In this scenario, queues 1, 2 and 3 receive 2/9, 3/9, and 4/9 of the bandwidth, respectively.



NOTE: Constant bit rate (CBR) traffic always enters the strictly **high** priority queue.

For more information about strict and alternate modes, see “Configuring VC CoS Mode” on page 348.

For general information about Layer 2 circuit trunk mode, see “Configuring Layer 2 Circuit Transport Mode” on page 300.

For interfaces configured in trunk mode, you can also configure dual EPD thresholds depending on packet loss priorities (PLPs). For more information, see “Configuring Two EPD Thresholds per Queue” on page 328.

Example: Configuring CoS Queues in Layer 2 Circuit Trunk Mode

Configure a scheduler map and trunk bandwidth:

```
[edit interfaces]
at-6/1/0 {
  encapsulation atm-ccc-cell-relay;
  atm-options {
    pic-type atm2;
    scheduler-maps {
      cos0 {
        vc-cos-mode strict;
        forwarding-class cbr-class {
          priority high;
          transmit-weight percent 10;
        }
        forwarding-class rtvbr-class {
          priority low;
          transmit-weight percent 20;
        }
        forwarding-class nrtvbr-class {
          priority low;
          transmit-weight percent 30;
        }
        forwarding-class ubr-class {
          priority low;
          transmit-weight percent 40;
        }
      }
    }
  }
  unit 0 {
    trunk-id 0;
    trunk-bandwidth 10m;
    atm-scheduler-map cos0;
  }
}
```

Configuring the Layer 2 Circuit Cell-Relay Cell Maximum

By default, each frame contains one cell. For ATM interfaces with Layer 2 circuit cell-relay transport mode configured, you can configure the maximum number of ATM cells per frame on the physical or logical interface. To set the maximum number of cells per frame, include the `cell-bundle-size` statement:

```
cell-bundle-size cells;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* atm-options]
- [edit interfaces *interface-name* unit *logical-unit-number*]

- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The cell bundle size can be from 1 through 176.

After 125 microseconds, cell bundling times out. This means that after 125 microseconds if the frame does not contain the configured value, the frame is transmitted anyway.

If you include the **cell-bundle-size** statement at the [edit interfaces *interface-name* atm-options] hierarchy level, then the configured value becomes the default for all the logical interface units configured for that physical interface. If you include the **cell-bundle-size** statement for a logical interface, the logical interface configuration overrides the value configured at the physical interface level.

The transmit rates you configure on the routers at each end of the connection must be the same value.

Class-Based Cell Bundling

For Layer 2 circuit trunk mode only, cell bundling is enhanced by a set of CoS and traffic shaping rules, as follows:

- CBR and real-time variable bit rate (RTVBR) cells are not bundled. They are always sent as single-cell packets.
- Cells with the same CLP bits are bundled together. This means all the cells in a bundle contain the same CLP value.
- Cells with the same CoS bits are bundled together. This means all the cells in a bundle belong to the same class of service.
- As alluded to in the previous rules, several triggers cause early packet transmission, meaning that the packet is transmitted before the number of cells received is equal to the value configured with the **cell-bundle-size** statement. These triggers are as follows:
 - The next cell is of type CBR or RTVBR.
 - The next cell has a different CLP bit.
 - The next cell has different CoS bits.
 - The 125-microsecond timer expires.

CoS-based cell bundling optimizes the release of a bundle by sending out the cell that triggers early packet transmission as a single-cell packet. This means that when a cell triggers early packet transmission, that cell is not bundled. Consequently, certain input data patterns might cause primarily single-cell packets to be transmitted. For example, say the output interface receives a steady pattern of two cells from a non-RTVBR queue, followed by two cells from a UBR queue. In this case, all transmitted packets contain a single cell because the first cell triggers a transition and is transmitted by itself. The second cell is also transmitted by itself because the third cell triggers another transition, and so on. This effect might not be dramatic

with a mix of traffic; it is most evident with steady traffic patterns, as generated by ATM test equipment programmed to emit regular sequences of CoS queue transitions.

Configuring the OAM F4 Cell Flows

For ATM2 IQ interfaces, the F4 flow cell is used for management of the VP level. If your router is equipped with an ATM2 IQ PIC, you can configure OAM F4 cell flows to identify and report VPC defects and failures. The JUNOS Software supports three types of OAM F4 cells in end-to-end F4 flows:

- Virtual Path Alarm Indication Signal (VP-AIS)
- Virtual Path Remote Defect Indication (VP-RDI)
- Virtual Path Loopback

The JUNOS Software does not support segment F4 flows, VPC continuity check, or VP performance management functions.

On each VP, you can configure an interval during which to transmit loopback cells by including the `oam-period` statement at the `[edit interfaces interface-name atm-options vpi vpi-identifier]` hierarchy level:

```
[edit interfaces interface-name atm-options vpi vpi-identifier]
oam-period (disable | seconds);
```

When you add a VPI at the `atm-options` hierarchy, an end-to-end F4 VCI is automatically opened to send and receive OAM F4, VP-AIS, and VP-RDI cells. If you enable OAM by including the `oam-period` statement in the configuration, the router sends and receives OAM F4 loopback cells.

If the physical ATM interface is configured with encapsulation type `atm-ccc-cell-relay`, then F4 VCIs are not created, and F4 OAM processing is not performed for the VPIs configured on that interface.

To modify OAM liveness values on a VP, include the `oam-liveness` statement at the `[edit interfaces interface-name atm-options vpi vpi-identifier]` hierarchy level:

```
[edit interfaces interface-name atm-options vpi vpi-identifier]
oam-liveness {
  up-count cells;
  down-count cells;
}
```

up-count is the minimum number of consecutive OAM F4 loopback cells received on a VPI before it is declared up.

down-count is the minimum number of consecutive OAM F4 loopback cells lost before a VPI is declared down.

When a VP-AIS or VP-RDI cell is received, the VPI is marked down. When a VP-AIS cell is received on a VPI, a VP-RDI is generated and transmitted on the same VPI. When an OAM F4 loopback request cell is received, the router sends a loopback reply cell, even if the `oam-period` statement is not included in the configuration of the VPI.

When a VPI is marked down because the VPI receives VP-AIS, VP-RDI, VC-AIS, or VC-RDI cells, or because the VPI does not receive down-count consecutive OAM F4 loopback replies, all the VCIs that belong to the VPI are marked down. When a VPI is marked up, all the VCIs that belong to the VPI are marked up. The status of logical interfaces is also changed when the status of the last VCI on that interface is changed.

For a configuration example, see “Example: Configuring ATM2 IQ Interfaces” on page 352.



NOTE: For interfaces that are configured for cell-relay promiscuous virtual path identifier (VPI) mode, the **show interfaces** command output does not show (OAM) F4 cell statistics.

Defining Virtual Path Tunnels

For ATM2 IQ interfaces, you can configure shaping on a VPI. When you do this, the VPI is called a VP tunnel. If your router is equipped with an ATM2 IQ PIC, you can configure VP tunnels and a weight for each VC. Each VC is serviced in WRR mode. When VCs have data to send, they send the number of cells equal to their weight before passing control to the next active VC. This allows proportional bandwidth sharing between multiple VCs within a rate-shaped VP tunnel. VP tunnels are not supported on point-to-multipoint interfaces.

If you change or delete VP tunnel traffic shaping, all logical interfaces on a VP are deleted and re-added.

All VPIs you configure on logical interfaces must also be configured on the physical interface, at the [edit interfaces *interface-name* atm-options] hierarchy level.

When you configure a VPI without shaping parameters, the VPI is a regular VPI; no shaping is attached. VCIs that belong to non-shaped VPIs can have VCI shaping.

For point-to-point interfaces, include the **shaping** statement at the [edit interfaces *interface-name* atm-options vpi *vpi-identifier*] hierarchy level:

```
[edit interfaces interface-name atm-options vpi vpi-identifier]
shaping {
  (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
  burst length);
  queue-length number;
}
```

For **cbr**, **vbr**, and **burst** statement usage guidelines, see “Defining the ATM Traffic-Shaping Profile” on page 319. For information about ATM2 IQ shaping values, see “Specifying ATM2 IQ Shaping Values” on page 325.

Configuring a Point-to-Point ATM1 or ATM2 IQ Connection

When you use ATM encapsulation on an interface, you must map each logical interface to a VCI. You can optionally map logical interfaces to a VPI.

For ATM1 and ATM2 IQ interfaces, you can configure a VCI and a VPI on a point-to-point ATM interface by including the `vci` statement:

```
vci vpi-identifier.vci-identifier;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For each VCI, configure the VCI and VPI identifiers. The default VPI identifier is 0. For ATM1 interfaces, the VCI identifier cannot exceed the highest-numbered VC configured for the interface with the `vpi` statement, as described in “Configuring the Maximum Number of ATM1 VCs on a VP” on page 300.

VCIs 0 through 31 are reserved for specific ATM values designated by the ATM Forum.

ATM2 IQ interfaces support only one invalid VC counter for all ports. The invalid VC counter is recorded at port 0 only.

When you are configuring point-to-point connections, the maximum transmission unit (MTU) sizes on both sides of the connections must be the same.

Configuring a Point-to-Multipoint ATM1 or ATM2 IQ Connection

An ATM interface can be a point-to-point interface or a point-to-multipoint (also called a multipoint nonbroadcast multiaccess [NBMA]) connection.

For ATM1 and ATM2 IQ interfaces, you can configure an NBMA ATM connection by including the following statements:

```

multipoint-destination address {
  epd-threshold cells;
  inverse-arp;
  oam-liveness {
    up-count cells;
    down-count cells;
  }
  oam-period (disable | seconds);
  shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
     rate burst length);
    queue-length number;
  }
  vci vpi-identifier.vci-identifier;
}

```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address*]

address is the interface's address. The address must include the destination prefix (for example, /24).

For each destination, include one **multipoint-destination** statement. *address* is the address of the remote side of the connection, and *vci-identifier* and *vpi-identifier* are the VCI and optional VPI identifiers for the connection.

When you configure point-to-multipoint connections, all interfaces in the subnet must use the same MTU size.

Configuring a Multicast-Capable ATM1 or ATM2 IQ Connection

For ATM1 and ATM2 IQ interfaces, you can configure a multicast-capable connection. By default, ATM connections assume unicast traffic. If your ATM switch performs multicast replication, you can configure the connection to support multicast traffic by including the **multipoint-vci** statement:

```
multipoint-vci vpi-identifier.vci-identifier;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

vci-identifier and *vpi-identifier* are the VCI and VPI identifiers, which define the ATM VCI over which the switch is expecting to receive multicast packets for replication.

You can configure multicast support only on point-to-multipoint ATM connections.

Configuring Inverse ATM1 or ATM2 ARP

For ATM1 and ATM2 IQ interfaces, you can configure inverse ATM Address Resolution Protocol (ARP), as described in RFC 2225, *Classical IP and ARP over ATM*. When inverse ATM ARP is enabled, the router responds to received inverse ATM ARP requests by providing IP address information to the requesting ATM device.

The router does not initiate inverse ATM ARP requests.

By default, inverse ATM ARP is disabled. To configure a VC to respond to inverse ATM ARP requests, include the **inverse-arp** statement:

```
inverse-arp;
```

For a list of hierarchy levels at which you can include this statement, see **inverse-arp**.

You must configure ATM LLC subnetwork attachment point (SNAP) encapsulation on the logical interface to support inverse ARP. No other ATM encapsulation types are allowed. For more information, see “Configuring ATM Interface Encapsulation” on page 330.

Defining the ATM Traffic-Shaping Profile

When you use an ATM encapsulation on ATM1 and ATM2 IQ interfaces, you can define bandwidth utilization, which consists of either a constant rate or a peak cell rate, with sustained cell rate and burst tolerance.

These values are used in the ATM generic cell-rate algorithm, which is a leaky bucket algorithm that defines the short-term burst rate for ATM cells, the maximum number of cells that can be included in a burst, and the long-term sustained ATM cell traffic rate.

If your router is equipped with an ATM2 IQ PIC, each VC can have independent shaping parameters. For more information, see “Defining Virtual Path Tunnels” on page 316.



NOTE: When the DS3 or E3 port parameters are not identical on all ports of a multiport ATM DS3 or E3 PIC, the ATM PIC driver might not always use the minimum port shaping rate (of all the ports on a multiport ATM DS3 or E3 PIC) selected for cell transmission shaping. The PIC's shaping rate is always updated to conform to the last port setting updated by the PIC software driver, rather than use the minimum port (shaping) rate. There is no syslog message to inform the user of the shaping rate decision applied by the software driver.

By default, the bandwidth utilization is unlimited; that is, unspecified bit rate (UBR) is used. Also, by default, buffer usage by VCs is unregulated.

To define limits to bandwidth utilization, include the **shaping** statement:

```
shaping {
  (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
   burst length);
  queue-length number;
}
```

For a list of hierarchy levels at which you can include this statement, see **shaping**.

The **rtvbr** statement is supported on ATM2 IQ PICs only. The **queue-length** statement is supported on ATM1 PICs only.

To configure VP tunnels on ATM2 IQ interfaces, include the **shaping** statement at the [edit interfaces *interface-name* atm-options vpi *vpi-identifier*] hierarchy level:

```
[edit interfaces interface-name atm-options vpi vpi-identifier]
shaping {
  (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
   burst length);
}
```

When configuring ATM traffic shaping, you can do the following:

- Configuring ATM CBR on page 320
- Configuring ATM2 IQ Real-Time VBR on page 321
- Configuring ATM VBR on page 321
- Specifying ATM1 Shaping Values on page 322
- Specifying ATM2 IQ Shaping Values on page 325

Configuring ATM CBR

For traffic that does not require the ability to periodically burst to a higher rate, you can specify a constant bit rate (CBR).

To specify a CBR on ATM1 and ATM2 IQ interfaces, include the **cbr** statement:

```
cbr rate;
```

For a list of hierarchy levels at which you can include this statement, see **cbr**. On J Series routers, ATM CBR shaping is not supported.

For ATM1 OC3 interfaces, the rate can be from 33 Kbps through 135.6 Mbps; for ATM1 OC12 interfaces, the rate can be from 33 Kbps through 276 Mbps.

For ATM2 IQ OC3 and OC12 interfaces, the rate can be from 33 Kbps through 542,526,792 bps.

For ATM2 IQ OC48 interfaces, the rate can be from 33 Kbps through 2,170,107,168 bps.

For ATM2 IQ DS3 and E3 interfaces, the rate can be from 33 Kbps to the maximum rate. The maximum rate varies depending on the ATM encapsulation and framing you configure, as shown in Table 27 on page 320.

Table 27: Shaping Rate Range by Interface Type

Interface Type	Maximum Rate
DS3 with direct ATM encapsulation	40,038,968 bps
DS3 with PLCP ATM encapsulation	36,864,000 bps
E3 with g.751 framing and direct ATM encapsulation	30,801,509 bps
E3 with g.751 framing PLCP ATM encapsulation	27,648,000 bps
E3 with g.832 framing	30,720,000 bps

Configuring ATM2 IQ Real-Time VBR

By default, ATM interfaces use UBR; that is, bandwidth utilization is unlimited. For ATM2 IQ interfaces only, you can configure RTVBR, which supports variable bit rate data traffic with average and peak traffic parameters. Compared to non-real-time VBR, RTVBR data is serviced at a higher priority with a relatively small sustainable cell rate (SCR) limit to minimize the delay. Real-time VBR is suitable for carrying packetized video and audio.

To configure RTVBR, include the **rtvbr** statement:

```
rtvbr peak rate sustained rate burst length;
```

For a list of hierarchy levels at which you can include this statement, see **rtvbr**.

When configuring RTVBR, you can define the following shaping properties:

- Peak rate—Top rate at which traffic can burst.
- Sustained rate—Normal traffic rate averaged over time.
- Burst length—Maximum number of cells that a burst of traffic can contain. It can be a value from 1 through 4000 cells.

The peak and sustained rates can be from 33 Kbps through 542,526,792 bps.

Configuring ATM VBR

By default, ATM interfaces use UBR; that is, bandwidth utilization is unlimited. For ATM1 and ATM2 IQ interfaces, you can configure non-real-time VBR, which supports variable bit rate data traffic with average and peak traffic parameters. Compared to RTVBR, non-real-time VBR is scheduled with a lower priority and with a larger SCR limit, allowing it to recover bandwidth if it falls behind. Non-real-time VBR is suitable for packet data transfers.

To define VBR on ATM1 and ATM2 IQ interfaces, include the **vbr** statement:

```
vbr peak rate sustained rate burst length;
```

For a list of hierarchy levels at which you can include this statement, see **vbr**.

When configuring VBR, you can define the following shaping properties:

- Peak rate—Top rate at which traffic can burst.
- Sustained rate—Normal traffic rate averaged over time.
- Burst length—Maximum number of cells that a burst of traffic can contain. It can be a value from 1 through 4000 cells.

Specifying ATM1 Shaping Values

For ATM1 interfaces, you can specify the rates in bits per second or cells per second. For OC3c interfaces, the highest rate is 135,631,698 bps (353,207.55 cps), which corresponds to 100 percent of the available line rate. For OC12c interfaces, the highest rate is 271,263,396 bps (706,415.09 cps), which corresponds to 50 percent of the available line rate. Table 28 on page 323 lists some of the other rates you can specify. If you specify a rate that is not listed, it is rounded to the nearest rate.

The exact number of values differs between OC12c and OC3c interfaces. OC12c interfaces have about four times as many value increments as OC3c interfaces.

For OC12c rates between 1/2 of the line rate and 1/128 of the line rate, there are 128 steps between each 1/*n* value. This means that there is 128 steps between the 1/2 and 1/3 line rate values, and another 128 steps between 1/3 and 1/4 and so on. For rates smaller than 1/127, there are (16,384 minus 127) or 16,257 values. The reason for this is that fractional shaping is ignored at rates below 1/127. This results in a total of about 32,384 distinct rates for OC12c. When *n* is larger than or equal to 127, the steps are 1/*n*.

For OC3c, the starting point is full line rate, the fraction/integer breakpoint is about 1/31, and there is a maximum of 4096 scheduler slots for use after 1/31 of line rate, producing about 8032 total distinct rates. When *n* is larger than or equal to 31, the steps are 1/*n*.

For ATM1 interfaces, the following formula can be used to predict the actual shaping rate:

- OC3 shaping settings between 135,631,698 bps (OC3 ATM cell line rate) and 4,375,216 bps (1/31 of OC3 ATM cell line rate).
- OC12 shaping settings between 271,263,396 bps (half OC12 ATM cell line rate – the highest rate supported) and 4,271,864 bps (1/127 of OC12 ATM cell line rate).

$$\text{actual-rate} = (128 * \text{line-rate}) / (\text{trunc} ((128 * \text{line-rate}) / \text{desired-rate}))$$

line-rate is the maximum available rate on the interface (in bits per second) after factoring out the overhead for SONET/SDH and ATM (per-cell) overheads. For OC3c interfaces, the line rate is calculated as follows:

$$\text{line-rate} = 155,520,000 \text{ bps} \times (26/27) \times (48/53) = 135,631,698.1 \text{ bps}$$

For OC12c interfaces, the line rate is calculated as follows:

$$\text{line-rate} = 622,080,000 \text{ bps} \times (26/27) \times (48/53) = 542,526,792.45 \text{ bps}$$

desired-rate is the rate you enter in the **vbr** statement, in bits per second.

The **trunc** operator indicates that all digits to the right of the decimal point should be dropped.

For shaping settings smaller than 1/31 of OC3 ATM cell line rate (4,375,216 bps) and 1/127 of OC12 ATM cell line rate (4,271,864 bps), you can predict the actual shaping rate using the following formula:

$$\text{actual-rate} = (1 / (\text{trunc} (\text{line-rate} / \text{desired-rate}) + 1)) * \text{line-rate}$$

For example, for OC12 interfaces, the actual rates for shaping below 4,271,864 bps are calculated as follows:

$$1 / 127 * 542,526,792.45 \text{ bps} = 4,271,864 \text{ bps (11124 cells/second)}$$

$$1 / 128 * 542,526,792.45 \text{ bps} = 4,238,490 \text{ bps (11038 cells/second)}$$

$$1 / 129 * 542,526,792.45 \text{ bps} = 4,205,634 \text{ bps (10952 cells/second)}$$

...

Buffers are shared among all VCs, and by default, there is no limit to the buffer size for a VC. If a VC is particularly slow, it might use all the buffer resources.

Table 28 on page 323 shows ATM1 traffic-shaping rates.

Table 28: ATM1 Traffic-Shaping Rates

Interface Type	Line Rate (bps)	Line Rate (cps)	Percentage of Total Line Rate
OC3			
	135,600,000	353,125	100.00
	134,542,320	350,370.66	99.22
	133,511,760	347,686.88	98.46
	132,494,760	345,038.44	97.71
	131,491,320	342,425.31	96.97
	130,501,440	339,847.5	96.24
	129,525,120	337,305	95.52
	128,562,360	334,797.81	94.81
	127,626,720	332,361.25	94.12
	126,691,080	329,924.69	93.43
OC12			
	271,263,396	706,415.09	50.00
	270,207,897	703,666.40	49.81
	269,160,579	700,939.01	49.61
	268,121,349	698,232.68	49.42
	267,090,113	695,547.17	49.23

Table 28: ATM1 Traffic-Shaping Rates (continued)

Interface Type	Line Rate (bps)	Line Rate (cps)	Percentage of Total Line Rate
	266,066,779	692,882.24	49.04
	265,051,257	690,237.65	48.85
	264,043,458	687,613.17	48.67
	263,043,293	685,008.58	48.48
	262,050,677	682,423.64	48.30

Example: Specifying ATM1 Shaping Values

Determine the actual rate in ATM1 interfaces when the desired rate is 80 percent of the maximum rate:

- OC3c:

$$135,600,000 \text{ bps} * 0.8 = 108,480,000 \text{ bps}$$

Because 108,480,000 bps is greater than 1/31 of OC3 ATM cell line rate:

$$\begin{aligned} \text{actual-rate} &= (128 * 135,600,000.1) / (\text{trunc} ((128 * 135,600,000.1) / \\ &\quad 108,480,000)) \\ \text{actual-rate} &= 17,356,800,013 / (\text{trunc} (17,356,800,013 / 108,480,000)) \\ \text{actual-rate} &= 17,356,800,013 / 160 \\ \text{actual-rate} &= 108,480,000 \text{ bps} \end{aligned}$$

- OC12c:

$$271,263,396 \text{ bps} * 0.8 = 217,010,716.8 \text{ bps}$$

Because 217,010,716.8 bps is greater than 1/127 of OC12 ATM cell line rate:

$$\begin{aligned} \text{actual-rate} &= (128 * 542,526,792.45) / (\text{trunc} ((128 * \\ &\quad 542,526,792.45) / 217,010,716.8)) \\ \text{actual-rate} &= 69,443,429,434 / (\text{trunc} (69,443,429,434 / 217,010,716.8)) \\ \text{actual-rate} &= 69,443,429,434 / 320 \\ \text{actual-rate} &= 217,010,717 \text{ bps} \end{aligned}$$

Determine the actual rate in ATM1 interfaces when the desired rate is 3,000,000 bps:

- OC3c:

Because 3,000,000 bps is smaller than 1/31 of OC3 ATM cell line rate:

$$\begin{aligned} \text{actual-rate} &= (1 / (\text{trunc} (\text{line-rate} / \text{desired-rate}) + 1)) * \text{line-rate} \\ \text{actual-rate} &= (1 / (\text{trunc} (135,631,698 / 3,000,000) + 1)) * 135,631,698 \\ \text{actual-rate} &= (1 / (45 + 1)) * 135,631,698 \\ \text{actual-rate} &= (1 / 46) * 135,631,698 \end{aligned}$$

actual-rate = 2,948,515 bps

■ OC12c:

Because 3,000,000 bps is smaller than 1/127 of OC12 ATM cell line rate:

actual-rate = (1 / (trunc (line-rate / desired-rate) + 1)) * line-rate
actual-rate = (1 / (trunc (542,526,792 / 3,000,000) + 1)) * 542,526,792
actual-rate = (1 / (180 + 1)) * 542,526,792
actual-rate = (1 / 181) * 542,526,792
actual-rate = 2,997,386 bps

Specifying ATM2 IQ Shaping Values

For ATM2 IQ OC3c interfaces, the maximum available rate is 100 percent of line rate, or 135,600,000 bps. For ATM2 IQ OC12c interfaces, the maximum available rate is 50 percent of line rate, or 271,273,396 bps. You can specify the rates in bits per second or cells per second. Fractional shaping is accurate within 0.5 percent of the desired rate.

Configuring the ATM1 Queue Length

ATM1 PICs contain a transmit buffer pool of 16,382 buffers, which are shared by all the PVCs that you configure on the PIC. Even multiple-port ATM PICs have a single buffer pool shared by all the ports.

By default, the ATM1 PIC allows PVCs to consume all the buffers they require. If the sustained traffic rate for a PVC exceeds its shaped rate, buffers are consumed. Eventually, all buffers on the PIC are consumed, and the other PVCs are underserved. This results in head-of-line blocking.

For each PVC, you prevent this situation by configuring the queue length of the PVC. The queue length is a limit on the number of transmit packets that can be queued. Packets that exceed the limit are dropped.

To limit the queue size of a PVC, include the **queue-length** statement:

queue-length *number*;

For a list of hierarchy levels at which you can include this statement, see **queue-length**.

The length can be from 1 through 16,383 packets. The default is 16,383 packets. You should include the **queue-length** statement in the configuration of all the PVCs that you configure on an ATM1 PIC. The **queue-length** statement performs two functions:

- It prevents head-of-line blocking because it limits the number of packets and therefore buffers that can be consumed by each configured PVC.
- It sets the maximum lifetime that can be sustained by packets over the PVC when traffic has oversubscribed the configured shaping contract.

The total value of all the queue lengths must not exceed the total number of packets that can be held in the buffer space available on the PIC. The total number of packets the buffers can hold depends on the size of the physical interface MTU, including all encapsulation overhead. You can use the following formula to calculate the total number of packets the buffer space can hold:

$$16,382 / (\text{Round Up} (\text{MTU} / 480))$$

For example, assuming default MTU settings for all ATM1 interfaces on a PIC, the total number of packets that can be held is:

$$16,382 / (\text{Round Up} (4482 / 480)) = 1638 \text{ packets}$$

Thus, you can configure up to 1638 for the combined queue length of all the PVCs on an ATM1 PIC that uses default MTU settings for all interfaces.

If you set a queue length to a very low value, small bursts in packets transiting the PVC might not be buffered.

The maximum lifetime that packets can sustain while transiting a PVC depends on the shaping rate you configure for the PVC, the setting for the `queue-length` statement, and the physical interface MTU. You can use the following formula to calculate the maximum lifetime that packets can sustain while transiting a PVC:

$$(\text{PVC queue-length in packets} \times \text{MTU}) / (\text{PVC shaping in bps} / 8)$$

For example, if you configure a PVC on an ATM1 interface with the default MTU, a CBR shaping rate of 3,840,000 bps (10,000 cps), and a queue length of 25 packets. The maximum lifetime is:

$$(25 \times 4482) / (3,840,000 / 8) = 233 \text{ ms}$$

This is the worst-case lifetime assuming all packets in the queue are MTU sized, and the traffic using the PVC is oversubscribing its configured shaping contract.

In general, we recommend that you use a maximum lifetime under 500 ms.

If you add or change the queue-length setting on the VC, the logical interface associated with the VC is deleted and re-added.

Configuring the ATM2 IQ EPD Threshold

The EPD threshold is a limit on the number of transmit cells that can be queued. Cells that exceed the limit are discarded. When a beginning of packet (BOP) cell is received, the VC's queue depth is checked against the EPD threshold. If the VC's queue depth exceeds the EPD threshold, the BOP cell and all subsequent cells in the packet are discarded. This prevents a single queue from draining all the buffers on the PIC.

By default, for UBR the EPD threshold is approximately 1 percent of the available cell buffers. If shaping is enabled, the default EPD threshold is proportional to the shaping rate according to the following formula:

$$\text{default epd-threshold} = \text{number of buffers} * \text{shaping rate} / \text{line rate}$$

By default, the software estimates how much buffer space is needed for each PVC. However, you can configure the per-VC buffer space. In general, ATM PVCs need larger buffers for data traffic and smaller buffers for time-sensitive applications. Unnecessarily deep buffers might cause excessive delays on congested PVCs. Overly shallow buffers might cause premature random early detection (RED) or tail packet drops in bursty conditions.

The minimum EPD threshold value is 48 cells. If the default EPD threshold formula results in an EPD threshold of less than 48 cells, the result will be ignored, and the minimum value of 48 cells will be used.

To set the EPD threshold of a PVC, include the **epd-threshold** statement:

```
epd-threshold cells;
```

For a list of hierarchy levels at which you can include this statement, see **epd-threshold**.

The allowable range for EPD threshold varies by interface type, as shown in Table 29 on page 327.

Table 29: EPD Threshold Range by Interface Type

Interface Type	EPD Range
1-port OC48	48 through 425,984 cells
1-port and 2-port OC12	48 through 425,984 cells
2-port OC3, DS3, and E3	48 through 212,992 cells
4-port DS3 and E3	48 through 106,496 cells

You should include the **epd-threshold** statement in the configuration of all the PVCs that you configure on an ATM2 IQ PIC. The **epd-threshold** statement performs two functions:

- It prevents head-of-line blocking because it limits the number of packets and therefore buffers that can be consumed by each configured PVC.
- It sets the maximum lifetime that can be sustained by packets over the PVC when traffic has oversubscribed the configured shaping contract.

If you add or change the EPD threshold on the VC, the logical interface associated with the VC is deleted and re-added.

On ATM2 IQ DS3 and E3 interfaces, you might be able to enter an EPD threshold or shaping parameter that exceeds the maximum threshold for these interfaces. If the configuration commits, the physical interface might indicate that it is up, but the logical interface fails. As a workaround, configure shaping parameters and EPD thresholds that do not exceed the bandwidth of the interface.

For information about configuring dual EPD thresholds on interfaces configured to use Layer 2 circuit trunk mode, see “Configuring Two EPD Thresholds per Queue” on page 328.

Example: Configuring the ATM2 IQ EPD Threshold

Configure the EPD threshold for a point-to-point ATM2 interface and a point-to-multipoint ATM2 interface.

On a Point-to-Point ATM2 Interface

```
[edit interfaces at-1/0/0]
unit 0 {
  vci 0.123;
  epd-threshold 1300;
  ...
}
```

On a Point-to-Multipoint ATM2 Interface

```
[edit interfaces at-1/0/1]
unit 0 {
  multipoint;
  family inet address 10.0.12.12/24 {
    multipoint-destination 10.0.12.14 vci 0.123 epd-threshold 1300;
    ...
  }
}
```

Configuring Two EPD Thresholds per Queue

For ATM2 IQ interfaces configured to use Layer 2 circuit trunk mode, you can set two EPD thresholds that depend on the PLPs of the packets. When you set a threshold with the `epd-threshold` statement, it applies to packets that have a PLP of 0. When you set a threshold with the `plp1` statement, it applies to packets that have a PLP of 1. If you include the `plp1` statement in the configuration, you must also include the `epd-threshold` statement.

To configure two EPD thresholds, include the `epd-threshold` and `plp1` statements:

```
epd-threshold cells plp1 cells;
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* atm-options scheduler-maps *map-name* forwarding-class *class-name*]
- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The value you set with the `epd-threshold` statement (for PLP0) should be equal to or greater than the value you set with the `plp1` statement. EPD threshold ranges vary by interface type. See Table 29 on page 327.

For general information about EPD thresholds, see “Configuring the ATM2 IQ EPD Threshold” on page 326.

Configuring the ATM2 IQ Transmission Weight

For ATM2 IQ interfaces configured with VPI shaping, you can control the number of cells a VCI can send each time the VCI has a turn to transmit by including the `transmit-weight` statement:

```
transmit-weight cells;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

VPI traffic shaping is not supported on point-to-multipoint interfaces.

The number of cells can be from 1 through 32,000. For a configuration example, see “Example: Configuring ATM2 IQ Interfaces” on page 352.

Defining the ATM OAM F5 Loopback Cell Period

For ATM1 and ATM2 IQ interfaces with an ATM encapsulation, you can configure the OAM F5 loopback cell period on virtual circuits. This is the interval at which OAM F5 loopback cells are transmitted.

By default, no OAM F5 loopback cells are sent. To send OAM F5 loopback cells, include the `oam-period` statement:

```
oam-period (disable | seconds);
```

For a list of hierarchy levels at which you can include this statement, see `oam-period`.

The period can be from 1 through 900 seconds. You can also choose the `disable` option to disable the OAM loopback cell transmit feature.

OAM VC-AIS and VC-RDI defect indication cells are used for identifying and reporting VC defects end-to-end. When a physical link or interface failure occurs, intermediate nodes insert OAM AIS cells into all the downstream VCs affected by the failure. Upon receiving an AIS cell on a VC, the router marks the logical interface down and sends an RDI cell on the same VC to notify the remote end of the error status. When an RDI cell is received on a VC, the router sets the logical interface status to down. When no AIS or RDI cells are received for 3 seconds, the router sets the logical interface status to up. You do not need to configure anything to enable defect indication.

Configuring the ATM OAM F5 Loopback Cell Threshold

For ATM1 and ATM2 IQ interfaces with an ATM encapsulation, you can configure the OAM F5 loopback cell threshold on VCs. This is the minimum number of

consecutive OAM F5 loopback cells received before a VC is declared up, or the minimum number of consecutive OAM F5 loopback cells lost before a VC is declared down.

By default, when five consecutive OAM F5 loopback cells are received, the VC is considered to be up, and when five consecutive cells are lost, the VC is considered to be down. To modify these values, include the **oam-liveness** statement:

```
oam-liveness {
  up-count cells;
  down-count cells;
}
```

For a list of hierarchy levels at which you can include this statement, see **oam-liveness**.

The cell count can be a value from 1 through 255.

Configuring ATM Interface Encapsulation

To configure ATM encapsulation on a physical interface, include the **encapsulation** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
encapsulation (atm-ccc-cell-relay | atm-pvc | ethernet-over-atm);
```

For ATM interfaces, the physical interface encapsulation can be one of the following:

- ATM cell-relay—This encapsulation connects two remote virtual circuits or ATM physical interfaces with an LSP. Traffic on the circuit is ATM cells.
- ATM PVC—ATM PVC encapsulation is defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*.
- Ethernet over ATM—As defined in RFC 1483 (the previous version of RFC 2684), this encapsulation type allows ATM interfaces to connect to devices that support only bridged-mode protocol data units (BPDUs). The JUNOS Software does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and media access control (MAC) header, and the packet is forwarded to the ATM interface.

Generally, you configure an interface's encapsulation at the [edit interfaces *interface-name*] hierarchy level. However, for ATM encapsulations, you can also configure the encapsulation type that is used inside the ATM cell itself. To do this, include the **encapsulation** statement:

```
encapsulation (atm-ccc-cell-relay | atm-ccc-vc-mux | atm-cisco-nlpid | atm-mlppp-llc |
  atm-nlpid | atm-ppp-llc | atm-ppp-vc-mux | atm-snap | atm-tcc-snap | atm-vc-mux |
  atm-tcc-vc-mux | ether-over-atm-llc | ether-vpls-over-atm-llc);
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Table 30 on page 331 shows the logical interface encapsulation types for ATM interfaces.

Table 30: ATM Logical Interface Encapsulation Types

Encapsulation Types	Comments
ATM CCC cell relay	<p>This encapsulation type connects two remote virtual circuits or ATM physical interfaces with an LSP.</p> <p>This encapsulation type carries traffic in ATM cells.</p> <p>When you use this encapsulation type, you can configure the <code>ccc</code> family only.</p>
ATM CCC VC multiplex	<p>This encapsulation type is for CCC circuits.</p> <p>When you use this encapsulation type, you can configure the <code>ccc</code> family only.</p>
ATM network layer protocol identifier (NLPID)	When you use this encapsulation type, you can configure the <code>inet</code> family only.
ATM SNAP	
ATM SNAP encapsulation on translational cross-connect (TCC) circuits	When you use this encapsulation type, you can configure the <code>tcc</code> family only.
ATM VC multiplex	When you use this encapsulation type, you can configure the <code>inet</code> family only.
ATM VC multiplex on TCC circuits	When you use this encapsulation type, you can configure the <code>tcc</code> family only.
Cell-relay accumulation mode (CAM)	<p>In this mode, the incoming 1 to 8 cells are packaged into a single packet and forwarded to the LSP. To configure CAM, include the <code>atm-cell-relay-accumulation</code> statement at the [edit chassis fpc <i>slot-number</i> pic <i>pic-number</i>] hierarchy level.</p> <p>This encapsulation type is for ATM1 interfaces only.</p> <p>For more information about CAM, see the <i>JUNOS System Basics Configuration Guide</i>.</p>
Cisco ATM NLPID	When you use this encapsulation type, you can configure the <code>inet</code> family only.
Ethernet over ATM	<p>This encapsulation type is for interfaces that carry IPv4 traffic.</p> <p>When you use this encapsulation type, you cannot configure point-to-multipoint interfaces.</p>

Table 30: ATM Logical Interface Encapsulation Types (*continued*)

Encapsulation Types	Comments
Ethernet VPLS over ATM	<p>This encapsulation type enables a VPLS instance to support bridging between Ethernet interfaces and ATM interfaces, as described in RFC 2684.</p> <p>Use this encapsulation type to support IEEE 802.1p classification binding on ATM VCs.</p> <p>This encapsulation type is for ATM2 IQ interfaces only.</p> <p>When you use this encapsulation type, you cannot configure point-to-multipoint interfaces.</p>
Multilink PPP over AAL5 LLC	<p>This encapsulation type is for ATM2 IQ interfaces only.</p> <p>When you use this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC.</p>
PPP over AAL5 LLC	<p>This encapsulation type is for ATM2 IQ interfaces only.</p> <p>When you use this encapsulation type, you cannot configure point-to-multipoint interfaces.</p>
PPP over AAL5 multiplex	<p>This encapsulation type is for ATM2 IQ interfaces only.</p> <p>When you use this encapsulation type, you cannot configure point-to-multipoint interfaces.</p>

Configuring an ATM1 Cell-Relay Circuit

For ATM1 interfaces, you can create an ATM cell-relay circuit by configuring an entire ATM physical device or an individual VC. When you configure an entire device, only cell-relay encapsulation is allowed on the logical interfaces; for ATM1 PICs, you use the **atm-options** statement to control the number and location of VCs. The configuration of allowed VCs on both ingress and egress ATM interfaces should be the same. For most interfaces, you can define a maximum of 4090 VCs per interface. The highest-numbered VC value you can configure is 4089. Promiscuous mode removes these limits. For more information, see “Configuring ATM Cell-Relay Promiscuous Mode” on page 296.

For ATM1 interfaces, if you are dedicating the entire device to a cell-relay circuit, include the **allow-any-vci** statement in the configuration of **unit 0**:

```
allow-any-vci;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit 0]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit 0]

Once you include this statement, you cannot configure other logical interfaces in the same physical interface.



NOTE: When you use ATM CCC cell-relay encapsulation, you must configure the logical encapsulation as `atm-ccc-cell-relay`. You cannot mix different logical encapsulation types on an interface that you have configured with ATM CCC cell-relay physical encapsulation.

Example: Configuring an ATM1 Cell-Relay Circuit

Configure an ATM1 cell-relay circuit:

	<pre>[edit interfaces at-1/2/0] encapsulation atm-ccc-cell-relay; atm-options { pic-type atm1; vpi 0 maximum-vcs 256; } unit 0 { point-to-point; encapsulation atm-ccc-cell-relay; allow-any-vci; }</pre>
Configuring an Individual VC on a Logical Interface	<pre>[edit interfaces at-1/1/0] encapsulation atm-ccc-cell-relay; atm-options { pic-type atm1; vpi 0 maximum-vcs 256; } unit 120 { encapsulation atm-ccc-cell-relay; vci 0.120; }</pre>
Configuring Nonpromiscuous Port Mode	<pre>[edit interfaces at-0/0/1] encapsulation atm-ccc-cell-relay; atm-options { pic-type atm1; vpi 0 { maximum-vcs 100; } vpi 1 { maximum-vcs 300; } vpi 4 { maximum-vcs 200; } } unit 0 { encapsulation atm-ccc-cell-relay; allow-any-vci; }</pre>

Configuring Nonpromiscuous VPI Mode	<pre>[edit interfaces at-0/0/1] encapsulation atm-ccc-cell-relay; atm-options { pic-type atm1; vpi 0 { maximum-vcs 100; } } unit 0 { encapsulation atm-ccc-cell-relay; vpi 0; }</pre>
Configuring Nonpromiscuous VCI Mode	<pre>[edit interfaces at-0/0/1] encapsulation atm-ccc-cell-relay; atm-options { pic-type atm1; vpi 0 { maximum-vcs 100; } } unit 0 { encapsulation atm-ccc-cell-relay; vci 0.50 }</pre>

Configuring PPP over ATM2 Encapsulation

For ATM2 IQ interfaces, you can configure PPP over AAL5 encapsulation, as described in RFC 2364, *PPP over AAL5*. PPP over ATM encapsulation associates a PPP link with an ATM AAL5 PVC.

The JUNOS Software supports three PPP over ATM encapsulation types:

- **atm-ppp-llc**—PPP over AAL5 LLC.
- **atm-ppp-vc-mux**—PPP over ATM AAL5 multiplex.
- **atm-mlppp-llc**—Multilink PPP over ATM AAL5 LLC. For this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.

To enable PPP over ATM encapsulation, include the **encapsulation** statement, specifying the **atm-mlppp-llc**, **atm-ppp-llc**, or **atm-ppp-vc-mux** encapsulation type:

```
encapsulation (atm-mlppp-llc | atm-ppp-llc | atm-ppp-vc-mux);
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

When you configure PPP over ATM encapsulation, you can enable PPP Challenge Handshake Authentication Protocol (CHAP) and keepalives on the logical interface. For more information about PPP CHAP and keepalives, see “Configuring the PPP Challenge Handshake Authentication Protocol” on page 112 and “Configuring Keepalives” on page 126.



NOTE: When you use PPP over ATM encapsulation, we recommend that you not include the `oam-period` statement in the configuration. Instead, we recommend that you enable keepalives to detect connection failures.

Example: Configuring PPP over ATM2 IQ Encapsulation

Configure three logical interfaces with PPP over ATM encapsulation:

```
[edit interfaces]
at-0/1/0 {
  atm-options {
    pic-type atm2;
    vpi 0;
    vpi 2;
  }
  unit 0 {
    encapsulation atm-ppp-llc;
    ppp-options {
      chap {
        access-profile pe-B-ppp-clients;
        local-name "pe-A-at-0/1/0";
      }
    }
    keepalives interval 5 up-count 6 down-count 4;
    vci 0.120;
    family inet address 192.168.13.13/30;
  }
  unit 1 {
    encapsulation atm-ppp-vc-mux;
    vci 2.120;
    keepalives interval 6 up-count 6 down-count 4;
    family inet address 192.168.14.13/30;
  }
  unit 2 {
    encapsulation atm-ppp-vc-mux;
    ppp-options {
      chap {
        passive;
        access-profile pe-A-ppp-clients;
        local-name "pe-A-at-0/1/0";
      }
    }
    keepalives interval 5 up-count 6 down-count 4;
    vci 2.121;
    family inet address 192.168.15.13/30;
  }
}
```

**Configuring Multilink
PPP over ATM2 IQ
Encapsulation**

```
[edit interfaces]
at-0/0/0 {
  atm-options {
    pic-type atm2;
    vpi 10;
  }
  unit 0 {
    encapsulation atm-mlppp-llc;
    ppp-options {
      chap {
        access-profile pe-B-ppp-clients;
        local-name " pe-A-at-0/0/0";
      }
    }
    keepalive interval 5 up-count 6 down-count 4;
    vci 10.120;
    family mlppp {
      bundle ls-0/3/0.0;
    }
  }
}
at-0/0/1 {
  atm-options {
    pic-type atm2;
    vpi 11;
  }
  unit 1 {
    encapsulation atm-mlppp-llc;
    ppp-options {
      chap {
        access-profile pe-B-ppp-clients;
        local-name " pe-A-at-0/0/0";
      }
    }
    keepalive interval 5 up-count 6 down-count 4;
    vci 11.120;
    family mlppp {
      bundle ls-0/3/0.0;
    }
  }
}
at-1/2/3 {
  atm-options {
    pic-type atm2;
    vpi 12;
  }
  unit 2 {
    encapsulation atm-mlppp-llc;
    ppp-options {
      chap {
        access-profile pe-B-ppp-clients;
        local-name " pe-A-at-0/0/0";
      }
    }
    keepalive interval 5 up-count 6 down-count 4;
    vci 12.120;
```



```

        family mlppp {
            bundle ls-0/3/0.0;
        }
    }
    ...
ls-0/3/0 {
    encapsulation multilink-ppp;
    interleave-fragments;
    keepalive;
    unit 0 {
        mrru 4500;
        short-sequence;
        fragment-threshold 16320;
        drop-timeout 2000;
        encapsulation multilink-ppp;
        interleave-fragments;
        minimum-links 8;
        family inet {
            address 10.10.0.1/32 {
                destination 10.10.0.2;
            }
        }
        family iso;
        family inet6 {
            address 8090::0:1/128 {
                destination 8090::0:2;
            }
        }
    }
    ...
}

```

Configuring E3 and T3 Parameters on ATM Interfaces

For ATM1 and ATM2 IQ interfaces, you can configure ATM E3 and T3 interfaces by including the following statements at the [edit interfaces *at-fpc/pic/port*] hierarchy level:

```

[edit interfaces at-fpc/pic/port]
e3-options {
    atm-encapsulation (direct | plcp);
    buildout feet;
    framing (g.751 | g.832);
    loopback (local | remote);
    (payload-scrambler | no-payload-scrambler);
}
t3-options {
    atm-encapsulation (direct | plcp);
    buildout feet;
    (cbit-parity | no-cbit-parity);
    loopback (local | payload | remote);
    (payload-scrambler | no-payload-scrambler);
}

```

The following options and default values differ from those described in the chapters “Configuring E3 Interfaces” on page 551 and “Configuring T3 Interfaces” on page 569:

- **atm-encapsulation**—PLCP is the default value. The E3 **line-format** option g.832 supports the **direct** ATM-encapsulation option only.
- **buildout**—The default value is 10 feet. The number of feet can be any integer value. The range is from 0 through 450 feet (about 137 meters).
- **cbit-parity**—The default option is to enable cbit parity.
- **framing**—There is no default option for E3 interfaces; T3 interfaces use the **cbit-parity** statement in place of the **framing** statement.
- **loopback**—By default, loopback is disabled.
- **payload-scrambler**—The default option is to enable payload scrambling.

In addition, the ATM E3 and T3 PICs support the **clocking** statement at the interface level, as do the SONET/SDH PICs. For more information about E3- and T3-specific parameters, see “Configuring E3 Interfaces” on page 551 and “Configuring T3 Interfaces” on page 569.



NOTE: You must configure all the ports on an ATM E3 or T3 PIC with the same framing and encapsulation. Otherwise, the system will set all the ports on the PIC to the slowest framing and encapsulating configuration. For ATM T3, this is PLCP. For ATM E3, this is G.751 PLCP.

Configuring SONET/SDH Parameters on ATM Interfaces

When configuring ATM1 and ATM2 IQ SONET/SDH interfaces, you can also include the following statements in the **sonet-options** statement to set SONET/SDH parameters on ATM interfaces:

```
[edit interfaces at-fpc/pic/port]
sonet-options {
  aps {
    advertise-interval milliseconds;
    authentication-key key;
    force;
    hold-time milliseconds;
    lockout;
    neighbor address;
    paired-group group-name;
    protect-circuit group-name;
    request;
    revert-time seconds;
    working-circuit group-name;
  }
  bytes {
    e1-quiet value;
    f1 value;
    f2 value;
    s1 value;
  }
}
```

```

        z3 value;
        z4 value;
    }
    loopback (local | remote);
    (payload-scrambler | no-payload-scrambler);
    rfc-2615;
    trigger {
        defect ignore {
            hold-time up milliseconds down milliseconds;
        }
    }
    (z0-increment | no-z0-increment);
}

```

For information about configuring specific SONET/SDH statements, see “Configuring SONET/SDH Interfaces” on page 843.

Configuring ATM2 IQ VC Tunnel CoS Components

The ATM2 IQ interface allows multiple IP queues into each VC. On M Series routers (except the M320 and M120 router), a VC tunnel can support four CoS queues. On the M320, M120, and T Series routers for all ATM2 IQ PICs except the OC48 PIC, a VC tunnel can support eight CoS queues. Within a VC tunnel, the WRR algorithm schedules the cell transmission of each queue. You can configure the queue admission policies, such as EPD or WRED, to control the queue size during congestion.

For information about CoS components that apply generally to all interfaces, see the *JUNOS Class of Service Configuration Guide*.

To configure ATM2 IQ VC tunnel CoS components, include the following statements at the [edit interfaces at-*fpc/pic/port*] hierarchy level:

```

[edit chassis fpc slot-number pic pic-number]
max-queues-per-interface number;
[edit interfaces at-fpc/pic/port]
atm-options {
    linear-red-profiles profile-name {
        high-plp-max-threshold percent;
        low-plp-max-threshold percent;
        queue-depth cells high-plp-threshold percent low-plp-threshold percent;
    }
    plp-to-clp;
    scheduler-maps map-name {
        forwarding-class class-name {
            epd-threshold cells plp1 cells;
            linear-red-profile profile-name;
            priority (high | low);
            transmit-weight (cells number | percent number);
        }
        vc-cos-mode (alternate | strict);
    }
}
unit 0 {
    atm-scheduler-map (map-name | default);
}

```

```

family family {
    address address {
        destination address;
    }
}
plp-to-clp;
shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate burst length);
}
vci vpi-identifier.vci-identifier;
}

```

This section contains the following topics:

- Configuring Linear RED Profiles on page 340
- Configuring an ATM Scheduler Map on page 341
- Enabling Eight Queues on ATM2 IQ Interfaces on page 342
- Configuring VC CoS Mode on page 348
- Enabling the PLP Setting to Be Copied to the CLP Bit on page 348
- Configuring ATM CoS on the Logical Interface on page 349
- Example: Configuring ATM2 IQ VC Tunnel CoS Components on page 349

Configuring Linear RED Profiles

Linear RED profiles define CoS virtual circuit drop profiles. You can configure up to 32 linear RED profiles per port. When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.

To configure linear RED profiles, include the `linear-red-profiles` statement at the `[edit interfaces at-fpc/pic/port atm-options]` hierarchy level:

```

[edit interfaces at-fpc/pic/port atm-options]
linear-red-profiles profile-name {
    high-plp-max-threshold percent;
    low-plp-max-threshold percent;
    queue-depth cells high-plp-threshold percent low-plp-threshold percent;
}

```

The `queue-depth`, `high-plp-threshold`, and `low-plp-threshold` statements are mandatory.

You can define the following options for each RED profile:

- **high-plp-max-threshold**—Define the drop profile fill-level for the high PLP CoS VC. When the fill level exceeds the defined percentage, all packets with high PLP are dropped.
- **low-plp-max-threshold**—Define the drop profile fill-level for the low PLP CoS VC. When the fill level exceeds the defined percentage, all packets with low PLP are dropped.

- **queue-depth**—Define maximum queue depth in the CoS VC drop profile. Packets are always dropped beyond the defined maximum. The range you can configure is from 1 through 64,000 cells.
- **high-plp-threshold**—Define CoS VC drop profile fill-level percentage when linear RED is applied to cells with high PLP. When the fill level exceeds the defined percentage, packets with high PLP are randomly dropped by RED.
- **low-plp-threshold**—Define CoS VC drop profile fill-level percentage when linear RED is applied to cells with low PLP. When the fill level exceeds the defined percentage, packets with low PLP are randomly dropped by RED.

Configuring an ATM Scheduler Map

To define a scheduler map, you associate it with a forwarding class. Each class is associated with a specific queue, as follows:

- **best-effort**—Queue 0
- **expedited-forwarding**—Queue 1
- **assured-forwarding**—Queue 2
- **network-control**—Queue 3



NOTE: For M320, M120, and T Series routers only, you can configure more than four forwarding classes and queues. For more information, see “Enabling Eight Queues on ATM2 IQ Interfaces” on page 342.

When you configure an ATM scheduler map, the JUNOS Software creates these CoS queues for a VC. The JUNOS Software prefixes each packet delivered to the VC with the next-hop rewrite data associated with each queue.

To configure an ATM scheduler map, include the **scheduler-maps** statement at the [edit interfaces at-*fpc/pic/port* atm-options] hierarchy level:

```
edit interfaces at-fpc/pic/port atm-options]
scheduler-maps map-name {
  forwarding-class class-name {
    epd-threshold cells plp1 cells;
    linear-red-profile profile-name;
    priority (high | low);
    transmit-weight (cells number | percent number);
  }
}
```

You can define the following options for each forwarding class:

- **epd-threshold** or **linear-red-profile**—An EPD threshold provides a queue of cells that can be stored with tail drop. When a BOP cell is received, the VC's queue depth is checked against the EPD threshold. If the VC's queue depth exceeds the EPD threshold, the BOP cell and all subsequent cells in the packet are discarded.

A linear RED profile defines the number of cells using the **queue-depth** statement within the RED profile. (You configure the **queue-depth** statement at the [edit interfaces *at-fpc/pic/port* atm-options linear-red-profiles *profile-name*] hierarchy level.)

By default, if you include the **scheduler-maps** statement at the [edit interfaces *at-fpc/pic/port* atm-options] hierarchy level, the interface uses an EPD threshold that is determined by the JUNOS Software based on the available bandwidth and other parameters. You can override the default EPD threshold by setting an EPD threshold or a linear RED profile.

- **priority**—By default, queue 0 is high-priority, and the remaining queues are low-priority. You can configure high or low queuing priority for each queue.
- **transmit-weight**—By default, the transmit weight is 95 percent for queue 0, and 5 percent for queue 3. You can configure the transmission weight in number of cells or percentage. Each CoS queue is serviced in WRR mode. When CoS queues have data to send, they send the number of cells equal to their weight before passing control to the next active CoS queue. This allows proportional bandwidth sharing between multiple CoS queues within a rate-shaped VC tunnel. A CoS queue can send from 1 through 32,000 cells or from 5 through 100 percent of queued traffic before passing control to the next active CoS queue within a VC tunnel.

The AAL5 protocol prohibits cells from being interleaved on a VC; therefore, a complete packet is always sent. If a CoS queue sends more cells than its assigned weight because of the packet boundary, the deficit is carried over to the next time the queue is scheduled to transmit. If the queue is empty after the cells are sent, the deficit is waived, and the queue's assigned weight is reset.



NOTE: If you include the **scheduler-maps** statement at the [edit interfaces *at-fpc/pic/port* atm-options] hierarchy level, the **epd-threshold** statement at the [edit interfaces *interface-name* unit *logical-unit-number*] or [edit interfaces *interface-name* unit *logical-unit-number* address *address* family *family* multipoint-destination *address*] hierarchy level has no effect because either the default EPD threshold, the EPD threshold setting in the forwarding class, or the linear RED profile takes effect instead.

For more information about forwarding classes, see the *JUNOS Class of Service Configuration Guide*.

Enabling Eight Queues on ATM2 IQ Interfaces

By default, ATM2 IQ PICs on T Series, M120, and M320 routers are restricted to a maximum of four egress queues per interface. You can enable eight egress queues on ATM2 IQ interfaces by including the **max-queues-per-interface** statement at the [edit chassis *fpc slot-number* *pic pic-number*] hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
max-queues-per-interface number;
```

The numerical value can be 4 or 8.

If you include the `max-queues-per-interface` statement, all ports on the ATM2 IQ PIC use the configured mode.

When you include the `max-queues-per-interface` statement and commit the configuration, all physical interfaces on the ATM2 IQ PIC are deleted and re-added. Also, the PIC is taken offline and then brought back online immediately. You do not need to manually take the PIC offline and online. You should change modes between four queues and eight queues, or vice versa, only when there is no active traffic going to the ATM2 IQ PIC.

To configure up to eight queues on the ATM2 IQ interface, you must also include the statements described in “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.

For general information about configuring up to eight forwarding classes and queues on PICs other than ATM2 IQ PICs, see the *JUNOS Class of Service Configuration Guide*.




NOTE: When you are considering enabling eight queues on an ATM2 IQ interface, you should note the following:

- ATM2 IQ interfaces using Layer 2 circuit trunk transport mode support only four CoS queues.
- ATM2 IQ OC48 interfaces support only four CoS queues.
- ATM2 IQ interfaces with MLPPP encapsulation support only four CoS queues.
- You can configure only four RED profiles for the eight queues. Thus, queue 0 and queue 4 share a single RED profile, as do queue 1 and queue 5, queue 2 and queue 6, and queue 3 and queue 7. There is no restriction on EPD threshold per queue.
- The default chassis scheduler allocates resources for queue 0 through queue 3, with 25 percent of the bandwidth allocated to each queue. When you configure the chassis to use more than four queues, you must configure and apply a custom chassis scheduler to override the default. To apply a custom chassis scheduler, include the `scheduler-map-chassis` statement at the `[edit class-of-service interfaces at-fpc/pic/*]` hierarchy level. For more information about configuring and applying a custom chassis scheduler, see the *JUNOS Class of Service Configuration Guide*.

Example: Enabling Eight Queues on T Series, M120, and M320 Platforms

In Figure 21 on page 344, Router A generates IP packets with different IP precedence settings. Router B is an M320, M120, or T Series router with two ATM2 IQ interfaces. On Router B, interface `at-6/1/0` receives traffic from Router A, while interface `at-0/1/0` sends traffic to Router C. This example shows the CoS configuration for Router B.

Figure 21: Example Topology for Router with Eight Queues

On Router B:

```

[edit chassis]
fpc 0 {
  pic 1 {
    max-queues-per-interface 8;
  }
}
fpc 6 {
  pic 1 {
    max-queues-per-interface 8;
  }
}
[edit interfaces]
at-0/1/0 {
  atm-options {
    linear-red-profiles {
      red_1 queue-depth 1k high-plp-threshold 50 low-plp-threshold 80;
      red_2 queue-depth 2k high-plp-threshold 40 low-plp-threshold 70;
      red_3 queue-depth 3k high-plp-threshold 30 low-plp-threshold 60;
      red_4 queue-depth 4k high-plp-threshold 20 low-plp-threshold 50;
    }
    scheduler-maps {
      sch_red {
        vc-cos-mode strict;
        forwarding-class fc_q0 {
          priority high;
          transmit-weight percent 5;
          linear-red-profile red_1;
        }
        forwarding-class fc_q1 {
          priority low;
          transmit-weight percent 10;
          linear-red-profile red_2;
        }
        forwarding-class fc_q2 {
          priority low;
          transmit-weight percent 15;
          linear-red-profile red_3;
        }
        forwarding-class fc_q3 {
          priority low;
          transmit-weight percent 20;
          linear-red-profile red_4;
        }
        forwarding-class fc_q4 {
          priority low;
          transmit-weight percent 5;
          linear-red-profile red_1;
        }
      }
    }
  }
}
  
```



```

}
forwarding-class fc_q5 {
    priority low;
    transmit-weight percent 10;
    linear-red-profile red_2;
}
forwarding-class fc_q6 {
    priority low;
    transmit-weight percent 15;
    linear-red-profile red_3;
}
forwarding-class fc_q7 {
    priority low;
    transmit-weight percent 20;
    linear-red-profile red_4;
}
}
sch_epd {
    vc-cos-mode alternate;
    forwarding-class fc_q0 {
        priority high;
        transmit-weight percent 5;
        epd-threshold 1024;
    }
    forwarding-class fc_q1 {
        priority low;
        transmit-weight percent 10;
        epd-threshold 2048;
    }
    forwarding-class fc_q2 {
        priority low;
        transmit-weight percent 15;
        epd-threshold 3072;
    }
    forwarding-class fc_q3 {
        priority low;
        transmit-weight percent 20;
        epd-threshold 4096;
    }
    forwarding-class fc_q4 {
        priority low;
        transmit-weight percent 5;
        epd-threshold 2048;
    }
    forwarding-class fc_q5 {
        priority low;
        transmit-weight percent 10;
        epd-threshold 3072;
    }
    forwarding-class fc_q6 {
        priority low;
        transmit-weight percent 15;
        epd-threshold 4096;
    }
    forwarding-class fc_q7 {
        priority low;

```

```

        transmit-weight percent 20;
        epd-threshold 5120;
    }
}
}
atm-options {
    vpi 0;
}
unit 0 {
    vci 0.100;
    shaping {
        cbr 1920000;
    }
    atm-scheduler-map sch_red;
    family inet {
        address 172.16.0.1/24;
    }
}
unit 1 {
    vci 0.101;
    shaping {
        vbr peak 1m sustained 384k burst 256;
    }
    atm-scheduler-map sch_epd;
    family inet {
        address 172.16.1.1/24;
    }
}
}
at-6/1/0 {
    atm-options {
        vpi 0;
    }
    unit 0 {
        vci 0.100;
        family inet {
            address 10.10.0.1/24;
        }
    }
    unit 1 {
        vci 0.101;
        family inet {
            address 10.10.1.1/24;
        }
    }
}
[edit class-of-service]
classifiers {
    inet-precedence inet_classifier {
        forwarding-class fc_q0 {
            loss-priority low code-points 000;
        }
        forwarding-class fc_q1 {
            loss-priority low code-points 001;
        }
    }
}

```

```

forwarding-class fc_q2 {
    loss-priority low code-points 010;
}
forwarding-class fc_q3 {
    loss-priority low code-points 011;
}
forwarding-class fc_q4 {
    loss-priority low code-points 100;
}
forwarding-class fc_q5 {
    loss-priority low code-points 101;
}
forwarding-class fc_q6 {
    loss-priority low code-points 110;
}
forwarding-class fc_q7 {
    loss-priority low code-points 111;
}
}
forwarding-classes {
    queue 0 fc_q0;
    queue 1 fc_q1;
    queue 2 fc_q2;
    queue 3 fc_q3;
    queue 4 fc_q4;
    queue 5 fc_q5;
    queue 6 fc_q6;
    queue 7 fc_q7;
}
interfaces {
    at-6/1/0 {
        unit * {
            classifiers {
                inet-precedence inet_classifier;
            }
        }
    }
}
}
[edit routing-options]
static {
    route 10.10.20.2/32 {
        next-hop at-0/1/0.0;
        retain;
        no-readvertise;
    }
    route 10.10.1.2/32 {
        next-hop at-0/1/0.1;
        retain;
        no-readvertise;
    }
}
}

```

Verifying the Configuration

To see the results of this configuration, you can issue the following operational mode commands:

- show interfaces at-0/1/0 extensive

- `show interfaces queue at-0/1/0`
- `show class-of-service forwarding-class`

Configuring VC CoS Mode

VC CoS mode defines the CoS queue scheduling priority. By default, the VC CoS mode is alternate. When it is a queue's turn to transmit, the queue transmits up to its weight in cells as specified by the `transmit-weight` statement at the `[edit interfaces at-fpc/pic/port atm-options scheduler-maps map-name forwarding-class class-name]` hierarchy level. The number of cells transmitted can be slightly over the configured or default transmit weight, because the transmission always ends at a packet boundary.

To configure the VC CoS mode, include the `vc-cos-mode` statement at the `[edit interfaces at-fpc/pic/port atm-options scheduler-maps]` hierarchy level:

```
edit interfaces at-fpc/pic/port atm-options scheduler-maps]
vc-cos-mode (alternate | strict);
```

Two modes of CoS scheduling priority are supported:

- **alternate**—Assign high priority to one queue. The scheduling of the queues alternates between the high priority queue and the remaining queues. Every other scheduled packet is from the high priority queue.
- **strict**—Assign strictly high priority to one queue. A queue with strictly high priority is always scheduled before the remaining queues. The remaining queues are scheduled in round-robin fashion.

Enabling the PLP Setting to Be Copied to the CLP Bit

For a PE router with customer edge (CE)-facing, egress, ATM2 IQ interfaces configured with standard AAL5 encapsulation, you can enable the PLP setting to be copied into the CLP bit.



NOTE: This configuration setting is not applicable to Layer 2 circuit encapsulations because the control word captures and preserves CLP information. For more information about Layer 2 circuit encapsulations, see “Configuring Layer 2 Circuit Transport Mode” on page 300.

By default, at egress ATM2 IQ interfaces configured with standard AAL5 encapsulation, the PLP information is not copied to the CLP bit. This means the PLP information is not carried beyond the egress interface onto the CE router.

You can enable the PLP information to be copied into the CLP bit by including the `plp-to-clp` statement:

```
plp-to-clp;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* atm-options]
- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Configuring ATM CoS on the Logical Interface

To apply the ATM scheduler map to a logical interface, include the `atm-scheduler-map` statement:

```
atm-scheduler-map (map-name | default);
```

For ATM CoS to take effect, you must configure the VCI and VPI identifiers and traffic shaping on each VC by including the following statements:

```
vci vpi-identifier.vci-identifier;
shaping {
  (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
   burst length);
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For more information, see “Configuring a Point-to-Point ATM1 or ATM2 IQ Connection” on page 316 and “Defining the ATM Traffic-Shaping Profile” on page 319.

You can also apply a scheduler map to the chassis traffic that feeds the ATM interfaces. For more information, see the *JUNOS Class of Service Configuration Guide*.

Example: Configuring ATM2 IQ VC Tunnel CoS Components

Configure ATM2 IQ VC tunnel CoS components:

```
[edit interfaces]
at-1/2/0 {
  atm-options {
    vpi 0;
    linear-red-profiles red-profile-1 {
      queue-depth 35000 high-plp-threshold 75 low-plp-threshold 25;
    }
  }
  scheduler-maps map-1 {
    vc-cos-mode strict;
    forwarding-class best-effort {
      priority low;
      transmit-weight percent 25;
      linear-red-profile red-profile-1;
    }
  }
}
```

```

    }
  }
  unit 0 {
    vci 0.128;
    shaping {
      vbr peak 20m sustained 10m burst 20;
    }
    atm-scheduler-map map-1;
    family inet {
      address 192.168.0.100/32 {
        destination 192.168.0.101;
      }
    }
  }
}

```

Example: Configuring ATM1 Interfaces

The following configuration is sufficient to get an ATM1 OC3 or OC12 interface up and running. By default, ATM interfaces use ATM PVC encapsulation.

```

[edit interfaces]
at-fpc/pic/port {
  atm-options {
    vpi vpi-identifier maximum-vcs maximum-vcs-value;
    unit 0 { # one unit per VC
      vci vpi-identifier.vci-identifier;
      family inet {
        address local-address {
          destination address;
        }
      }
    }
  }
  unit 1 {# second VC
    ...
  }
}

```

Complex Configuration Example

```

[edit interfaces]
at-0/0/0 {
  encapsulation atm-pvc;
  atm-options {
    vpi 0 maximum-vcs 1200;
  }
  unit 2 {
    encapsulation atm-snap;
    inverse-arp;
    vci 0.80;
    family inet {
      mtu 1500;
      address 192.168.0.3/32 {
        destination 192.168.0.1;
      }
    }
  }
  unit 3 {

```

```

        encapsulation atm-snap;
        vci 0.32;
        oam-period 60;
        family inet {
            mtu 1500;
            address 192.168.4.3/32 {
                destination 192.168.4.2;
            }
        }
    }
}
at-0/2/0 {
    encapsulation atm-pvc;
    atm-options {
        vpi 0 maximum-vcs 1200;
    }
    unit 2 {
        encapsulation atm-snap;
        inverse-arp;
        vci 0.82;
        family inet {
            mtu 1500;
            address 192.168.5.3/32 {
                destination 192.168.5.2;
            }
        }
    }
}
at-0/3/0 {
    encapsulation atm-pvc;
    atm-options {
        vpi 0 maximum-vcs 1200;
    }
    unit 140 {
        encapsulation atm-snap;
        multipoint;
        family inet {
            address 192.168.7.4/24 {
                multipoint-destination 192.168.7.5;
                vci 0.100;
                inverse-arp;
            }
        }
    }
}
at-7/3/0 {
    encapsulation atm-pvc;
    atm-options {
        vpi 0 maximum-vcs 1200;
    }
    unit 0 {
        encapsulation atm-snap;
        vci 0.32;
        family inet {
            address 192.168.12.3/32 {
                destination 192.168.12.2;
            }
        }
    }
}

```

```

    }
  }
}

```

Example: Configuring ATM2 IQ Interfaces

Configure VP tunnel-shaping and OAM F4 on an ATM2 IQ interface:

```

interfaces {
  at-5/2/0 {
    atm-options {
      vpi 0 {
        shaping {
          vbr peak 10m sustained 6m burst 12;
        }
        oam-period 10;
        oam-liveness {
          up-count 6;
          down-count 5;
        }
      }
      vpi 4 {
        shaping {
          vbr peak 7m sustained 4m burst 24;
        }
      }
      vpi 5 {
        oam-period 10;
        oam-liveness {
          up-count 6;
          down-count 5;
        }
      }
      vpi 6;
    }
    unit 0 {
      vci 0.128;
      transmit-weight 20;
      family inet {
        address 192.168.9.225/32 {
          destination 192.168.9.224;
        }
      }
    }
    unit 1 {
      vci 0.129;
      transmit-weight 30;
      family inet {
        address 192.168.9.226/32 {
          destination 192.168.9.227;
        }
      }
    }
  }
}

```



```
unit 2 {  
  vci 5.123;  
  shaping {  
    vbr peak 60m sustained 4m burst 24;  
  }  
  family inet {  
    address 192.168.9.227/32 {  
      destination 192.168.9.230;  
    }  
  }  
}
```


Chapter 14

Configuring ATM-over-ADSL Interfaces

This chapter includes the following topics:

- ATM-over-ADSL Overview on page 355
- Configuring Physical ATM Interfaces and Logical Interface Properties for ADSL on page 356
- Configuring the ATM-over-ADSL Virtual Path Identifier on page 356
- Configuring the ATM-over-ADSL Physical Interface Operating Mode on page 357
- Configuring the ATM-over-ADSL Physical Interface Encapsulation Type on page 358
- Configuring the ATM-over-ADSL Logical Interface Encapsulation Type on page 358
- Configuring the ATM-over-ADSL Protocol Family on page 359
- Configuring the ATM-over-ADSL Virtual Channel Identifier on page 360

ATM-over-ADSL Overview

J4300 and J6300 Services Routers with asymmetrical DSL (ADSL) Annex A or Annex B PIMs can use an ATM interface to send network traffic through a point-to-point connection to a DSLAM. ATM-over-ADSL interfaces are not supported on J2300 Services Routers.



NOTE: You can configure J4300 and J6300 Services Routers with ADSL PIMs for connections through DSL only, not for direct ATM connections.

You configure the underlying ADSL as an ATM interface with an interface name of **at-pim/0/port**. Multiple encapsulation types are supported on both the physical and logical ATM-over-ADSL interface.

You can configure Point-to-Point Protocol over Ethernet (PPPoE) over ATM to connect through DSL lines. For PPPoE on an ATM-over-ADSL interface, you must configure encapsulation on both the physical and logical interfaces. To configure encapsulation on an ATM-over-ADSL physical interface, use Ethernet over ATM encapsulation. To configure encapsulation on an ATM-over-ADSL logical interface, use the PPPoE over AAL5 LLC encapsulation. LLC encapsulation allows a single ATM virtual connection to transport multiple protocols.



NOTE: PPPoE encapsulation is not supported on an M120 router with ATM2 PICs.

When you configure a point-to-point encapsulation such as PPP on a physical interface, the physical interface can have only one logical interface (only one **unit** statement) associated with it.

For more information about configuring PPPoE, see “Configuring PPPoE” on page 788.

Configuring Physical ATM Interfaces and Logical Interface Properties for ADSL

To configure physical ATM interfaces for ADSL, include the **vpi 0** statement at the [edit interfaces at-*pim*/0/*port* atm-options] hierarchy level, the **operating-mode** statement at the [edit interfaces at-*pim*/0/*port* dsl-options] hierarchy level, and the **encapsulation** statement at the [edit interfaces at-*pim*/0/*port*] hierarchy level:

```
[edit interfaces at-pim/0/port]
atm-options {
  vpi 0;
}
dsl-options {
  operating-mode mode;
}
encapsulation (atm-pvc | ethernet-over-atm);
```

Configure logical interface properties by including the **encapsulation** statement, **family** statement, and **vci** statement:

```
unit logical-unit-number {
  encapsulation (atm-vc-mux | atm-nlpd | atm-cisco-nlpd | atm-snap | atm-ppp-vc-mux |
    atm-ppp-llc | ether-over-atm-llc | ppp-over-ether-over-atm-llc);
  family inet {
    vci vpi-identifier.vci-identifier;
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Configuring the ATM-over-ADSL Virtual Path Identifier

Set the ATM virtual path identifier (VPI) to 0 (zero) by including the **vpi 0** statement at the [edit interfaces at-*pim*/0/*port* atm-options] hierarchy level:

```
[edit interfaces at-pim/0/port atm-options]
vpi 0;
```

Configuring the ATM-over-ADSL Physical Interface Operating Mode

Configure the ADSL operating mode on the physical ATM interface by including the `operating-mode` statement at the [edit interfaces *at-pim/0/port* dsl-options] hierarchy level:

```
[edit interfaces at-pim/0/port dsl-options]
operating-mode (adsl2plus | ansi-dmt | auto | etsi | itu-annexb-non-ur2 | itu-annexb-ur2 |
itu-dmt | itu-dmt-bis);
```

By default, the mode is `auto`, which means the ADSL line autonegotiates the setting to match the setting of the DSLAM located at the central office.

Table 31 on page 357 shows the Annex A PIM and Annex B PIM operational modes for ATM-over-ADSL interfaces.

Table 31: ATM-over-ADSL Operational Modes

Encapsulation Types	Comments
Annex A PIMs	
adsl2plus	Set the ADSL line to train in the ITU G.992.5 mode.
ansi-dmt	Set the ADSL line to train in the ANSI T1.413 Issue 2 mode.
auto	Set the ADSL line to autonegotiate the setting to match the setting of the DSLAM located at the central office. The ADSL line trains in the ANSI T1.413 Issue 2 (<code>ansi-dmt</code>) or ITU G.992.1 (<code>itu-dmt</code>) mode.
itu-dmt	Set the ADSL line to train in the ITU G.992.1 mode.
itu-dmt-bis	Set the ADSL line to train in the ITU G.992.3 mode.
itu-lite	Set the ADSL line to train in the G.992.2 mode.
itu-lite-bis	Set the ADSL line to train in the G.992.4 mode.
Annex B PIMs	
adsl2plus	Set the ADSL line to train in the ITU G.992.5 mode.
auto	Set the ADSL line after autonegotiating the setting to match the setting of the DSLAM located at the central office.
etsi	Set the ADSL line to train in the ETSI TS 101 388 V1.3.1 mode.
itu-dmt	Set the ADSL line to train in the ITU G.992.1 mode.
itu-dmt-bis	Set the ADSL line to train in the ITU G.992.3 mode.

Table 31: ATM-over-ADSL Operational Modes (*continued*)

Encapsulation Types	Comments
itu-annexb-ur2	Set the ADSL line to train in the ITU G.992.1 Deutsche Telekom UR-2 mode.
itu-annexb-non-ur2	Set the ADSL line to train in the ITU G.992.1 non-UR-2 mode.
itu-dmt	Set the ADSL line to train in the ITU G.992.1 mode.

Configuring the ATM-over-ADSL Physical Interface Encapsulation Type

Configure the physical interface encapsulation type by including the `encapsulation` statement at the `[edit interfaces at-pim/0/port]` hierarchy level:

```
[edit interfaces at-pim/0/port]
encapsulation type;
```

Table 32 on page 358 shows the physical interface encapsulation types for ATM-over-ADSL interfaces.

Configuring the ATM-over-ADSL Logical Interface Encapsulation Type

Configure the logical interface encapsulation type by including the `encapsulation` statement:

```
[edit interfaces at-pim/0/port unit logical-unit-number]
encapsulation type;
```

You can include this statement at the following hierarchy levels:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]
```

Table 32 on page 358 shows the logical interface encapsulation types for ATM-over-ADSL interfaces.

Table 32: ATM-over-ADSL Encapsulation Types

Encapsulation Types	Comments
Physical Interface	
ether-over-atm	Ethernet over ATM encapsulation. Use this type of encapsulation for interfaces that carry IPv4 traffic.
atm-pvc	ATM permanent virtual circuits (PVCs).

Table 32: ATM-over-ADSL Encapsulation Types (*continued*)

Encapsulation Types	Comments
Logical Interface	
atm-vc-mux	Use ATM VC multiplex encapsulation. You can only configure the <code>inet</code> family when you use this type of encapsulation.
atm-nlpd	Use ATM network layer protocol ID (NLPD) encapsulation. You can only configure the <code>inet</code> family when you use this type of encapsulation.
atm-cisco-nlpd	Use Cisco NLPD encapsulation. You can only configure the <code>inet</code> family when you use this type of encapsulation.
atm-snap	Use ATM subnetwork attachment point (SNAP) encapsulation.
atm-ppp-vc-mux	Use PPP over ATM AAL5 multiplex encapsulation.
atm-ppp-llc	Use ATM PPP over AAL5 logical link control (LLC) encapsulation.
ether-over-atm-llc	Use Ethernet over LLC encapsulation for interfaces that carry IPv4 traffic. You cannot configure multipoint interfaces if you use this type of encapsulation.
ppp-over-ether-over-atm-llc	Use PPP over Ethernet over ATM LLC encapsulation. You cannot configure the interface address when you use this encapsulation type. Instead, you configure the interface address on the PPP interface.

Configuring the ATM-over-ADSL Protocol Family

Configure the protocol family type by including the `family` statement:

```
[edit interfaces at-pim/0/port unit logical-unit-number]
family family;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Configuring the ATM-over-ADSL Virtual Channel Identifier

Configure the virtual channel identifier (VCI) type and value by including the `vci` statement:

```
[edit interfaces at-pim/0/port unit logical-unit-number]  
vci vpi-identifier.vci-identifier;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Chapter 15

Configuring ATM-over-SHDSL Interfaces

This chapter includes the following topics:

- ATM-over-SHDSL Overview on page 361
- Configuring ATM Mode for SHDSL Overview on page 362
- Configuring ATM Mode on the PIM on page 363
- Configuring SHDSL Operating Mode on an ATM Physical Interface on page 364
- Configuring Encapsulation on the ATM Physical Interface on page 364
- Configuring Logical Interface Properties on page 365
- Example: Configuring an ATM-over-SHDSL Interface on page 366
- Verifying an ATM-over-SHDSL Interface Configuration on page 367

ATM-over-SHDSL Overview

The symmetric high-speed digital subscriber line (SHDSL) Physical Interface Module (PIM) is available for J Series Services Routers. The PIM supports multi-rate, high-speed, symmetrical digital subscriber line technology for data transfer between a single customer premises equipment (CPE) subscriber and a central office (CO). Unlike ADSL, which was designed for delivering more bandwidth downstream than upstream, SHDSL is symmetrical and delivers a bandwidth of 2.3 Mbps in both directions. The SHDSL PIM has 2 ports and supports ATM-over-SHDSL mode only.

SHDSL is defined in the following specifications from the ITU and the Internet Engineering Task Force (IETF):

- ITU G.991.2, *Single-pair High-speed Digital Subscriber Line (SHDSL) Transceiver*
- ITU G.994.1, *Handshake Procedures for Digital Subscriber Line (DSL) Transceivers*
- ITU G.997.1, *Physical Layer Management for Digital Subscriber Line (DSL) Transceivers*
- RFC 3276, *Definitions of Managed Objects for High Bit-Rate DSL - 2nd generation (HDSL2) and Single-Pair High-Speed Digital Subscriber Line (SHDSL) Lines*

J Series routers with SHDSL Annex A or Annex B PIMs act as a primary WAN link. They use an ATM interface to send network traffic through a point-to-point connection to a DSL-access multiplexer (DSLAM). You can configure Point-to-Point Protocol over Ethernet (PPPoE) over ATM to connect through DSL lines. For more information about configuring PPPoE, see “Configuring PPPoE” on page 788.

ATM-over-SHDSL interfaces are not supported on J2300 Services Routers.



NOTE: You can configure J Series routers with SHDSL PIMs for connections through SHDSL only, not for direct ATM connections.

Configuring ATM Mode for SHDSL Overview

To configure the ATM mode for SHDSL, include the `pic-mode` statement at the [edit chassis `fpc fpc-number` pic 0 shdsl] hierarchy level:

```
[edit chassis]
fpc fpc-number {
  pic 0 {
    shdsl {
      pic-mode (1-port-atm | 2-port-atm);
    }
  }
}
```

For more information about configuring the ATM mode, see the *JUNOS System Basics Configuration Guide* and the *J-series Services Router Advanced WAN Access Configuration Guide*.

To configure SHDSL operating mode on the physical ATM interface and set the encapsulation, include the `shdsl-options` statement and the `encapsulation` statement at the [edit interfaces `at-pim/0/port`] hierarchy level:

```
[edit interfaces at-pim/0/port]
shdsl-options {
  annex (annex-a | annex-b);
  line-rate line-rate;
  loopback (local remote);
  snr-margin {
    current margin;
    snext margin;
  }
  encapsulation (atm-pvc | ethernet-over-atm)
}
```

To configure ATM virtual path identifier (VPI) options for the interface, include the `vpi` statement at the [edit interfaces `interface-name` atm-options] hierarchy level:

```
[edit interfaces interface-name]
atm-options {
  vpi vpi-identifier {
    maximum-vcs maximum-vcs;
    oam-liveness {
      up-count cells;
      down-count cells;
    }
    oam-period (disable | seconds);
  }
}
```

```
}
```

For more information about configuring ATM VPI options, see “Configuring the Maximum Number of ATM1 VCs on a VP” on page 300.

To configure logical interface properties, include the **encapsulation** statement, **family** statement, and **vci** statement:

```
unit logical-unit-number {
  encapsulation type;
  family inet {
    vci vpi-identifier.vci-identifier;
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Configuring ATM Mode on the PIM

The J Series routers with an SHDSL PIM installed support the 2-port, two-wire mode (Annex A or Annex B). You can configure only one mode on each 2-port SHDSL PIM.



NOTE: G.SHDSL interfaces on a J Series router only support 2-port, two-wire mode. This is enabled by default. The 1-port, 4-wire mode is not supported.

The two-wire mode supports autodetection of the line rate or fixed line rate and network speeds from 192 Kbps to 2.3 Kbps in 64-Kbps increments.

For information about configuring Annex A or Annex B, see “Configuring SHDSL Operating Mode on an ATM Physical Interface” on page 364.

To configure the ATM mode for SHDSL, include the **pic-mode** statement at the [edit chassis *fpc fpc-number* pic 0 shdsl] hierarchy level:

```
[edit chassis]
fpc fpc-number {
  pic 0 {
    shdsl {
      pic-mode (1-port-atm | 2-port-atm);
    }
  }
}
```

The default is 2-wire (two-port ATM) mode. To set the default explicitly, specify the **2-port-atm** option. For 4-wire (single-port ATM) mode, specify the **1-port-atm** option.

For more information about configuring the **pic-mode** statement, see the *JUNOS System Basics Configuration Guide*. For information about configuring the ATM mode, see the *J-series Services Router Basic LAN and WAN Access Configuration Guide*.

Configuring SHDSL Operating Mode on an ATM Physical Interface

To configure the SHDSL operating mode on the physical ATM interface, include the **shdsl-options** statement at the **[edit interfaces at-pim/0/port]** hierarchy level:

```
[edit interfaces at-pim/0/port]
shdsl-options {
  annex (annex-a | annex-b);
  line-rate line-rate;
  loopback (local | remote);
  snr-margin {
    current margin;
    snext margin;
  }
}
```

Configure the following SHDSL options:

- **annex**—The type of annex:
 - **annex-a**—Use for North American SHDSL network implementations.
 - **annex-b**—Use for European SHDSL network implementations.
- **line-rate**—The SHDSL line rate. The default for 2-wire mode is auto. The default for 4-wire mode is 4608 Kbps.
- **loopback**—A loopback connection, **local** or **remote**.
 - **local**—Use to troubleshoot physical PIC errors. A local loopback loops packets, including both data and timing information, back on the local router's PIM.
 - **remote**—Use to troubleshoot physical circuit problems between the local router and the remote router. A remote loopback loops packets, including both data and timing information, back on the remote router's PIC.
- **snr-margin**—The SHDSL signal-to-noise ratio (SNR) margin, **current** or **snext**. The SNR margin is the difference between the desired SNR and the actual SNR.
 - **current**—Current SNR is the difference between desired SNR and the actual SNR. When configured, the line trains at higher than current noise margin plus SNR threshold.
 - **snext**—Self-near-end crosstalk (SNEXT) SNR margin line trains the line at higher than SNEXT threshold.

Configuring Encapsulation on the ATM Physical Interface

To configure the type of encapsulation for the physical ATM interface, include the **encapsulation** statement at the **[edit interfaces at-pim /0/port]** hierarchy level:

```
[edit interfaces at-pim/0/port]
```

```
encapsulation (atm-pvc | ether-over-atm);
```

Configure one of the following:

- **atm-pvc**—ATM permanent virtual circuits (PVCs), used for PPP over ATM over SHDSL interfaces. This is the default encapsulation.
- **ether-over-atm**—Ethernet over ATM encapsulation. For interfaces that carry IPv4 traffic, use this type of encapsulation.

Configuring Logical Interface Properties

To configure logical interface properties, include the **encapsulation** statement, **family** statement, and **vci** statement:

```
unit logical-unit-number {
  encapsulation type;
  family inet {
    vci vpi-identifier.vci-identifier;
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

To configure the logical link-layer encapsulation type, include the **encapsulation** statement.

ATM-over-SHDSL interfaces that use **inet** (IP) protocols support the following encapsulations on the logical interface:

- **atm-vc-mux**—Use ATM VC multiplex encapsulation. You can only configure the **inet** family when you use this type of encapsulation.
- **atm-nlpd**—Use ATM network layer protocol ID (NLPD) encapsulation. You can only configure the **inet** family when you use this type of encapsulation.
- **atm-cisco-nlpd**—Use Cisco NLPD encapsulation. You can only configure the **inet** family when you use this type of encapsulation.

ATM-over-SHDSL for PPP over ATM interfaces support the following encapsulations on the logical interface:

- **atm-ppp-llc**—Use ATM PPP over AAL5 logical link control (LLC) encapsulation.
- **atm-ppp-vc-mux**—Use PPP over ATM AAL5 multiplex encapsulation.

ATM-over-SHDSL interfaces also support the following encapsulations on the logical interface:

- **atm-snap**—Use ATM subnetwork attachment point (SNAP) encapsulation.
- **atm-mlppp-llc**—For ATM2 IQ interfaces only, use Multilink PPP (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.
- **ppp-over-ether-over-atm-llc**—Use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, you configure the interface address on the PPP interface.
- **family**—The family protocol type.
- **vci**—The virtual channel identifier (VCI) type and value.
- **vci-identifier**—ATM virtual circuit identifier. Unless you configure the interface to use promiscuous mode, this value cannot exceed the largest numbered VC configured for the interface with the **maximum-vcs** option of the **vpi** statement. Specify a VCI identifier from 0 through 4089 or 0 through 65,535 with promiscuous mode. VCIs from 0 through 31 are reserved.
- **vpi-identifier**—ATM virtual path identifier. Specify a VPI from 0 through 255. The default is 0.

Example: Configuring an ATM-over-SHDSL Interface

The following example illustrates an ATM-over-SHDSL interface configuration.

Configuration for the ATM Mode on the PIM	<pre>[edit chassis] fpc 6 { pic 0 { shdsl { pic-mode 2-port-atm; } } }</pre>
Configuration for the SHDSL Operating Mode on the Physical ATM Interface	<pre>[edit interfaces at-6/0/0/0] shdsl-options { annex annex-b; line-rate 192; loopback local; snr-margin { current 1; snext 2; } }</pre>
Configuration for the Encapsulation on the Physical ATM Interface	<pre>[edit interfaces at-6/0/0/0] encapsulation ethernet-over-atm;</pre>
Configuration for the Logical Interface	<pre>[edit interfaces at-6/0/0/0 unit 3] encapsulation atm-nlpid; family inet {</pre>

```
vci 25;  
}
```

Verifying an ATM-over-SHDSL Interface Configuration

To verify an ATM-over-SHDSL interface configuration, you can issue the following operational mode command:

```
user@host> show interfaces at-pim/0/port extensive
```


Part 6

Configuring Frame Relay

- Configuring Frame Relay on page 371

Chapter 16

Configuring Frame Relay

This chapter discusses configuration of the following Frame Relay properties:

- Frame Relay Overview on page 371
- Configuring Frame Relay Interface Encapsulation on page 372
- Configuring Frame Relay Control Bit Translation on page 375
- Configuring the Media MTU on Frame Relay Interfaces on page 377
- Setting the Protocol MTU with Frame Relay Encapsulation on page 377
- Configuring Frame Relay Keepalives on page 378
- Configuring Inverse Frame Relay ARP on page 379
- Configuring the Router as a DCE with Frame Relay Encapsulation on page 380
- Configuring Frame Relay DLCIs on page 380

Frame Relay Overview

The Frame Relay protocol allows network designers to reduce costs by using shared facilities that are managed by a Frame Relay service provider. Users pay fixed charges for the local connections from each site in the Frame Relay network to the first point of presence (POP) in which the provider maintains a Frame Relay switch. The portion of the network between the endpoint switches is shared by all the customers of the service provider, and individual data-link connection identifiers (DLCIs) are assigned to ensure each customer receives only their own traffic.

Users contract with their providers for a specific minimum portion of the shared bandwidth Committed Information Rate (CIR) and for a maximum allowable peak rate, Burst Information Rate (BIR). Depending on the terms of the contract, traffic exceeding the CIR can be marked as eligible for discard, in the event of network congestion, or a best effort term can apply up to the BIR rate.

Frame Relay does not require private and permanently connected wide area network facilities, unlike some older WAN protocols.

Frame Relay was developed as a replacement for the older and much slower X.25 protocol. It scales to much higher data rates because it does not require explicit acknowledgment of each frame of data.

You can configure the Frame Relay protocol on SONET/SDH, E1/E3, and T1/T3 physical router interfaces, and on the channelized DS3, channelized OC12,

channelized T3 intelligent queuing (IQ), channelized OC12 IQ, and channelized E1 IQ interfaces.

Configuring Frame Relay Interface Encapsulation

Point-to-Point Protocol (PPP) encapsulation is the default encapsulation type for physical interfaces. You need not configure encapsulation for any physical interfaces that support PPP encapsulation. If you do not configure encapsulation, PPP is used by default. For physical interfaces that do not support PPP encapsulation, you must configure an encapsulation to use for packets transmitted on the interface. You can optionally configure an encapsulation on a logical interface, which is the encapsulation used within certain packet types.

For more information, see the following sections:

- Configuring the Frame Relay Encapsulation on a Physical Interface on page 372
- Configuring the Frame Relay Encapsulation on a Logical Interface on page 375

Configuring the Frame Relay Encapsulation on a Physical Interface

For Frame Relay interfaces, configure Frame Relay encapsulation on the physical interface. This encapsulation is defined in RFC 1490, *Multiprotocol Interconnect over Frame Relay*. SONET/SDH and T3 interfaces can use Frame Relay encapsulation.

To configure Frame Relay encapsulation on a physical interface, include the encapsulation statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]  
encapsulation type;
```

When you configure a multipoint encapsulation (such as Frame Relay), the physical interface can have multiple logical units, and the units can be either point-to-point or multipoint.

The encapsulation type can be one of the following:

- Flexible Frame Relay (**flexible-frame-relay**)—IQ interfaces can use flexible Frame Relay encapsulation. You use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.
- Frame Relay (**frame-relay**)—Defined in RFC 1490. E1, E3, link services, SONET/SDH, T1, T3, and voice services interfaces can use Frame Relay encapsulation. Five related versions are supported:

- Circuit cross-connect (CCC) version (**frame-relay-ccc**)—The same as standard Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to CCC. The logical interface must also have **frame-relay-ccc** encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.
- Translational cross-connect (TCC) version (**frame-relay-tcc**)—Similar to Frame Relay CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- Extended CCC version (**extended-frame-relay-ccc**)—This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC. The logical interface must have **frame-relay-ccc** encapsulation. When you use this encapsulation type, you can configure the **ccc** family only.
- Extended TCC version (**extended-frame-relay-tcc**)—Similar to extended Frame Relay CCC, this encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC, which is used for circuits with different media on either side of the connection.
- Port CCC version (**frame-relay-port-ccc**)—Defined in the Internet Engineering Task Force (IETF) document, *Frame Relay Encapsulation over Pseudo-Wires* (expired December 2002). This encapsulation type allows you to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. The connection between the two CE routers can be either user-to-network interface (UNI) or network-to-network interface (NNI); this is completely transparent to the PE routers. The logical interface does not require an encapsulation statement. When you use this encapsulation type, you can configure the **ccc** family only.
- Frame Relay Ether Type (**frame-relay-ether-type**)—Physical interfaces can use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. IETF frame relay encapsulation identifies the payload format using NLPID and SNAP formats. Cisco-compatible Frame Relay encapsulation uses the Ethernet type to identify the type of payload. Two related versions are supported:
 - TCC version (**frame-relay-ether-type-tcc**)—Cisco-compatible Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to TCC. This encapsulation is used for circuits with different media on either side of the connection.
 - Extended TCC version (**extended-frame-relay-ether-type-tcc**)—This encapsulation allows you to dedicate Cisco-compatible Frame Relay TCC for DLCIs 1 through 1022. This encapsulation is used for circuits with different media on either side of the connection. Extended Frame Relay ether type TCC encapsulation is supported on the same PICs as extended Frame Relay TCC encapsulation.



NOTE: When the encapsulation type is set to Cisco-compatible Frame Relay encapsulation, ensure that the LMI type is set to ANSI or Q933-A.

Support for extended Frame Relay and flexible Frame Relay differs by PIC type, as shown in Table 33 on page 374.

Table 33: PIC Support for Enhanced Frame Relay Encapsulation Types

PIC Type	Extended Frame Relay CCC	Extended Frame Relay TCC	Flexible Frame Relay
Intelligent Queuing			
1-port Channelized CHOC12 IQ	Yes	Yes	Yes
4-port Channelized DS3 IQ	Yes	Yes	Yes
10-port Channelized E1 IQ	Yes	Yes	Yes
4-port E3 IQ	Yes	Yes	Yes
1-port Channelized STM1 IQ	Yes	Yes	Yes
SONET/SDH			
1-port OC12	Yes	Yes	No
2-port OC3	Yes	Yes	No
1-port OC48	Yes	Yes	No
1-port OC192	Yes	Yes	No
1-port STM16 SDH, SMSR	Yes	Yes	No
Others			
4-port E1	No	No	No
4-port T1	No	No	No
4-port T3	No	No	No
10-port Channelized E1	No	No	No
2-port Channelized DS3	No	No	No
1-port Channelized OC12, SMIR	No	No	No
4-port Channelized DS3	No	No	No
1-port Channelized STM1, SMIR	No	No	No
2-port Serial	No	No	No

Example: Configuring the Encapsulation on a Physical Interface

Configure Frame Relay encapsulation on a SONET/SDH interface. The second and third family statements allow Intermediate System-to-Intermediate System (IS-IS) and Multiprotocol Label Switching (MPLS) to run on the interface.

```
[edit interfaces]
so-7/0/0 {
```

```

encapsulation frame-relay;
unit 0 {
    point-to-point;
    family inet {
        address 192.168.1.113/32 {
            destination 192.168.1.114;
        }
    }
    family iso;
    family mpls;
}
}

```

Configuring the Frame Relay Encapsulation on a Logical Interface

Generally, you configure an interface's encapsulation at the [edit interfaces *interface-name*] hierarchy level. However, for Frame Relay encapsulation, you can also configure the encapsulation type that is used inside the Frame Relay packet itself. To do this, include the `encapsulation` statement, specifying the `frame-relay-ccc`, `frame-relay-ppp`, `frame-relay-tcc`, `frame-relay-ether-type`, or `frame-relay-ether-type-tcc` option:

```

encapsulation (frame-relay-ccc | frame-relay-ppp | frame-relay-tcc | frame-relay-ether-type
| frame-relay-ether-type-tcc);

```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Configuring Frame Relay Control Bit Translation

On interfaces with Frame Relay CCC encapsulation, you can configure Frame Relay control bit translation, as defined in the IETF documents:

- Internet draft draft-martini-frame-encap-mpls-00.txt, *Frame Relay Encapsulation over Pseudo-Wires* (expired December 2002)
- Internet draft draft-martini-l2circuit-encap-mpls-07.txt, *Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks* (expired December 2004)

To support Frame Relay services over IP and MPLS backbones using Layer 2 VPNs and Layer 2 circuits, you can configure translation of the Frame Relay control bits. When you configure translation of Frame Relay control bits, the bits are mapped into the Layer 2 circuit control word and preserved across the IP or MPLS backbone.

The JUNOS Software allows you to translate the following Frame Relay control bits:

- Discard eligibility (DE)—A header bit used to identify lower priority traffic that can be dropped during periods of congestion.

- Forward explicit congestion notification (FECN)—A header bit transmitted by the source router requesting that the destination router slow down its requests for data.
- Backward explicit congestion notification (BECN)—A header bit transmitted by the destination router requesting that the source router send data more slowly.

By default, translation of Frame Relay control bits is disabled. If you enable Frame Relay control bit translation, the bits are translated in both directions (CE to PE and PE to CE):

- From CE to PE—At ingress, the DE, FECN, and BECN header bits from the incoming Frame Relay header are mapped to the control word.
- From PE to CE—At egress, the DE, FECN, and BECN header bits from the control word are mapped to the outgoing Frame Relay header.

The Frame Relay control bits do not map to MPLS EXP labels, and do not affect class-of-service (CoS) behavior inside the provider network.

You enable or explicitly disable translation of Frame Relay control bits by including the `translate-discard-eligible` and `translate-fecn-and-becn` statements:

```
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family *ccc*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family *ccc*]

If you enable or disable Frame Relay control bit translation on one CE-facing interface, you must configure the same Frame Relay control bit translation settings on the other CE-facing interface.

If you change the Frame Relay control bit translation settings, the circuit goes down and comes back up, which might result in traffic loss for a few seconds.

If you enable Frame Relay control bit translation, the number of supportable Layer 2 virtual private networks (VPNs) and Layer 2 circuits is reduced to one eighth of what the router can support without Frame Relay control bit translation enabled.

For ATM2 IQ interfaces, the control word contains a field to carry ATM cell loss priority (CLP) information by default. For more information, see “Configuring Layer 2 Circuit Transport Mode” on page 300.

For more information about Layer 2 circuits, see the *JUNOS VPNs Configuration Guide* and the *JUNOS Routing Protocols Configuration Guide*. For a comprehensive example, see the *JUNOS Feature Guide*.

Configuring the Media MTU on Frame Relay Interfaces

For Frame Relay interfaces, the default media maximum transmission unit (MTU) is 4482 bytes. (For a complete list of MTU values, see “Configuring the Media MTU” on page 98.)

To modify the default media MTU size for a physical interface, include the `mtu` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
mtu bytes;
```

If you change the size of the media MTU, you must ensure that the size is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. You can include the `mtu` statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number family family]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family]`

For more information, see “Setting the Protocol MTU with Frame Relay Encapsulation” on page 377.

Setting the Protocol MTU with Frame Relay Encapsulation

For each interface, you can configure an interface-specific MTU by including the `mtu` statement at the `[edit interfaces interface-name]` hierarchy level. If you need to modify this MTU for a particular protocol family, include the `mtu` statement:

```
mtu mtu;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number family family]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family]`

For Frame Relay encapsulation, the default protocol MTU is 4470 bytes.

If you increase the size of the protocol MTU, you must ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. (You configure the media MTU by including the `mtu` statement at the `[edit interfaces interface-name]` hierarchy level, as discussed in “Configuring the Media MTU on Frame Relay Interfaces” on page 377.)

When the family is `mpls`, the default protocol MTU is 1488 bytes. MPLS packets are 1500 bytes and have 4 to 12 bytes of overhead.

Configuring Frame Relay Keepalives

By default, physical interfaces configured with Cisco High-level Data Link Control (HDLC) or Point-to-Point Protocol (PPP) encapsulation send keepalive packets at 10-second intervals. The Frame Relay term for keepalives is Local Management Interface (LMI) packets; note that the JUNOS Software supports both ANSI T1.617 Annex D LMIs and ITU Q933 Annex A LMIs.

To disable the sending of keepalives on a physical interface, include the `no-keepalives` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]
no-keepalives;
```

For back-to-back Frame Relay connections, either disable the sending of keepalives on both sides of the connection, or configure one side of the connection as a data terminal equipment (DTE) (the default JUNOS configuration) and the other as a data circuit-terminating equipment (DCE).

If keepalives are enabled, the number of possible DLCI configurations on a multipoint or multicast connection is limited by the MTU size selected for the interface. To calculate the available DLCIs, use the formula $(MTU - 12) / 5$. To increase the number of possible DLCIs, disable keepalives.

Configuring Tunable Keepalives for Frame Relay LMI

On interfaces configured with Frame Relay connections, you can tune the keepalive settings by using the `lmi` statement. A Frame Relay interface can be either DCE or DTE (the default JUNOS configuration). DTE acts as a master, requesting status from the DCE part of the link.

By default, the JUNOS Software uses ANSI T1.617 Annex D LMIs. To change to ITU Q933 Annex A LMIs, include the `lmi-type itu` statement at the `[edit interfaces interface-name lmi]` hierarchy level:

```
[edit interfaces interface-name lmi]
lmi-type itu;
```

To configure Frame Relay keepalive parameters, include the `lmi` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]
lmi (Frame Relay) {
  lmi-type (ansi | itu);
  n391dte number;
  n392dce number;
  n392dte number;
  n393dce number;
  n393dte number;
  t391dte seconds;
  t392dce seconds;
}
```

You can include the following statements:

- **n391dte**—DTE full status polling interval. The DTE sends a status inquiry to the DCE at the interval specified by **t391dte**. **n391dte** specifies the frequency at which these inquiries expect a full status report; for example, a **n391dte** value of 10 would specify a full status report in response to every tenth inquiry. The intermediate inquiries ask for a keepalive exchange only. The range is from 1 through 255, with a default value of 6.
- **n392dce**—DCE error threshold. The number of errors required to bring down the link, within the event-count specified by **n393dce**. The range is from 1 through 10, with a default value of 3.
- **n392dte**—DTE error threshold. The number of errors required to bring down the link, within the event-count specified by **n393dte**. The range is from 1 through 10, with a default value of 3.
- **n393dce**—DCE monitored event-count. The range is from 1 through 10, with a default value of 4.
- **n393dte**—DTE monitored event-count. The range is from 1 through 10, with a default value of 4.
- **t391dte**—DTE keepalive timer. Period at which the DTE sends out a keepalive response request to the DCE and updates status depending on the DTE error threshold value. The range is from 5 through 30 seconds, with a default value of 10 seconds.
- **t392dce**—DCE keepalive timer. Period at which the DCE checks for keepalive responses from the DTE and updates status depending on the DCE error threshold value. The range is from 5 through 30 seconds, with a default value of 15 seconds.

Configuring Inverse Frame Relay ARP

Frame Relay interfaces support inverse Frame Relay ARP, as described in RFC 2390, *Inverse Address Resolution Protocol*. When inverse Frame Relay ARP is enabled, the router responds to received inverse Frame Relay ARP requests by providing IP address information to the requesting router on the other end of the Frame permanent virtual circuit (PVC).

The router does not initiate inverse Frame Relay ARP requests.

By default, inverse Frame Relay ARP is disabled. To configure a router to respond to inverse Frame Relay ARP requests, include the **inverse-arp** statement:

```
inverse-arp;
```

For a list of hierarchy levels at which you can include this statement, see **inverse-arp**.

You must configure Frame Relay encapsulation on the logical interface to support inverse ARP. For more information, see “Configuring Frame Relay Interface Encapsulation” on page 372.

Configuring the Router as a DCE with Frame Relay Encapsulation

By default, when you configure an interface with Frame Relay encapsulation, the routing platform is assumed to be DTE. That is, the routing platform is assumed to be at a terminal point on the network. To configure the routing platform to be DCE, include the `dce` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
dce;
```

When you configure the router to be a DCE, keepalives are disabled by default.

For back-to-back Frame Relay connections, either disable the sending of keepalives on both sides of the connection, or configure one side of the connection as a DCE and the other as a DTE by removing the `dce` statement from the configuration (the default JUNOS configuration).

Configuring Frame Relay DLCIs

When you are using Frame Relay encapsulation on an interface, each logical interface corresponds to one or more permanent virtual circuits (PVCs) or switched virtual circuits (SVCs). For each PVC or SVC, you must configure one data-link connection identifier (DLCI).

A Frame Relay interface can be a point-to-point interface or a point-to-multipoint (also called a multipoint nonbroadcast multiaccess [NBMA]) connection.

To configure Frame Relay DLCIs, you can do the following:

- Configuring a Point-to-Point Frame Relay Connection on page 380
- Configuring a Point-to-Multipoint Frame Relay Connection on page 381
- Configuring a Multicast-Capable Frame Relay Connection on page 381

Configuring a Point-to-Point Frame Relay Connection

To configure a point-to-point Frame Relay connection, include the `dlci` statement:

```
dlci dlci-identifier;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number]`

The DLCI identifier is a value from 16 through 1022. Numbers 1 through 15 are reserved for future use. A point-to-point interface can have one DLCI.



NOTE: For information about Frame Relay DLCI limitations for channelized interfaces, see “Data-Link Connection Identifiers on Channelized Interfaces” on page 388.

You configure the router to use DLCI sparse mode by including the `sparse-dlcis` statement at the `[edit chassis fpc slot-number pic pic-number]` hierarchy level. For more information about DLCI sparse mode, see the *JUNOS System Basics Configuration Guide*.

For more information about Frame Relay DLCIs, see “Configuring a Point-to-Point Frame Relay Connection” on page 380.

When you are configuring point-to-point connections, the MTU sizes on both sides of the connection must be the same.

Configuring a Point-to-Multipoint Frame Relay Connection

To configure a point-to-multipoint Frame Relay connection (also called a multipoint NBMA connection), include the `multipoint-destination` statement:

```
multipoint-destination address dlci dlci-identifier;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number family family address address]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family address address]`

For each destination, include one `multipoint-destination` statement. `address` is the address of the remote side of the connection, and `dlci-identifier` is the DLCI identifier for the connection.

When you are configuring point-to-multipoint connections, all interfaces in the subnet must use the same MTU size.

If keepalives are enabled, causing the interface to send LMI messages during idle times, the number of possible DLCI configurations is limited by the MTU selected for the interface. For more information, see “Configuring Frame Relay Keepalives” on page 378.

Configuring a Multicast-Capable Frame Relay Connection

By default, Frame Relay connections assume unicast traffic. If your Frame Relay switch performs multicast replication, you can configure the connection to support multicast traffic by including the `multicast-dlci` statement:

```
multicast-dlci dlci-identifier;
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number]`

- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The DLCI identifier is a value from 16 through 1022 that defines the Frame Relay DLCI over which the switch expects to receive multicast packets for replication.

You can configure multicast support only on point-to-multipoint Frame Relay connections.

If keepalives are enabled, causing the interface to send LMI messages during idle times, the number of possible DLCI configurations is limited by the MTU selected for the interface. For more information, see “Configuring Frame Relay Keepalives” on page 378.

Part 7

Configuring Channelized Interfaces

- Channelized Interfaces on page 385
- Configuring Channelized OC48/STM16 IQE Interfaces on page 405
- Configuring Channelized OC12/STM4 Interfaces on page 423
- Configuring Channelized OC3 IQ and IQE Interfaces on page 455
- Configuring Channelized STM1 Interfaces on page 465
- Configuring Channelized T3 Interfaces on page 479
- Configuring Channelized T1 Interfaces on page 495
- Configuring Channelized E1 Interfaces on page 501
- Configuring Channelized E1 PRI and T1 PRI Interfaces on page 509

Chapter 17

Channelized Interfaces

This chapter provides a high-level overview of channelized interfaces, focusing mainly on the capabilities, properties, and structure of channelized IQ and IQE interfaces:

- Channelized Interfaces Overview on page 385
- Channelized Interface Capabilities on page 386
- Data-Link Connection Identifiers on Channelized Interfaces on page 388
- Clock Sources on Channelized Interfaces on page 390
- Channelized E1 and T1 PIM Properties on page 393
- Channelized IQ and IQE Interfaces Properties on page 393
- Structure of Channelized IQ and Channelized IQE PICs on page 396

Channelized Interfaces Overview

Channelized interfaces enable you to configure a number of individual channels that subdivide the bandwidth of a larger interface and minimize the number of Physical Interface Cards (PICs) that an installation requires.



NOTE: Channelized intelligent queuing (IQ) and channelized enhanced intelligent queuing (IQE) interfaces require M Series Enhanced Flexible PIC Concentrators (FPCs) and MX Series Enhanced Flexible PIC Concentrators (FPCs).

Wherever JUNOS configuration guides refer to channelized interfaces and PICs without the “intelligent queuing IQ or IQE” descriptor, they are referring to the original channelized interfaces and PICs.

On M40e routers, all supported interface types support a maximum number of 784 traffic-bearing interfaces that can be created per interface port and includes ports on channelized PICs.

MX Series routers support two Type 2 Channelized IQ PICs: OC12/STM4 IQE PIC with SFP and OC48/STM16 IQE PIC with SFP. Each channelized OC12/STM4 PIC supports 4 ports, and the channelized OC48/STM16 PIC supports one port.

Channelized 4xCOC12 IQE PICs support deep-channelization of up to six OC slices (STS1) per port. For example, only six OC slices can be channelized to CT1/T1 or CE1/E1.

Channelized COC48 IQE PICs support deep-channelization of up to six OC slices (STS1) in a block of 12 contiguous OC slices. For example, only six OC slices out of OC slice 1-12 can be channelized to CT1/T1 or CE1/E1. The PIC supports deep-channelization of maximum 24 OC slices in this way.

Channelized OC48 IQE PICs do not support STS-48 clear-channel mode.

IQ and IQE PICs do not support aggregated SONET (link bonding).

For channelized IQ and IQE logical interfaces, you can configure class of service (CoS). For more information, see the *JUNOS Class of Service Configuration Guide*.

Channelized Interface Capabilities

You can configure each port of a channelized IQ PIC or channelized IQE PIC as a single interface that uses the entire available bandwidth, or partition each port into smaller data channels. In either case, you start with a channelized interface (designated by a **c** in the interface name, as in **coc12**). From the channelized interfaces, you configure data channels. Following are the channelized interface names and data-channel interface names associated with channelized IQ and IQE PICs.

Channelized Interface Names

This section lists the channelized interface names.

- **coc48-fpc/pic/port**—Channelized OC48 IQE interface. Configure on a Channelized OC48 IQE PIC.
- **coc12-fpc/pic/port**—Channelized OC12 interface. Configure on Channelized OC12 IQ or IQE PICs.

- **coc3-fpc/pic/port:channel**—Channelized OC3 interface. Configure on Channelized OC3 IQ or IQE, Channelized OC12 IQ or IQE PICs.
- **coc1-fpc/pic/port:channel**—Channelized OC1 interface. Configure on Channelized OC3 IQ or IQE, Channelized OC12 IQ or IQE, or Channelized OC48 IQE PICs.
- **ct3-fpc/pic/port:channel**—Channelized T3 interface. Configure on Channelized OC3 IQ or IQE, Channelized OC12 IQ or IQE, Channelized OC48 IQE, or Channelized DS3 IQ or IQE PICs.
- **cstm16-fpc/pic/port**—Channelized STM16 interface. Configure on a Channelized OC48 IQE PIC in SDH mode.
- **cstm4-fpc/pic/port**—Channelized STM4 interface. Configure on a Channelized OC12 IQ or IQE PIC in SDH mode.
- **cstm1-fpc/pic/port**—Channelized STM1 interface. Configure on a Channelized STM1 IQ or IQE PIC.
- **cau4-fpc/pic/port:channel**—Channelized AU-4 IQ interface. Configure on Channelized STM1 IQ or IQE, Channelized OC48 IQE, or Channelized OC12 IQE PICs.
- **ct1-fpc/pic/port:channel**—Channelized T1 interface. Configure on Channelized OC3 IQ or IQE, Channelized OC12 IQ or IQE, Channelized T1 IQ or IQE, Channelized OC48 IQE, or Channelized DS3 IQ or IQE PICs.
- **ce1-fpc/pic/port:channel**—Channelized E1 interface. Configure on Channelized E1 IQ or IQE, Channelized STM1 IQ or IQE, Channelized OC48/STM16 IQE, or Channelized OC12/STM4 IQE PICs.

Data-Channel Interface Names

This section lists the data-channel interface names.

- **e1-fpc/pic/port:channel**—E1 channel. Configure on Channelized E1 IQ or IQE, Channelized STM1 IQ or IQE, Channelized OC12/STM4 IQE, or Channelized OC48 IQE PICs.
- **e3-fpc/pic/port:channel**—E3 channel. Configure on Channelized OC3/STM1 IQE, or Channelized OC12/STM4 IQE, Channelized OC48 IQE, or Channelized/Clear channel DS3E3 IQE or E3 IQ PICs.
- **ds-fpc/pic/port:channel**—N×DS0 channel. Configure on Channelized OC3 IQ or IQE, Channelized OC12 IQ or IQE, Channelized OC48/STM16 IQE, Channelized STM1 IQ or IQE, Channelized DS3 IQ or IQE, Channelized T1 IQ, or Channelized E1 IQ or IQE PICs.
- **so-fpc/pic/port:channel**—SONET/SDH channel. Configure one OC3 channel on a Channelized OC3 IQ or IQE, four OC3 channels on a Channelized OC12 IQ or IQE, one OC12 channel on a Channelized OC12 IQ or IQE, four OC12 channels on Channelized OC48 IQE, or one STM1 channel on a Channelized STM1 IQ or IQE PICs.
- **t1-fpc/pic/port:channel**—T1 channel. Configure on Channelized T1 IQ or IQE, Channelized OC3 IQ or IQE, Channelized OC12 IQ or IQE, Channelized OC48 IQE, or Channelized DS3 IQ or IQE PICs.
- **t3-fpc/pic/port:channel**—T3 channel. Configure on Channelized OC3 IQ or IQE, Channelized OC12 IQ or IQE, Channelized OC48 IQE, Clear Channel DS3E3 IQE, or Channelized DS3 IQ or IQE PICs.

Data-Link Connection Identifiers on Channelized Interfaces

If you use Frame Relay encapsulation on a channelized interface, see Table 34 on page 388 for the maximum number of data-link connection identifiers (DLCIs) per channel that you can configure at each channel level for various channelized PICs.

If you use a per-unit-scheduler configuration on a channelized interface, see Table 35 on page 389 for the maximum number of data-link connection identifiers (DLCIs) per channel that you can configure at each channel level for various channelized PICs.



NOTE: The actual number of DLCIs you can configure for each channel is determined by the capabilities of your system, such as the number and types of PICs installed. If the number of DLCIs in the configuration exceeds the capabilities of your system, the router might not be able to support the maximum DLCI values shown in Table 34 on page 388. To determine the capabilities of your system, please contact Juniper Networks customer support.

Table 34: Frame Relay DLCI Limitations for Channelized Interfaces

PIC Types	Number of DLCIs per Level	Range
Original Channelized PICs		
DS0 level channels	3 for sparse mode	1–1022 for sparse mode (0 is reserved for the Local Management Interface [LMI])
T3 and T1 level channels	63 for regular mode	1–63 for regular mode
	3 for sparse mode	1–1022 for sparse mode (0 is reserved for the LMI)
Channelized IQ and IQE PICs		
DS0 level channels (Channelized DS3 IQ or IQE, Channelized STM1 IQ or IQE, Channelized E1 IQ or IQE, Channelized OC3 IQ or IQE, or Channelized OC12 IQ or IQE PICs)	16	1–1022 (0 is reserved for the LMI)
E1 level channels (Channelized E1 IQ or IQE PIC)	64	1–1022 (0 is reserved for the LMI)
E1 level channels (Channelized STM1 IQ or IQE PIC)	64	1–1022 (0 is reserved for the LMI)
OC3 level channels (Channelized OC3 IQ or IQE, or Channelized OC12 IQ or IQE PIC)	1022	1–1022 (0 is reserved for the LMI)
OC12 level channels (Channelized OC12 IQ or IQE, Channelized OC48/STM16 IQE PICs, and (per port on) OC12 ports on 4xOC12/STM4 IQE PICs)	1022	1–1022 (0 is reserved for the LMI)
STM1 level channel (Channelized STM1 IQ or IQE PIC)	1022	1–1022 (0 is reserved for the LMI)

Table 34: Frame Relay DLCI Limitations for Channelized Interfaces (continued)

PIC Types	Number of DLCIs per Level	Range
T1 level channels (Channelized DS3 IQ or IQE PIC)	64	1–1022 (0 is reserved for the LMI)
T1 level channels (Channelized OC3 IQ or IQE, or Channelized OC12 IQ or IQE PIC)	64	1–1022 (0 is reserved for the LMI)
T3 level channel (Channelized DS3 IQ or IQE, Channelized OC3 IQ or IQE, or Channelized OC12 IQ or IQE PIC)	1022	1–1022 (0 is reserved for the LMI)

Table 35: Per Unit Scheduler DLCI Limitations for Channelized Interfaces

PIC Types	Number of DLCIs per Level			
	Non M40e Platforms		M40e Platform Only	
	With Per-Unit-Scheduler	Without Per-Unit-Scheduler	With Per-Unit-Scheduler	Without Per-Unit-Scheduler
DS0 level channels	64	64	16	16
T1/E1 level channels	64	64	64	64
DS3/E3 level channels	975	† Protocol family combinations apply	256	256
SONET	975	† Protocol family combinations apply	975	† Protocol family combinations apply

† In these router, PIC, and scheduler configurations:
Combining multiple protocol families per PIC changes the number of Frame Relay DLCIs as shown in Table 36 on page 389.

Table 36: Protocol Family Combinations

Protocol Family Combinations	Number of DLCIs per PIC
inet	3600
inet6	3600
mpls	3000
inet, inet6	2400
inet, mpls	2000
inet6, mpls	2000
inet, inet6, mpls	1550

Clock Sources on Channelized Interfaces

Channelized interfaces and channelized IQ and IQE interfaces have different clocking capabilities. For channelized IQ and IQE interfaces, you can configure clocking on each interface independently by including the **clocking** (*internal* | *external*) statement at the [edit interfaces *interface-name*] hierarchy level.

For channelized IQ and IQE interfaces, clocking is provided as follows:

- For all channelized IQ and IQE PICs, the **clocking** statement is supported on all channels. To configure clocking on individual interfaces, include the **clocking** statement at the [edit interfaces *type-fpc/pic/port:channel*] hierarchy level. If you do not include the **clocking** statement, the individual interfaces use internal clocking by default.
- SONET/SDH-level clocking is provided at the root controller interface at the [edit interfaces *type-fpc/pic/port*] hierarchy level.
- Configure T3-level clocking by including the **clocking** statement at the [edit interfaces *ct3-fpc/pic/port*] hierarchy level.
- Configure T1-level clocking by including the **clocking** statement at the [edit interfaces *t1-fpc/pic/port:channel*] hierarchy level.
- Configure E1-level clocking by including the **clocking** statement at the [edit interfaces *ce1-fpc/pic/port*] hierarchy level.
- Configure clocking for all NxDS0 channels by including the **clocking** statement at the [edit interfaces *ct1-fpc/pic/port:channel*] or [edit interfaces *ce1-fpc/pic/port*] hierarchy level.
- The **clocking** statement is ignored if you include it at the [edit interfaces *coc1-fpc/pic/port:channel*] or [edit interfaces *cau4-fpc/pic/port:channel*] hierarchy level.
- SONET/SDH level clocking is applicable only at the controller interfaces for channelized IQ and IQE PICs. Clocking configuration is not effective at the *so-fpc/pic/port* or *so-fpc/pic/port:channel* for channelized IQ and IQE PICs.

For non-IQ and non-IQE channelized interfaces, clocking at each channel level is provided as follows:

- For Channelized OC12, DS3, and E1 PICs, the **clocking** statement is supported only for channel 0; it is ignored if included in the configuration of other channels. The clock source configured for channel 0 applies to all channels on these channelized interfaces.
- For the Channelized STM1 PIC, the **clocking** statement is supported on channels 0 through 62. To configure clocking on the STM1 interface, include the **loop-timing** statement at the [edit interfaces *e1-fpc/pic/port:0 sonet-options*] hierarchy level. To configure clocking on individual E1 interfaces, include the **clocking** statement at the [edit interfaces *e1-fpc/pic/port:channel*] hierarchy level. The channel number can be 0 through 62. If you do not include the **clocking** statement, the individual E1 interfaces use internal clocking by default.

- For channelized STM1 interfaces, you should configure the clock source at one side of the connection to be internal and configure the other side of the connection to be external.
- When you configure the clock source for a channelized interface—`t3-fpc/pic/port:0`, for example—you must also include the `channel-group` statement at the `[edit chassis]` hierarchy level, and specify channel group 0.

Table 37 on page 391 lists the clocking capabilities for each channelized PIC.

Table 37: Clocking Capabilities by Channelized PIC Type

PIC Type	SONET/SDH Level	DS3 Level	DS1/E1 Level
Channelized PICs			
Channelized DS3 and Multichannel DS3	Not applicable.	The <code>loop-timing</code> statement is supported at the <code>[edit interfaces t1-fpc/pic/port:0 t3-options]</code> or <code>[edit interfaces fpc/pic/port:0:0 t3-options]</code> hierarchy level.	The <code>clocking</code> statement is supported at the <code>[edit interfaces t1-fpc/pic/port:0]</code> or <code>[edit interfaces ds-fpc/pic/port:0:0]</code> hierarchy level.
Channelized E1	Not applicable.	Not applicable.	The <code>clocking</code> statement is supported at the <code>[edit interfaces e1-fpc/pic/port:0]</code> or <code>[edit interfaces ds-fpc/pic/port:0]</code> hierarchy level.
Channelized OC12	Not configurable.	The <code>clocking</code> statement is supported at the <code>[edit interfaces t3-fpc/pic/port:0]</code> hierarchy level.	Not applicable.
Channelized STM1	Not configurable.	Not applicable.	The <code>clocking</code> statement is supported at the <code>[edit interfaces e1-fpc/pic/port:[0-62]]</code> hierarchy level.
Channelized IQ and IQE PICs			
Channelized DS3 IQ or IQE	Not applicable.	<p>The <code>clocking</code> statement is supported at the <code>[edit interfaces ct3-fpc/pic/port]</code> hierarchy level.</p> <p>The <code>clocking</code> statement is ignored if you include it at the <code>[edit interfaces t3-fpc/pic/port]</code> hierarchy level.</p>	<p>For T1 channels, the <code>clocking</code> statement is supported at the <code>[edit interfaces t1-fpc/pic/port:[1-28]]</code> hierarchy level.</p> <p>For NxDS0 channels, the <code>clocking</code> statement is supported at the <code>[edit interfaces ct1-fpc/pic/port:[1-28]]</code> hierarchy level.</p>
Channelized E1 IQ	Not applicable.	Not applicable.	<p>For E1 and NxDS0 channels, the <code>clocking</code> statement is supported at the <code>[edit interfaces ce1-fpc/pic/port]</code> hierarchy level.</p> <p>The <code>clocking</code> statement is ignored if you include it at the <code>[edit interfaces e1-fpc/pic/port]</code> hierarchy level.</p>

Table 37: Clocking Capabilities by Channelized PIC Type (continued)

PIC Type	SONET/SDH Level	DS3 Level	DS1/E1 Level
Channelized OC3 IQ or IQE	<p>The clocking statement is supported at the [edit interfaces <i>coc3-fpc/pic/port</i>] hierarchy level.</p> <p>The clocking statement is ignored if you include it at the [edit interfaces <i>so-fpc/pic/port</i>] hierarchy level.</p>	<p>The clocking statement is supported at the [edit interfaces <i>t3-fpc/pic/port:[1-12]</i>] hierarchy level.</p> <p>The clocking statement is ignored if you include it at the [edit interfaces <i>coc1-fpc/pic/port:channel</i>] hierarchy level.</p>	<p>The clocking statement is supported at the [edit interfaces <i>ct1-fpc/pic/port:[1-12]:[1-28]</i>] and [edit interfaces <i>t1-fpc/pic/port:[1-12]:[1-28]</i>] hierarchy levels.</p>
Channelized OC12 IQ or IQE	<p>The clocking statement is supported at the [edit interfaces <i>coc12-fpc/pic/port</i>] hierarchy level.</p> <p>The clocking statement is ignored if you include it at the [edit interfaces <i>so-fpc/pic/port</i>] hierarchy level.</p>	<p>The clocking statement is supported at the [edit interfaces <i>t3-fpc/pic/port:[1-12]</i>] hierarchy level.</p> <p>The clocking statement is ignored if you include it at the [edit interfaces <i>coc1-fpc/pic/port:channel</i>] hierarchy level.</p>	<p>The clocking statement is supported at the [edit interfaces <i>ct1-fpc/pic/port:[1-12]:[1-28]</i>] and [edit interfaces <i>t1-fpc/pic/port:[1-12]:[1-28]</i>] hierarchy levels.</p>
Channelized OC48 IQE	<p>The clocking statement is supported at the [edit interfaces <i>coc48-fpc/pic/port</i>] hierarchy level.</p> <p>The clocking statement is ignored if you include it at the [edit interfaces <i>so-fpc/pic/port</i>] hierarchy level.</p>	<p>The clocking statement is supported at the [edit interfaces <i>t3-fpc/pic/port:[1-48]</i>] hierarchy level.</p> <p>The clocking statement is ignored if you include it at the [edit interfaces <i>coc1-fpc/pic/port:channel</i>] hierarchy level.</p>	<p>The clocking statement is supported at the [edit interfaces <i>ct1-fpc/pic/port:[1-48]:[1-28]</i>] and [edit interfaces <i>t1-fpc/pic/port:[1-48]:[1-28]</i>] hierarchy levels.</p>
Channelized STM1 IQ or IQE	<p>The clocking statement is supported at the [edit interfaces <i>cstm1-fpc/pic/port</i>] hierarchy level.</p> <p>The clocking statement is ignored if you include it at the [edit interfaces <i>cau4-fpc/pic/port:channel</i>] or [edit interfaces <i>so-fpc/pic/port</i>] hierarchy level.</p>	Not applicable.	<p>For E1 and NxDS0 channels, the clocking statement is supported at the [edit interfaces <i>ce1-fpc/pic/port[1-63]</i>] hierarchy level.</p> <p>The clocking statement is ignored if you include it at the [edit interfaces <i>e1-fpc/pic/port</i>] hierarchy level.</p>

Table 37: Clocking Capabilities by Channelized PIC Type (continued)

PIC Type	SONET/SDH Level	DS3 Level	DS1/E1 Level
Channelized STM4 IQ or IQE	<p>The clocking statement is supported at the <code>[edit interfaces cstm4-fpc/pic/port]</code> hierarchy level.</p> <p>The clocking statement is ignored if you include it at the <code>[edit interfaces cau4-fpc/pic/port:channel]</code> or <code>[edit interfaces so-fpc/pic/port]</code> hierarchy level.</p>	Not applicable.	<p>For E1 and NxDS0 channels, the clocking statement is supported at the <code>[edit interfaces ce1-fpc/pic/port[1-4]:[1-63]]</code> hierarchy level.</p> <p>The clocking statement is ignored if you include it at the <code>[edit interfaces e1-fpc/pic/port]</code> hierarchy level.</p>
Channelized STM16 IQE	<p>The clocking statement is supported at the <code>[edit interfaces cstm16-fpc/pic/port]</code> hierarchy level.</p> <p>The clocking statement is ignored if you include it at the <code>[edit interfaces cau4-fpc/pic/port:channel]</code> or <code>[edit interfaces so-fpc/pic/port]</code> hierarchy level.</p>	Not applicable.	<p>For E1 and NxDS0 channels, the clocking statement is supported at the <code>[edit interfaces ce1-fpc/pic/port[1-16]:[1-63]]</code> hierarchy level.</p> <p>The clocking statement is ignored if you include it at the <code>[edit interfaces e1-fpc/pic/port]</code> hierarchy level.</p>

Channelized E1 and T1 PIM Properties

Channelized E1 and T1 PIMs on J Series routers provide support for ISDN Primary Rate Interface (PRI) connectivity for dial-in and callback and for use as primary or backup network connections. You can configure up to 30 channelized E1 time slots (`ce1-pim/0/port`) or 23 channelized T1 time slots (`ct1-pim/0/port`) as an ISDN PRI group, with the 16th E1 time slot or the 24th T1 time slot operating as the D-channel to control the group of time slots as B-channels. These B-channels can operate unconfigured. The encapsulation type `multilink-ppp`, `cisco-hdlc`, or `ppp` is configured under the dialer interface.

For more information about configuring the dialer interface, see “Configuring ISDN Logical Interface Properties” on page 823.

E1 and T1 time slots unused by ISDN PRI can operate normally as DS0 interfaces. PRI B-channels run at 64 Kbps, but do not support the 56-Kbps line rate.

For more information about Channelized E1 PIMs, ISDN PRI connectivity, and the ISDN features they support, see the *J-series Services Router Basic LAN and WAN Access Configuration Guide*.

Channelized IQ and IQE Interfaces Properties

On channelized IQ and IQE interfaces, you can specify options that are globally applied to all interface types associated with channelized IQ and IQE interfaces. For

example, **e1-options** statements that you include at the `[edit interfaces ce1-fpc/pic/port]` hierarchy level apply globally to all E1 and NxDS0 interfaces that you create by partitioning `ce1-fpc/pic/port`. Likewise, **t3-options** statements that you include at the `[edit interfaces ct3-fpc/pic/port]` hierarchy level apply globally to all T1 and NxDS0 interfaces that you create by partitioning `ct3-fpc/pic/port`.

You can also apply interface options at the channel level. For example, you can include **t1-options** statements at the `[edit interfaces t1-fpc/pic/port <:channel>]` hierarchy level, and **ds0-options** statements at the `[edit interfaces ds0-0/1/1<:channel>]` hierarchy level.

Only a subset of the interface options is valid on each type of channelized IQ interface. You configure all HDLC information at the end-data channel level, not at the parent level. For example, configure HDLC information at the `[edit interfaces ds-fpc/pic/port<:channel>]` hierarchy level, not at the `[edit interfaces ct1-fpc/pic/port<:channel>]` hierarchy level.

Automatic Protection Switching (APS) is supported on channelized OC3, OC12, STM1, and STM4 IQ interfaces. To configure APS, include the **aps** statement with options at the `[edit interfaces interface-name sonet-options]` hierarchy level. For information about configuring APS, see “Configuring APS and MSP” on page 859.

In interchassis and intrachassis redundant LSQ configurations that use MLPPP and SONET APS, you can inhibit a router from sending PPP termination-request messages to the remote host if the link PIC fails. To inhibit the router from sending PPP termination-request messages to the remote host if the link PIC fails, include the **no-termination-request** statement at the `[edit interfaces interface-name ppp-options]` hierarchy level.

The **no-termination-request** statement is supported only with MLPPP and SONET APS configurations and works with PPP, PPP over Frame Relay, and MLPPP interfaces only. The supported PIC types are as follows:

- Channelized OC48/STM16 IQE PICs
- Channelized OC12/STM4 IQ and IQE PICs
- Channelized OC3 IQ and IQE PICs
- Channelized STM1 IQ and IQE PICs

Channelized IQ and IQE interfaces do not support receive buckets or transmit buckets.

For channelized IQ and IQE interfaces, there are some limitations on where you place certain statements in the configuration. When you configure clocking, bit error rate testing (BERT), C-bit parity, and loopback statements on T3, T1, or DS0 channels, you must follow these guidelines:

- For T3 IQ interfaces, you can include the **loopback payload** statement at the [edit interfaces *ct3-fpc/pic/port*] and [edit interfaces *t3-fpc/pic/port:channel*] hierarchy levels. For T1 interfaces, you can include the **loopback payload** statement at the [edit interfaces *t1-fpc/pic/port:channel*] hierarchy level; it is ignored if included at the [edit interfaces *ct1-fpc/pic/port*] hierarchy level. For NxDS0 interfaces, payload and remote loopback are the same. If you configure one, the other is ignored. NxDS0 IQ interfaces do not support local loopback.
- If you include clocking, BERT, and C-bit parity configurations at both the [edit interfaces *ct3-fpc/pic/port<:channel> t3-options*] and [edit interfaces *t3-fpc/pic/port<:channel> t3-options*] hierarchy levels, the channelized T3-level statements are valid, and the T3-level statements are ignored.
- If you include clocking, BERT, and C-bit parity configurations at both the [edit interfaces *ct3-fpc/pic/port<:channel> t3-options*] and [edit interfaces *t1-fpc/pic/port<:channel> t1-options*] hierarchy levels, the channelized T3-level statements are operational for the T3 connections and the T1-level statements are operational for the T1 connections.
- Because DS0 channels do not have clocking capability, you must configure clocking at the [edit interfaces *ct1-fpc/pic/port<:channel> t1-options*] or [edit interfaces *ce1-fpc/pic/port<:channel> e1-options*] hierarchy level for channelized NxDS0 IQ interfaces.
- You can set BERT at the [edit interfaces *t3-fpc/pic/port<:channel> t3-options*] hierarchy level or on any partitioned channel of the channelized T3 interface. There are 12 BERT patterns available for NxDS0 channels and 28 BERT patterns for T1, channelized T1, T3, and channelized T3 interfaces within channelized IQ interfaces.
- For channelized IQ and IQE PICs, SONET/SDH level, use the **sonet-options loopback** statement **local** and **remote** options at the controller interface (*coc48*, *cstm16*, *coc12*, *cstm4*, *coc3*, *cstm1*). It is ignored for path-level interfaces *so-fpc/pic/port* or *so-fpc/pic/port:channel*.
- For channelized interfaces that use Frame Relay encapsulation, the number of configurable DLCIs varies by channelized interface type.
- For channelized interfaces, you can configure class of service (CoS) on channels, but not at the controller level.
- For original Channelized OC12 PICs, limited CoS functionality is supported. For more information, contact Juniper Networks customer support.
- CoS is not configurable on controller interfaces.

Structure of Channelized IQ and Channelized IQE PICs

Figure 22 on page 396 through Figure 34 on page 400 show the structural organization of the channelized PICs, channelized IQ PICs, and channelized IQE PICs. Table 38 on page 401 through Table 40 on page 402 show the structure of channelized IQE PICs, channelized IQ PICs, and channelized PICs.

Figure 22: Channelized OC48/STM16 IQE PIC (in SONET Mode)

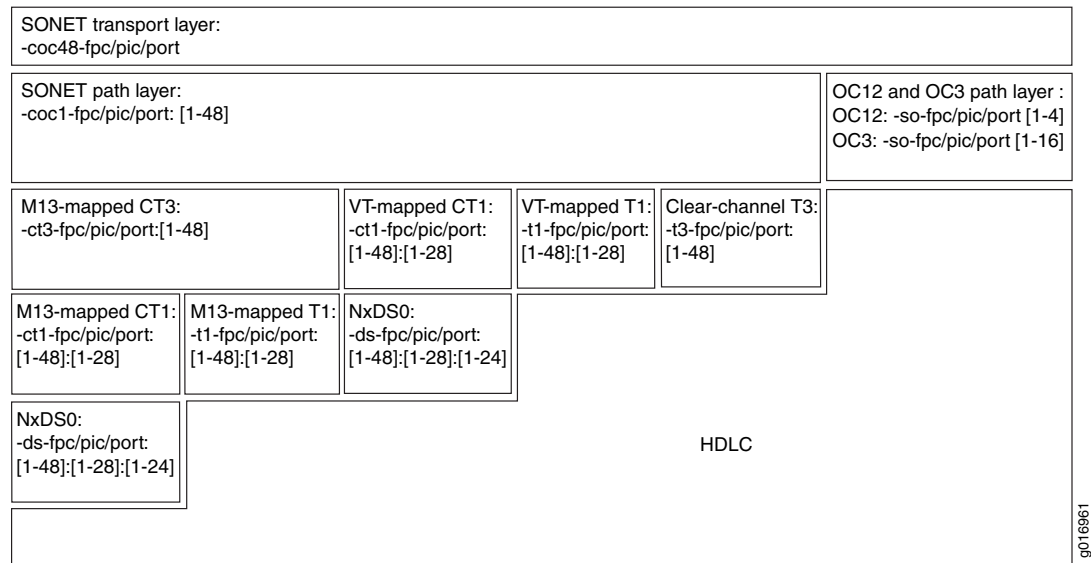


Figure 23: Channelized OC48/STM16 IQE PIC (in SDH Mode)

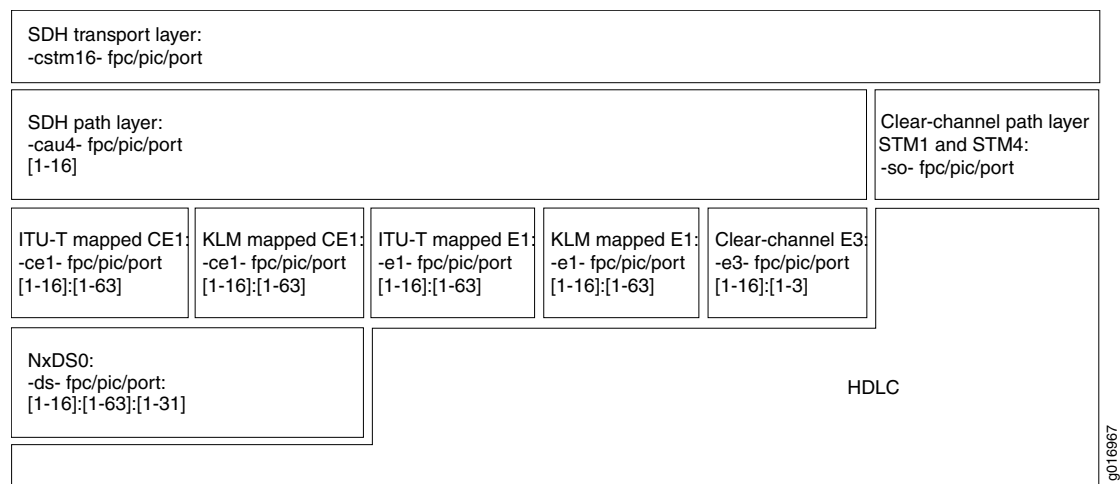


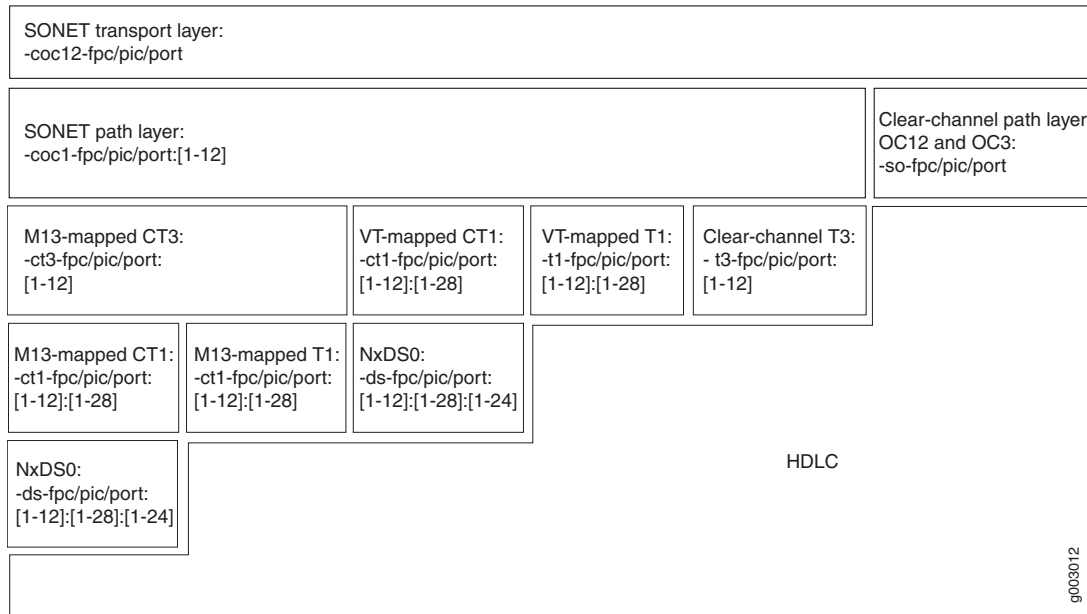
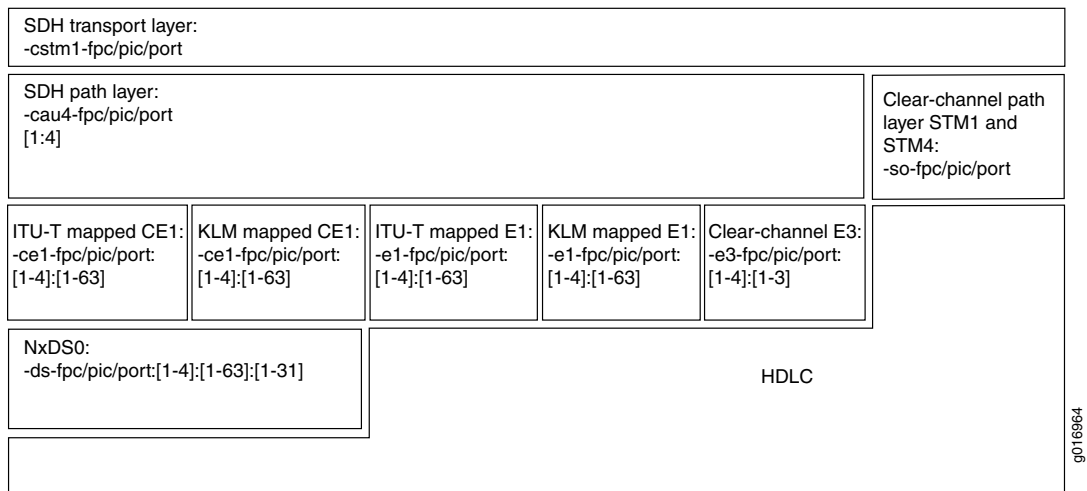
Figure 24: Channelized OC12 IQ PIC and Channelized OC12/STM4 IQE PIC (in SONET Mode)**Figure 25: Channelized OC12/STM4 IQE PIC (in SDH Mode)**

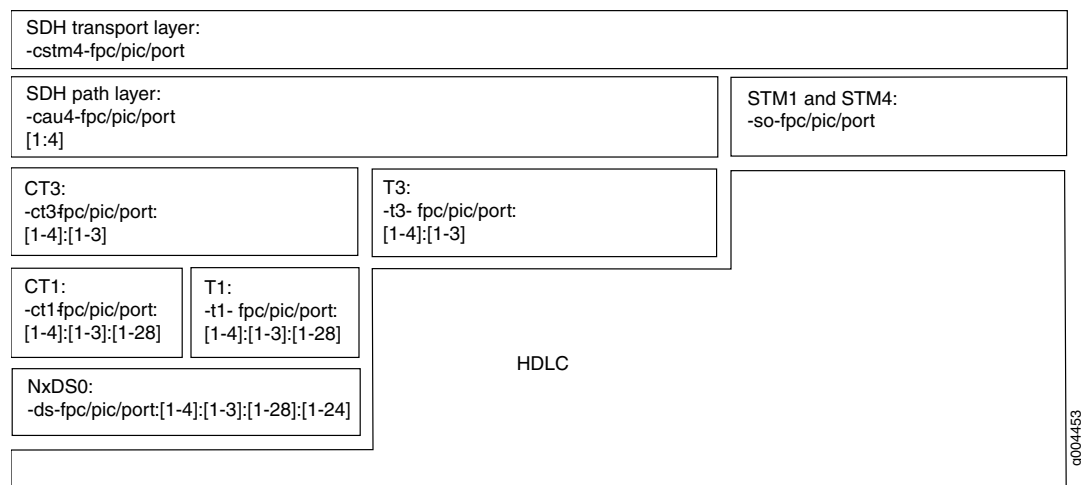
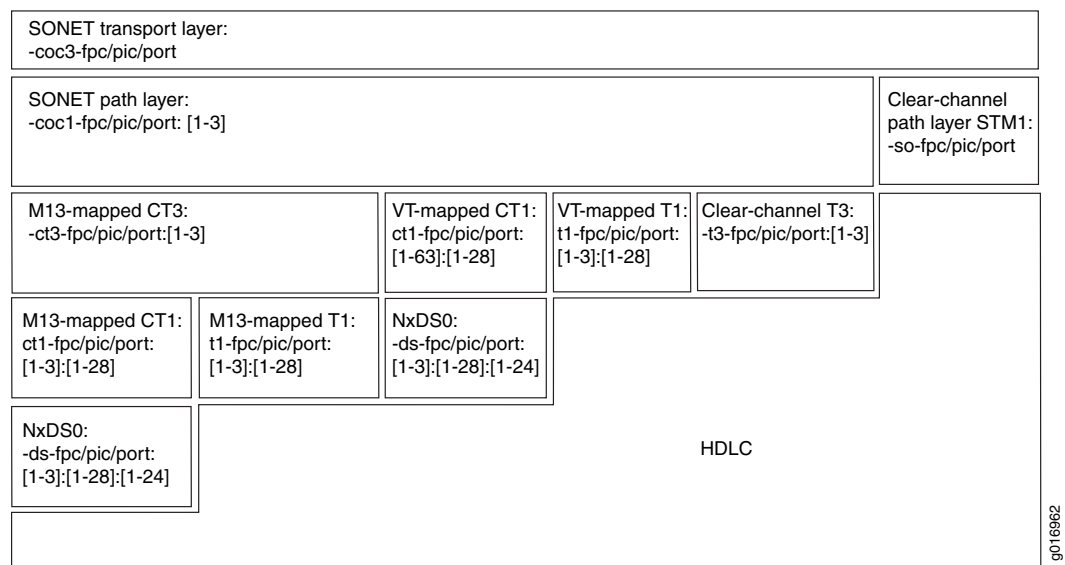
Figure 26: Channelized OC12/STM4 IQ PIC (in SDH Mode)**Figure 27: Channelized OC3 Ports (in SONET Mode) on Channelized OC3 IQ and Channelized OC3/STM1 IQE PICs**

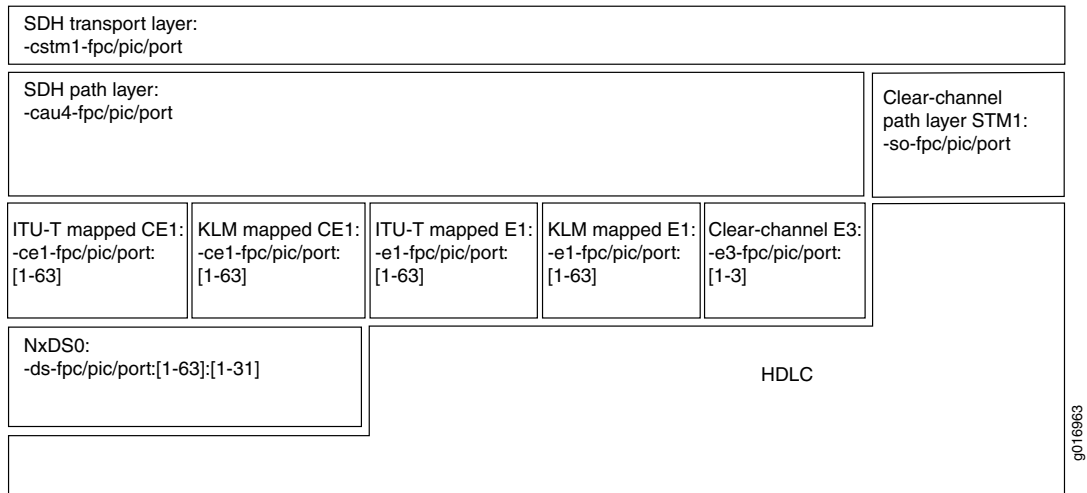
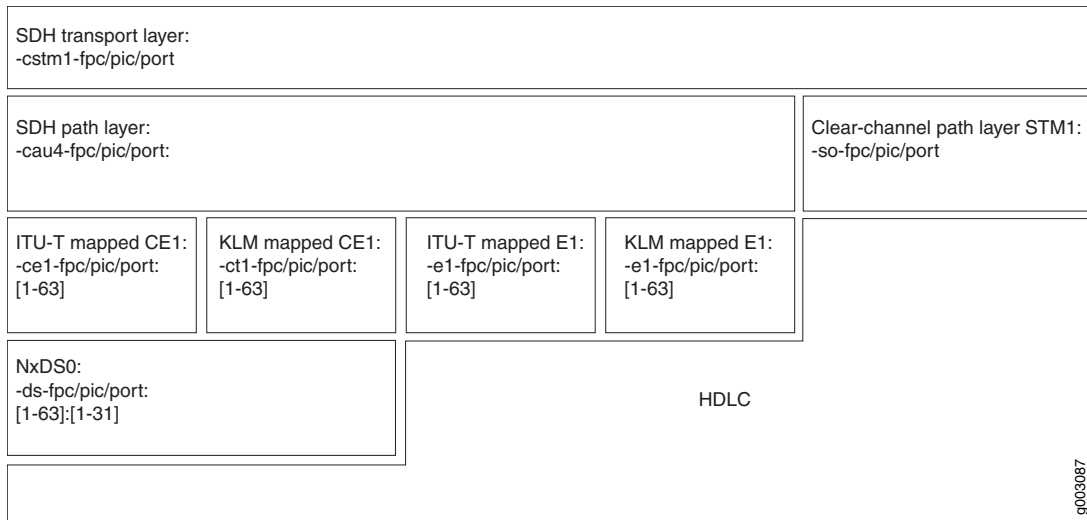

Figure 28: Channelized CSTM1 Ports (in SDH Mode) on Channelized OC3/STM1 IQE PIC**Figure 29: Channelized STM1 IQ PIC****Figure 30: Channelized CDS3/E3 IQE PIC (in DS3 Mode)**

Figure 31: Channelized CDS3/E3 IQE PIC (in E3 Mode)




Figure 32: Channelized DS3 IQ PIC




Figure 33: Channelized T1 IQ and IQE PIC




Figure 34: Channelized E1 IQ and IQE PIC




Table 38: Structural Differences: Channelized IQE PICs

PIC Type	Transport	Path	DS3	DS1/E1	E3
Channelized IQE PICs					
Channelized OC48/STM16 IQE (SONET Mode)	<i>coc48-fpc/pic/port</i>	<i>coc1-fpc/pic/port</i> :[1-48]	<i>ct3-fpc/pic/port</i> :[1-48]	<i>ct1-fpc/pic/port</i> :[1-48]:[1-28]	Not applicable.
		<i>so-fpc/pic/port</i>	<i>t3-fpc/pic/port</i> :[1-48]	<i>t1-fpc/pic/port</i> :[1-48]:[1-28]	
Channelized OC48/STM16 IQE (SDH Mode)	<i>cstm16-fpc/pic/port</i>	<i>cau4-fpc/pic/port</i> :[1-16]	Not applicable.	<i>ce1-fpc/pic/port</i> :[1-16]:[1-63]	<i>e3-fpc/pic/port</i> :[1-16]:[1-3]
		<i>so-fpc/pic/port</i>		<i>e1-fpc/pic/port</i> :[1-16]:[1-63]	
Channelized OC12 IQE (SONET Mode)	<i>coc12-fpc/pic/port</i>	<i>coc1-fpc/pic/port</i> :[1-12]	<i>ct3-fpc/pic/port</i> :[1-12]	<i>ct1-fpc/pic/port</i> :[1-12]:[1-28]	Not applicable.
		<i>so-fpc/pic/port</i>	<i>t3-fpc/pic/port</i> :[1-12]	<i>t1-fpc/pic/port</i> :[1-12]:[1-28]	
Channelized STM4 IQE (SDH Mode)	<i>cstm4-fpc/pic/port</i>	<i>cau4-fpc/pic/port</i> : [1-4]	Not applicable.	<i>ce1-fpc/pic/port</i> :[1-4]:[1-63]	<i>e3-fpc/pic/port</i> :[1-4]:[1-3]
		<i>so-fpc/pic/port</i>		<i>e1-fpc/pic/port</i> :[1-4]:[1-63]	
Channelized OC3 IQE (SONET)	<i>coc3-fpc/pic/port</i>	<i>coc1-fpc/pic/port</i> :[1-3]	<i>ct3-fpc/pic/port</i> :[1-3]	<i>ct1-fpc/pic/port</i> :[1-3]:[1-28]	Not applicable.
		<i>so-fpc/pic/port</i>	<i>t3-fpc/pic/port</i> :[1-3]	<i>t1-fpc/pic/port</i> :[1-3]:[1-28]	
Channelized STM1 IQE	<i>cstm1-fpc/pic/port</i>	<i>cau4-fpc/pic/port</i>	Not applicable.	<i>ce1-fpc/pic/port</i> :[1-63]	<i>e3-fpc/pic/port</i> :[1:3]]
		<i>so-fpc/pic/port</i>		<i>e1-fpc/pic/port</i> :[1-63]	
Channelized DS3 IQE	Not applicable.	Not applicable.	<i>ct3-fpc/pic/port</i> <i>t3-fpc/pic/port</i>	<i>ct1-fpc/pic/port</i> :[1-28] <i>t1-fpc/pic/port</i> :[1-28]	Not applicable.
Channelized E3 IQE	Not applicable.	Not applicable.	Not applicable.	Not applicable.	<i>e3-fpc/pic/port</i> :[1:4]
Channelized T1 IQE	Not applicable.	Not applicable.	Not applicable.	<i>ct1-fpc/pic/port</i> <i>t1-fpc/pic/port</i>	Not applicable.
Channelized E1 IQE	Not applicable.	Not applicable.	Not applicable.	<i>ce1-fpc/pic/port</i> <i>e1-fpc/pic/port</i>	Not applicable.

Table 39: Structural Differences: Channelized IQ PICs

PIC Type	Transport	Path	DS3	DS1/E1	E3
Channelized IQ PICs					
Channelized OC12/STM4 IQ (SONET Mode)	coc12-fpc/pic/port	coc1-fpc/pic/port :[1-12]	ct3-fpc/pic/port :[1-4]:[1-3]	ct1-fpc/pic/port :[1-3]:[1-28]	Not applicable.
		so-fpc/pic/port	t3-fpc/pic/port :[1-4]:[1-3]	ct1-fpc/pic/port :[1-4]:[1-3]:[1-28]	
Channelized OC12/STM4 IQ (SDH Mode)	cstm4-fpc/pic/port	cau4-fpc/pic/port:[1-4]	ct3-fpc/pic/port :[1-4]:[1-3]	ct1-fpc/pic/port :[1-3]:[1-28]	Not applicable.
		so-fpc/pic/port	t3-fpc/pic/port :[1-4]:[1-3]	t1-fpc/pic/port :[1-4]:[1-3]:[1-28]	
Channelized OC3 IQ (SONET)	coc3-fpc/pic/port	coc1-fpc/pic/port :[1-3]	ct3-fpc/pic/port :[1-3]	ct1-fpc/pic/port :[1-3]:[1-28]	Not applicable.
		so-fpc/pic/port	t3-fpc/pic/port :[1-3]	t1-fpc/pic/port :[1-3]:[1-28]	
Channelized STM1 IQ (SDH)	Not applicable.	cau4-fpc/pic/port so-fpc/pic/port	Not applicable.	ce1-fpc/pic/port :[1-63] e1-fpc/pic/port :[1-63]	Not applicable.
Channelized DS3 IQ	Not applicable.	Not applicable.	ct3-fpc/pic/port	ct1-fpc/pic/port :[1-28]	Not applicable.
			t3-fpc/pic/port	t1-fpc/pic/port :[1-28]	
Channelized E1 IQ	Not applicable.	Not applicable.	Not applicable.	ce1-fpc/pic/port e1-fpc/pic/port	Not applicable.

Table 40: Structural Differences: Channelized PICs

PIC Type	Transport	Path	DS3	DS1/E1	E3
Channelized PICs					
Channelized OC12	t3-fpc/pic/port :0	t3-fpc/pic/port :[0-11]	t3-fpc/pic/port :[0-11]	Not applicable.	Not applicable.
Channelized STM1	e1-fpc/pic/port :0	e1-fpc/pic/port :0	Not applicable.	e1-fpc/pic/port :[0-63]	Not applicable.
Channelized T3 and Multichannel T3	Not applicable.	Not applicable.	t1-fpc/pic/port :0	t1-fpc/pic/port :[0-27]	Not applicable.

Table 40: Structural Differences: Channelized PICs *(continued)*

PIC Type	Transport	Path	DS3	DS1/E1	E3
Channelized E1	Not applicable.	Not applicable.	Not applicable.	<i>e1-fpc/pic/port</i> <i>ds-fpc/pic/port</i> :0	Not applicable.

Chapter 18

Configuring Channelized OC48/STM16 IQE Interfaces

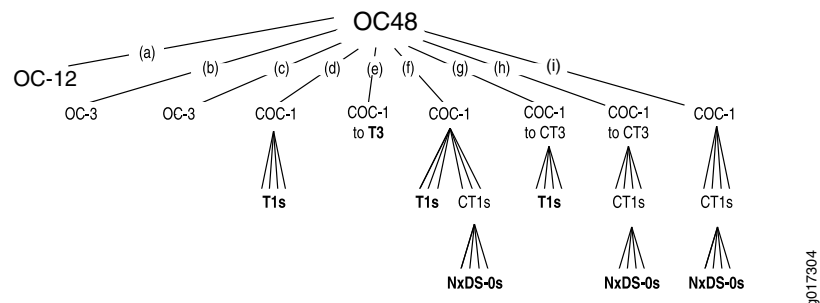
This section contains the following topics:

- Channelized OC48/STM16 IQE Interfaces Overview on page 405
- Configuring Channelized OC48/STM16 IQE Interfaces in SONET Mode on page 407
- Configuring Channelized OC48/STM16 IQE Interfaces (SDH Mode) on page 415
- Configuring Link PIC Failover on Channelized OC48/STM16 IQE Interfaces on page 419
- Example: Configuring Channelized OC48 Interfaces with Partitioned Channels on page 419

Channelized OC48/STM16 IQE Interfaces Overview

Channelized enhanced intelligent queuing (IQE) interfaces allow arbitrary and dynamic channelization of serial links, allowing greater flexibility than the channelized interfaces. Figure 35 on page 405, Figure 36 on page 406, and Figure 37 on page 407 illustrate the Channelized OC48/STM16 IQE Physical Interface Cards (PICs) in several examples of many possible configurations.

Figure 35: Sample Channelization of OC48/STM16 IQE PIC (SONET Mode)




Bold entries correspond to actual packet channels.

In the example in Figure 35 on page 405, a Channelized OC48/STM16 IQE PIC operating in SONET mode is partitioned into the following OC slices:

- a. A clear channel OC12 interface.
- b. An OC3 interface.
- c. A channelized COC1 partitioned into T1 interfaces.
- d. A channelized COC1 partitioned into a T3 interface.
- e. A channelized COC1 partitioned into CT3, partitioned into T1 interfaces, and CT1s partitioned into NxDS0 interfaces.
- f. A channelized COC1 partitioned into CT3, partitioned into T1 interfaces.
- g. A channelized COC1 partitioned into CT3, partitioned into CT1s, partitioned into NxDS0 interfaces.
- h. A channelized COC1 partitioned into CT1s, partitioned into NxDS0 interfaces.

This is one of thousands of ways to configure a Channelized OC48/STM16 IQE PIC. To configure the interfaces shown in Figure 36 on page 406, see “Configuring Channelized OC48/STM16 IQE Interfaces (SDH Mode)” on page 415.

Figure 36: Sample Channelization of OC48/STM16 IQE PIC (SDH Mode)



In Figure 36 on page 406, a Channelized OC48/STM16 IQE PIC operating in SDH mode results in a channelized STM16 interface, which can be partitioned as the following:

- a. Up to 4 STM4s.
- b. Up to 16 STM1s.
- c. Up to 16 CAU4s that can each be partitioned into up to 63 E1s, up to 3 E3s, or up to 63 CE1s. Each CE1 can be partitioned into up to 31 NxDS0s.

This is one of thousands of ways to configure a Channelized OC48/STM16 IQE PIC.


Figure 37: Sample Channelization of OC48/STM16 IQE PIC to E3 Channels

Figure 37 on page 407 shows five E3 channels configured on the Channelized OC48/STM16 IQE PIC. You can configure 43 additional E3 channels. For more information about configuring E3 channels on Channelized OC48/STM16 IQE PICs, see “Configuring E3 Interfaces” on page 417.

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring Channelized OC48/STM16 IQE Interfaces in SONET Mode

This section describes how to configure channelized OC48/STM16 IQE interfaces, discussing the following topics:

- Configuring OC12 Interfaces on page 407
- Configuring OC3 Interfaces on page 408
- Configuring T3 Interfaces on page 409
- Configuring T1 Interfaces on page 410
- Configuring Fractional T1 Interfaces on page 412
- Configuring NxDS0 Interfaces on page 413

Configuring OC12 Interfaces

You can configure up to four OC12 interfaces on a 1-port Channelized OC48/STM16 IQE PIC. To configure an OC12 interface, include the **partition**, **oc-slice**, and **interface-type** statements at the [edit interfaces coc48-fpc/pic/port] hierarchy level, specifying the **so** interface type:

```
[edit interfaces coc48-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type so;
```

The partition number is the sublevel interface partition index. For SONET/SDH interfaces, the partition number does not correlate with bandwidth size. For OC12 interfaces, the partition number can be from 1 through 4.



NOTE: For channelized OC48 IQE interfaces, channel numbering begins with 1 (:1).

The OC-slice range is the range of SONET/SDH slices. For SONET/SDH interfaces, the OC-slice range specifies the bandwidth size required for the interface type you are configuring. OC12 interfaces must occupy 12 consecutive OC slices per interface, in one of the following forms:

- 1–12
- 13–24
- 25–26
- 37–48

By contrast, the T3 and OC1 interfaces each occupy one OC slice per interface and OC3 interfaces occupy three slices per interface.

The interface type is the channelized interface type or data channel you are creating. For channelized OC48 IQE interfaces, the interface type can be `so`.

Example: Configuring OC12 Interfaces

Configure an OC12 interface, using partition 1 and OC slices 1 through 12. This configuration creates interface `so-1/1/0:1`.

```
[edit interfaces coc48-1/1/0]
partition 1 oc-slice 1-12 interface-type so;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring OC3 Interfaces

To configure an OC3 interface, include the `partition`, `oc-slice`, and `interface-type` statements at the `[edit interfaces coc48-fpc/pic/port]` hierarchy level, specifying the `so` interface type:

```
[edit interfaces coc48-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type so;
```

The partition number is the sublevel interface partition index. For SONET/SDH interfaces, the partition number does not correlate with bandwidth size. For OC3 interfaces, the partition number can be from 1 through 16.



NOTE: For channelized OC48 IQE interfaces, channel numbering begins with 1 (:1).

The OC-slice range is the range of SONET/SDH slices. For SONET/SDH interfaces, the OC-slice range specifies the bandwidth size required for the interface type you are configuring. OC3 interfaces must occupy three consecutive OC slices per interface, in one of the following forms:

- 1–3
- 4–6
- 7–9
- 10–12
- and so on (in groups of 3), up to 48

By contrast, the T3 and OC1 interfaces each occupy one OC slice per interface.

The interface type is the channelized interface type or data channel you are creating. For channelized OC48 IQE interfaces, the interface type can be **so**.

Example: Configuring OC3 Interfaces

Configure an OC3 interface, using partition 1 and OC slices 4 through 6. This configuration creates interface **so-1/1/0:1**.

```
[edit interfaces coc48-1/1/0]
partition 1 oc-slice 4-6 interface-type so;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring T3 Interfaces

To configure a T3 interface on an OC48/STM16 IQE PIC, include the **partition**, **oc-slice**, and **interface-type** statements at the **[edit interfaces coc48-fpc/pic/port]** hierarchy level, specifying the **coc1** interface type:

```
[edit interfaces coc48-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type coc1;
```

This configuration creates interface **coc1-fpc/pic/port:channel**.

Then, include the **no-partition interface-type** statement at the **[edit interfaces coc1-fpc/pic/port:channel]** hierarchy level, specifying the **t3** interface type:

```
[edit interfaces coc1-fpc/pic/port:channel]
no-partition interface-type t3;
```

This configuration creates interface **t3-fpc/pic/port:channel**.

The partition number is the sublevel interface partition index and is correlated with the channel number. For channelized OC1 interfaces, the partition number can be from 1 through 48. For channelized OC48/STM16 IQE interfaces, channel numbering begins with 1 (:1).

The OC-slice range is the range of SONET/SDH slices. For SONET/SDH interfaces, the OC-slice range specifies the bandwidth size required for the interface type you are configuring. For channelized OC1 interfaces, the OC slice can be from 1 through 12. You can configure only one OC slice per channelized OC1 interface.

The interface type is the channelized interface type or clear channel you are creating. For channelized OC48 interfaces, **type** can be **so** or **coc1**.



NOTE: Channelized OC48/STM16 IQE interfaces in M Series, MX Series, and T Series routers reserve channels 0 through 3 of each OC12 space for STS3C SONET channels.

When you configure E3 or T3 channels in OC12 spaces on the described PICs, the JUNOS Software allocates them starting from channel 4 because channels 0 through 3 are reserved for four STS3c SONET channels. Channel numbers are allocated sequentially in the following order: 4, 5, 6, 7, 8, 9, 11, 0, 1, 2, 3.

Only after channels 4 through 11 of the OC12 space are exhausted (all 4 through 11 configured) for E3 or T3 channels will the JUNOS Software then allocate channel 0 through 3 space for further E3 or T3 channels; thereby using up the 0 through 3 space previously reserved for four STS3c SONET channels.

If a subsequent reconfiguration of this OC12 space occurs, where you try to replace channels 4 through 6 or 7 through 9 with an OC3 SONET channel; the configuration fails because the channel 0 through 3 space is already occupied by the last E3 or T3 channels configured. This causes a failure in channel allocation and the Device Control Daemon (DCD) keeps retrying forever to configure the channel allocation on the interface. The only resolution is to reconfigure the last configured E3/T3 channels with OC3 channels, to free channels 0 through 3.

Example: Configuring T3 Interfaces

Configure a T3 interface using partition 3 and OC slice 3. This configuration creates interface t3-1/1/0:3:

```
[edit interfaces coc48-1/1/0]
partition 3 oc-slice 3 interface-type coc1;
[edit interfaces coc1-1/1/0:3]
no-partition interface-type t3;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring T1 Interfaces

To configure T1 interfaces on a Channelized OC48 IQE PIC, perform the following tasks:

1. Partition the channelized OC48 IQE interface into channelized OC1 interfaces by including the **partition**, **oc-slice**, and **interface-type** statements at the `[edit interfaces coc48-fpc/pic/port]` hierarchy level, specifying the **coc1** interface type:

```
[edit interfaces coc48-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type coc1;
```

2. If your network equipment uses VT mapping, partition the channelized OC1 interface into T1 interfaces by including the **partition** and **interface-type** statements at the `[edit interfaces coc1-fpc/pic/port]` hierarchy level, specifying the **t1** interface type:

```
[edit interfaces coc1-fpc/pic/port:channel]
partition partition-number interface-type t1;
```

3. If your network equipment uses M13 or C-bit parity, convert the channelized OC1 interface into a channelized T3 interface by including the **no-partition** and **interface-type** statements at the `[edit interfaces coc1-fpc/pic/port:channel]` hierarchy level, specifying the **ct3** interface type. Note that because the **no-partition** statement is included, this configuration does not create another level of channelization, as denoted by the number of colons in the resulting interface.

```
[edit interfaces coc1-fpc/pic/port:channel]
no-partition partition-number interface-type ct3;
```

4. Partition the channelized T3 interface into T1 interfaces by including the **partition** and **interface-type** statements at the `[edit interfaces ct3-fpc/pic/port:channel]` hierarchy level, specifying the **t1** interface type:


```
[edit interfaces ct3-fpc/pic/port:channel]
partition partition-number interface-type t1;
```



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQE interfaces. You can only apply CoS rules to the aggregate bit streams.

Figure 38 on page 411 shows VT-mapped and M13 or C-bit parity-mapped configurations of T1 interfaces.

Figure 38: T1 Interfaces on a Channelized OC48 PIC



Bold entries correspond to actual packet channels.

g017307

Example: Configuring T1 Interfaces

Configure the following T1 interfaces:

```
t1-0/0/0:1:1
t1-0/0/0:1:2
t1-0/0/0:1:3
t1-0/0/0:1:4
t1-0/0/0:1:5
```

VT-Mapped Configuration [edit interfaces coc48-0/0/0]
partition 1 oc-slice 1 interface-type coc1;

[edit interfaces coc1-0/0/0:1]
partition 1-5 interface-type t1;

M13 or C-bit Parity-Mapped Configuration [edit interfaces coc48-0/0/0]
partition 1 oc-slice 1 interface-type coc1;

[edit interfaces coc1-0/0/0:1]
no-partition interface-type ct3;

[edit interfaces ct3-0/0/0:1]
partition 1-5 interface-type t1;

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring Fractional T1 Interfaces

By default, all the time slots on a channelized T1 interface are used. To configure a fractional T1 interface on a Channelized OC48 IQE PIC, perform the following tasks:

1. Configure a T1 interface. For more information, see “Configuring T1 Interfaces” on page 410.
2. Configure the number of time slots allocated to the T1 interface by including the `timeslots` statement at the [edit interfaces t1-fpc/pic/port<:channel> t1-options] hierarchy level:

```
[edit interfaces t1-fpc/pic/port<:channel> t1-options]
timeslots time-slot-range;
```

For channelized T1 interfaces, the time-slot range is from 1 through 24. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces. For more information, see “Configuring Fractional T1 Time Slots” on page 567.

Example: Configuring Fractional T1 Interfaces

Configure a fractional T1 interface that uses time slots 1 through 5 and 10:

```
[edit interfaces coc48-0/0/0]
partition 1 oc-slice 1 interface-type coc1;
[edit interfaces coc1-0/0/0:1]
partition 1 interface-type t1;
[edit interfaces t1-0/0/0:1:1 t1-options]
timeslots 1-5,10;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring NxDS0 Interfaces

To configure NxDS0 interfaces on a Channelized OC48 IQE PIC, perform the following tasks:

1. Partition the channelized OC48 IQE interface into channelized OC1 interfaces by including the **partition**, **oc-slice**, and **interface-type** statements at the [edit interfaces coc48-fpc/pic/port:channel] hierarchy level, specifying the coc1 interface type:

```
[edit interfaces coc48-fpc/pic/port:channel]
partition partition-number oc-slice oc-slice-range interface-type coc1;
```

2. If your network equipment uses VT mapping, partition the channelized OC1 interface into channelized T1 interfaces by including the **partition** and **interface-type** statements at the [edit interfaces coc1-fpc/pic/port:channel] hierarchy level, specifying the ct1 interface type:

```
[edit interfaces coc1-fpc/pic/port:channel]
partition partition-number interface-type ct1;
```



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQE interfaces. You can only apply CoS rules to the aggregate bit streams.

If your network equipment uses M13 or C-bit parity, convert the channelized OC1 interface into a channelized T3 interface by including the **no-partition** and **interface-type** statements at the [edit interfaces coc1-fpc/pic/port] hierarchy level, specifying the ct3 interface type:

```
[edit interfaces coc1-fpc/pic/port:channel]
no-partition partition-number interface-type ct3;
```



NOTE: Because the **no-partition** statement is included, this configuration task does not create another level of channelization, as denoted by the number of colons in the resulting interface.

-
3. Partition the channelized T3 interface into channelized T1 interfaces by including the **partition** and **interface-type** statements at the [edit interfaces ct3-fpc/pic/port:channel] hierarchy level, specifying the ct1 interface type:

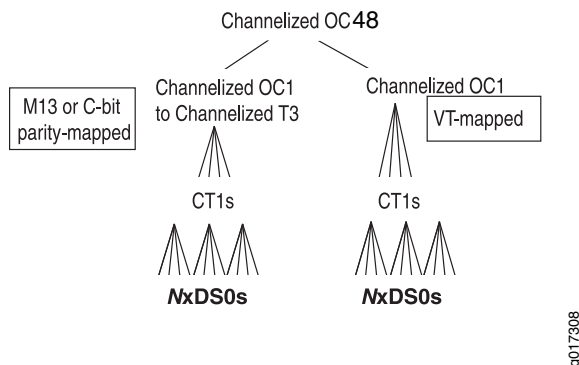
```
[edit interfaces ct3-fpc/pic/port:channel]
partition partition-number interface-type ct1;
```

4. Configure channelized NxDS0 interfaces on the channelized T1 interface by including the **partition**, **timeslots**, and **interface-type** statements at the [edit interfaces ct1-fpc/pic/port:channel] hierarchy level, specifying the ds interface type:

```
[edit interfaces ct1-fpc/pic/port:channel:channel]
partition partition-number timeslots time-slot-range interface-type ds;
```

Figure 39 on page 414 shows VT-mapped and M13 or C-bit parity-mapped configurations of NxDS0 interfaces.

Figure 39: Sample Channelization of OC48 IQE PIC



Bold entries correspond to actual packet channels.

Example: Configuring NxDS0 Interfaces

Configure the following two NxDS0 interfaces with 10 time slots and 4 time slots, respectively:

VT-Mapped Configuration

```
ds-0/0/0:1:2:1
ds-0/0/0:1:2:2

[edit interfaces coc48-0/0/0]
partition 1 oc-slice 1 interface-type coc1;

[edit interfaces coc1-0/0/0:1]
partition 2 interface-type ct1;

[edit interfaces ct1-0/0/0:1:2]
partition 1 timeslots 1-10 interface-type ds;
partition 2 timeslots 12-15 interface-type ds;
```

M13 or C-bit Parity-Mapped Configuration

```
[edit interfaces coc48-0/0/0]
partition 1 oc-slice 1 interface-type coc1;

[edit interfaces coc1-0/0/0:1]
no-partition interface-type ct3;

[edit interfaces ct3-0/0/0:1]
partition 2 interface-type ct1;

[edit interfaces ct1-0/0/0:1:2]
partition 1 timeslots 1-10 interface-type ds;
```

partition 2 timeslots 12-15 interface-type ds;

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring Channelized OC48/STM16 IQE Interfaces (SDH Mode)

The Channelized OC48 IQE PIC configured for SDH mode creates a single channelized STM16 interface. You can configure the STM16 interface using the **partition** statement at the [edit interfaces cstm16-fpc/pic/port] hierarchy level to partition it into the following OC slices:

- 16 channelized AU-4 interfaces or a path layer with 4 STM4 or 16 STM1 interfaces.
- 16 channelized AU-4 interfaces, each partitioned to 3 clear channel E3 interfaces or 63 CE1 or E1 (ITU-T or KLM) interfaces. Combination of E1, CE1 and E3 are not supported in a single cau4.
- 16 channelized AU-4 interfaces, each partitioned to 63 CE1 (ITU-T or KLM) interfaces each partitioned to 31 NxDS0 interfaces

This section describes how to configure the following channelized OC48 IQE interfaces on a Channelized OC48 IQE PIC configured in SDH mode:

- Configuring a Channelized OC48/STM16 IQE PIC for SDH Mode on page 415
- Configuring Clear Channel STM1 and STM4 Interfaces on page 416
- Configuring Channelized AU-4 Interfaces on page 416
- Configuring E3 Interfaces on page 417
- Configuring E1 or Channelized E1 Interfaces on page 418
- Configuring NxDS0 IQE Interfaces on page 418

Configuring a Channelized OC48/STM16 IQE PIC for SDH Mode

To configure a Channelized OC48/STM16 IQE PIC to operate in SDH mode, include the **framing sdh** statement at the [edit chassis fpc fpc/pic/port] hierarchy level:

```
[edit chassis ]
  fpc 0 {
    pic 2 {
      framing sdh;
    }
  }
}
```

This configuration creates interface cstm16-0/2/0.

For more information, see the *JUNOS System Basics Configuration Guide*.

Configuring Clear Channel STM1 and STM4 Interfaces

On a Channelized OC48/STM16 IQE PIC, you can partition the CSTM16 transport layer into 4 clear channel STM4 interfaces or 16 clear channel STM1 interfaces. Combinations of STM4 and STM1 are also permitted, but you must observe the OC-slice parameters.

To configure an STM4 interface, include the **partition** and **interface-type** statements at the [edit interfaces cstm16-fpc/pic/port] hierarchy level:

```
[edit interfaces cstm16-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type so;
```

This configuration creates interface *so-fpc/pic/port.channel*.

To configure an STM1 interface, include the **partition** and **interface-type** statements at the [edit interfaces cstm16-fpc/pic/port] hierarchy level:

```
[edit interfaces cstm16-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type so;
```

This configuration creates interface *so-fpc/pic/port.channel*.

Configuring Channelized AU-4 Interfaces

To configure a channelized AU-4 interface, include the **partition**, **oc-slice**, and **interface-type** statements at the [edit interfaces cstm16-fpc/pic/port:channel] hierarchy level, specifying the **cau4** interface type:

```
[edit interfaces cstm16-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type cau4;
```

This configuration creates interface *cau4-fpc/pic/port:channel*.

The partition number is the sublevel interface partition index. For SDH interfaces, the partition number is not correlated with bandwidth size. For channelized OC48/STM16 IQE interfaces, channelized STM16 interface can have from 1 through 16 partition numbers and channel numbering begins with 1 (:1).

The OC-slice range is the range of SONET/SDH slices. For SDH interfaces, the OC-slice range specifies the bandwidth size required for the interface type you are configuring. The interface type is the channelized interface type or data channel you are creating.

Example: Configuring Channelized AU-4 Interfaces

Configure channelized AU-4 interfaces:

```
[edit interfaces cstm16-0/2/0]
partition 1 oc-slice 1-3 interface-type cau4;
```


Configuring E3 Interfaces

To configure E3 interfaces, include the `partition` and `interface-type` statements at the `[edit interfaces cau4-fpc/pic/port]` hierarchy level, specifying the `e3` interface type:

```
[edit interfaces]
cau4-fpc/pic/port {
  partition partition-number interface-type e3;
}
```

This configuration creates the interfaces `e3-fpc/pic/port:channel`.



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQE interfaces. You can only apply CoS rules to the aggregate bit streams.



NOTE: Channelized OC48/STM16 IQE interfaces in M Series, MX Series, and T Series routers reserve channels 0-3 of each OC12 space for STS3C SONET channels.

When you configure E3 or T3 channels in OC12 spaces on the described PICs, JUNOS Software allocates them starting from channel 4 because channels 0-3 are reserved for four STS3c SONET channels. Channel numbers are allocated sequentially in the following order: 4, 5, 6, 7, 8, 9, 11, 0, 1, 2, 3.

Only after channels 4 through 11 of the OC12 space are exhausted (all 4 through 11 configured) for E3 or T3 channels will JUNOS Software then allocate channel 0-3 space for further E3 or T3 channels; thereby using up the 0-3 space previously reserved for four STS3c SONET channels.

If a subsequent reconfiguration of this OC12 space occurs, where you try to replace channels 4-6 or 7-9 with an OC3 SONET channel; it fails because the channel 0-3 space is already occupied by the last E3 or T3 channels configured. This causes a failure in channel allocation and the Device Control Daemon (DCD) keeps retrying forever to configure the channel allocation on the interface. The only resolution is to reconfigure the last configured E3/T3 channels with OC3 channels, to free channels 0-3.

Example: Configuring E3 Interfaces

Configure E3 interfaces, using partition 1:

```
[edit interfaces]
cau4-0/2/0:1 {
  partition 1 interface-type e3;
}
e3-0/2/0:1:1;
```

Configuring E1 or Channelized E1 Interfaces

To configure E1 or channelized E1 interfaces, include the `partition` and `interface-type` statements at the `[edit interfaces cau4-fpc/pic/port]` hierarchy level, specifying the `e1` or `ce1` interface type:

```
[edit interfaces]
cau4-fpc/pic/port {
  partition partition-number interface-type e1;
}
cau4-fpc/pic/port {
  partition partition-number interface-type ce1;
}
```

This configuration creates the interfaces `e1-fpc/pic/port:channel` and `ce1-fpc/pic/port:channel`.



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQE interfaces. You can only apply CoS rules to the aggregate bit streams.

Example: Configuring E1 and Channelized E1 Interfaces

Configure E1 or channelized E1 interfaces, using partition 3 and partition 4:

```
[edit interfaces]
cau4-0/2/0:1 {
  partition 3 interface-type e1;
}
cau4-0/2/0:1 {
  partition 4 interface-type ce1;
}
```

This configuration creates interfaces `e1-0/2/0:1:3` and `ce1-0/2/0:1:4`.

Configuring NxDS0 IQE Interfaces

Configure channelized NxDS0 IQE interfaces on the channelized E1 IQE interface by including the `partition`, `timeslots`, and `interface-type` statements at the `[edit interfaces ce1-fpc/pic/port:channel]` hierarchy level, specifying the `ds` interface type:

```
[edit interfaces ce1-fpc/pic/port:channel:channel]
partition partition-number timeslots time-slot-range interface-type ds;
```

This configuration creates the interface `ds-fpc/pic/port:channel`.

The time-slot range is from 1 through 31. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. You can use a combination of ranges and discontinuous time slots, for example:

1,9-18,21

Example: Configuring NxDS0 IQE Interfaces

Configure channelized NxDS0 interfaces, using partition 4 and time slots 1 through 10:

```
[edit interfaces]
ce1-0/2/0:1:2:3 {
  partition 4 interface-type ds0 timeslots 1-10;
}
```

This configuration creates interface ds0-0/2/0:1:2:3:4.

Configuring Link PIC Failover on Channelized OC48/STM16 IQE Interfaces

For Channelized OC48 IQE PICs used as linking PICs in redundant LSQ configurations, you can inhibit the router from sending PPP termination-request messages to the remote host if the link PIC fails. To do this, include the **no-termination-request** statement at the [edit interfaces *interface-name* ppp-options] hierarchy level:

```
no-termination-request;
```

The **no-termination-request** statement is supported only with MLPPP and SONET APS configurations and works with PPP, PPP over Frame Relay, and MLPPP interfaces only.

For information about interchassis and intrachassis LSQ failover, see the *JUNOS Services Interfaces Configuration Guide*.

Example: Configuring Channelized OC48 Interfaces with Partitioned Channels

The following configuration is sufficient to get the channelized OC48 interface up and running. The OC48 interface can be divided into up to 4 OC12 channels, up to 16 OC3 channels, or up to 48 OC1 channels and combinations are permitted; for example, 1 OC12, 4 OC3s, and 24 OC1s. There are 48 OC1 slices available on the OC48 IQE interface. An OC48 configuration uses all 48 slices, each OC12 uses 12 slices, each OC1 uses 1 slice. Permissible combinations must fit within the 48 available OC1 slices. DS1 channels can use the following encapsulation types:

- PPP, PPP CCC, and PPP TCC
- Frame Relay, Frame Relay CCC, and Frame Relay TCC
- Cisco HDLC, Cisco HDLC CCC, and Cisco HDLC TCC

The channels can also have logical interfaces.

```
[edit interfaces]
t3-fpc/pic/port:0 {
  encapsulation cisco-hdlc;
  t3-options {
    compatibility-mode larscom;
  }
}
```

```

        payload-scrambler;
    }
    unit 0 {
        family inet {
            address 10.11.30.1/30;
        }
        family iso;
    }
}
t3-fpc/pic/port:1 {
    encapsulation ppp;
    t3-options {
        compatibility-mode larscom;
        payload-scrambler;
    }
    unit 0 {
        family inet {
            address 10.11.30.5/30;
        }
        family iso;
    }
}
t3-fpc/pic/port:2 {
    encapsulation frame-relay;
    t3-options {
        compatibility-mode larscom;
        payload-scrambler;
    }
    unit 0 {
        dlci 100;
        family inet {
            address 10.11.30.9/30;
        }
        family iso;
    }
    unit 1 {
        dlci 101;
        family inet {
            address 10.11.31.9/30;
        }
        family iso;
    }
}
t3-1fpc/pic/port:3 {
    encapsulation cisco-hdlc-ccc;
    t3-options {
        compatibility-mode larscom;
        payload-scrambler;
    }
    unit 0;
}
t3-fpc/pic/port:4 {
    encapsulation ppp-ccc;
    t3-options {
        compatibility-mode larscom;
        payload-scrambler;
    }
}

```

```

    }
    unit 0;
}
t3-fpc/pic/port:5 {
    dce;
    encapsulation frame-relay-ccc;
    t3-options {
        compatibility-mode larscom;
        payload-scrambler;
    }
    unit 0 {
        encapsulation frame-relay-ccc;
        dlc1 1000;
    }
    unit 1 {
        encapsulation frame-relay-ccc;
        dlc1 1001;
    }
}

```


Chapter 19

Configuring Channelized OC12/STM4 Interfaces

This section contains the following topics:

- Channelized OC12/STM4 IQ and IQE Interfaces Overview on page 423
- Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode) on page 427
- Configuring Channelized OC12/STM4 IQE Interfaces (SDH Mode) on page 434
- Configuring Channelized OC12/STM4 IQ Interfaces (SDH Mode) on page 440
- Configuring Channelized OC12 Interfaces on page 445
- Configuring Link PIC Failover on Channelized OC12/STM4 IQ and IQE Interfaces on page 447
- Example: Configuring a Channelized OC12 IQ Interface as an Unpartitioned, Clear Channel on page 448
- Example: Configuring Channelized OC12 Interfaces with Partitioned Channels on page 451

Channelized OC12/STM4 IQ and IQE Interfaces Overview

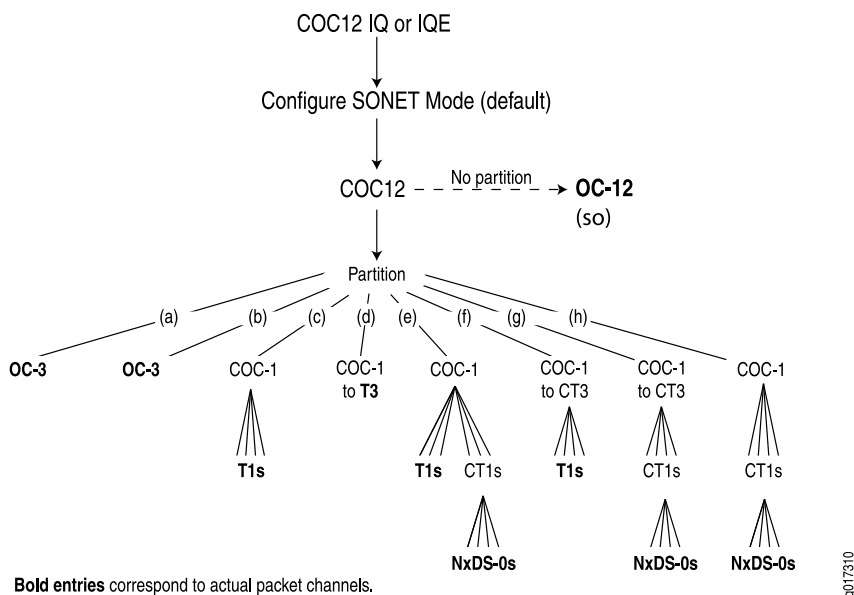
Channelized IQ and channelized IQE interfaces allow arbitrary and dynamic channelization of serial links, allowing greater flexibility than the channelized interfaces. Channelized OC12/STM4 IQ and IQE Physical Interface Cards (PICs) can be configured to operate in SONET or SDH mode. Each physical port on a multiple-port IQE PIC can be configured to operate in either SONET or SDH mode for increased granularity. The following sections describe the different modes of operation and channelization possibilities.

- Channelization of OC12/STM4 IQ and Channelized OC12/STM4 IQE PICs (SONET Mode) on page 424
- Channelization of OC12/STM4 IQE PIC (SDH Mode) on page 425
- Channelization of OC12/STM4 IQ PIC (SDH Mode) on page 425
- Channelization of OC12 PIC (SONET Mode) on page 426

Channelization of OC12/STM4 IQ and Channelized OC12/STM4 IQE PICs (SONET Mode)

Channelized OC12/STM4 IQ PICs and Channelized OC12/STM4 IQE PICs can be configured to operate in SONET or SDH mode and partitioned into various partitions. Figure 40 on page 424 illustrates one possible channelization configuration for Channelized OC12/STM4 IQ and IQE PICs operating in SONET mode.

Figure 40: Sample Channelization of OC12/STM4 IQ or IQE PIC (SONET Mode)



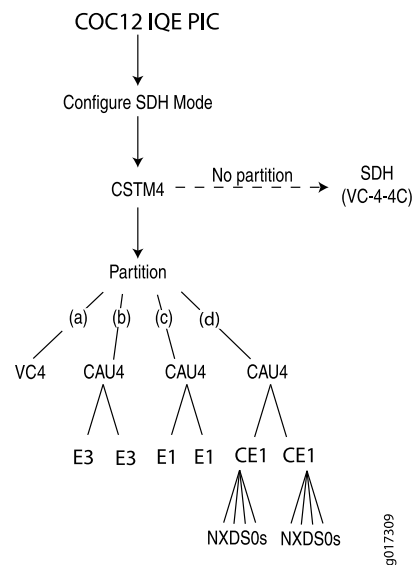
In the example in Figure 40 on page 424, a Channelized OC12/STM4 IQ PIC operating in SONET mode is partitioned into the following OC slices:

- An OC3 interface.
- Another OC3 interface.
- A channelized OC1 partitioned into T1 interfaces.
- A channelized OC1 converted into a T3 interface.
- A channelized OC1 partitioned into T1 interfaces and channelized T1s, which are partitioned into NxDS0 interfaces.
- A channelized OC1 converted into a channelized T3, which is partitioned into T1 interfaces.
- A channelized OC1 converted into a channelized T3, which is partitioned into T1 interfaces and a channelized T1, which is partitioned into NxDS0 interfaces.
- A channelized OC1 partitioned into channelized T1s, which are partitioned into NxDS0 interfaces.

Channelization of OC12/STM4 IQE PIC (SDH Mode)

Channelized OC12/STM4 IQE PICs can be configured to operate in SONET or SDH mode and partitioned to various smaller partitions. Figure 41 on page 425 illustrates one possible channelization configuration for Channelized OC12/STM4 IQE PICs operating in SDH mode.

Figure 41: Sample Channelization of OC12/STM4 IQE PIC (SDH Mode)



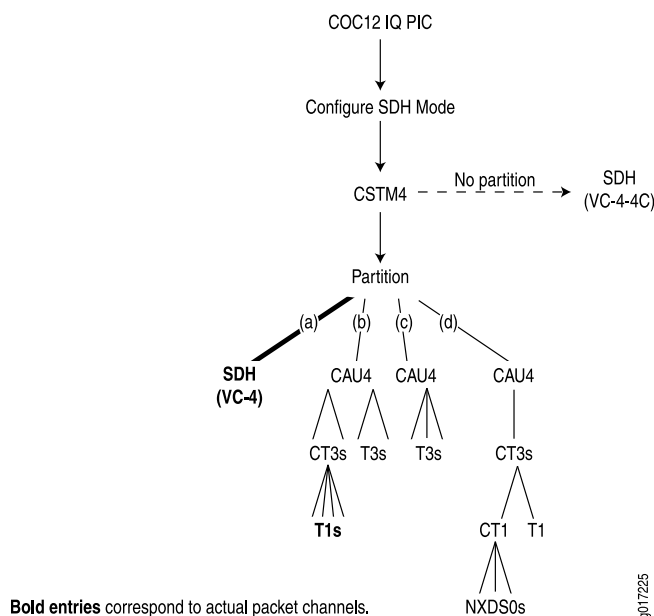
In Figure 41 on page 425, a Channelized OC12/STM4 IQE PIC operating in SDH mode results in a channelized STM4 interface, which can be nonpartitioned into one SDH VC-4-VC interface or partitioned into the following OC slices:

- An SDH VC-4 interface.
- A channelized AU-4 partitioned into E3 interfaces.
- A channelized AU-4 interface partitioned into E1 interfaces.
- A channelized AU-4 interface partitioned into CE1 interfaces partitioned into NxDS0 interfaces.

This is one of thousands of ways to configure a Channelized OC12/STM4 IQE PIC.

Channelization of OC12/STM4 IQ PIC (SDH Mode)

Channelized OC12/STM4 IQ PICs can be configured to operate in SONET or SDH mode and partitioned into various smaller partitions. Figure 42 on page 426 illustrates one possible channelization configuration for Channelized OC12/STM4 IQ PICs operating in SDH mode.

Figure 42: Sample Channelization of OC12/STM4 IQ PIC (SDH Mode)

In Figure 42 on page 426, a Channelized OC12/STM4 IQ PIC operating in SDH mode results in a channelized STM4 interface, which can be nonpartitioned into one SDH VC-4-VC interface or partitioned into the following OC slices:

- An SDH VC-4 interface.
- A channelized AU-4 partitioned into channelized T3 interfaces and T3 interfaces.
- Another channelized AU-4 interface converted into T3 interfaces.
- Another channelized AU-4 interface converted into a channelized T3 interface, which is partitioned further into a channelized T1 and a T1 interface. The channelized T1 interface is further partitioned into NxDS0 interfaces.

This is one of thousands of ways to configure a Channelized OC12/STM4 IQ PIC.

Channelization of OC12 PIC (SONET Mode)

OC12 PICs can be configured to various smaller partitions, such as T3s.


Figure 43: Sample Channelization of OC12 PIC (non IQ and IQE)

Figure 43 on page 426 shows five T3 channels configured on the Channelized OC12 PIC. You can configure seven additional T3 channels. For more information about configuring Channelized OC12 PICs, see “Configuring Channelized OC12 Interfaces” on page 445.

Configuring Channelized OC12/STM4 IQ and IQE Interfaces (SONET Mode)

This section describes how to configure channelized OC12/STM4 IQ and IQE interfaces, discussing the following topics:

- Configuring an OC12/STM4 Interface on page 427
- Configuring T3 Interfaces on page 427
- Configuring OC3 Interfaces on page 429
- Configuring T1 Interfaces on page 429
- Configuring NxDS0 Interfaces on page 431
- Configuring Fractional T1 Interfaces on page 433

Configuring an OC12/STM4 Interface

You can configure one OC12 interface on a one-port Channelized OC12/STM4 IQ or IQE PIC. On a 4-port OC12/STM4 IQ or IQE PIC, you can configure one OC12 interface per port. To configure an OC12 interface, include the **no-partition** and **interface-type** statements at the [edit interfaces coc12-fpc/pic/port] hierarchy level:

```
[edit interfaces coc12-fpc/pic/port]
no-partition interface-type so;
```

This configuration creates interface *so-fpc/pic/port*.



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQ and IQE interfaces. You can only apply CoS rules to the aggregate bit streams.

Configuring T3 Interfaces

To configure a T3 interface on an OC12 PIC, include the **partition**, **oc-slice**, and **interface-type** statements at the [edit interfaces coc12-fpc/pic/port] hierarchy level, specifying the **coc1** interface type:

```
[edit interfaces coc12-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type coc1;
```

This configuration creates interface *coc1-fpc/pic/port:channel*.

Then, include the **no-partition interface-type** statement at the [edit interfaces coc1-fpc/pic/port:channel] hierarchy level, specifying the **t3** interface type:

```
[edit interfaces coc1-fpc/pic/port:channel]
no-partition interface-type t3;
```

This configuration creates interface *t3-fpc/pic/port:channel*.

The partition number is the sublevel interface partition index and is correlated with the channel number. For channelized OC1 interfaces, the partition number can be from 1 through 12.



NOTE: For channelized OC12 interfaces, channel numbering begins with 0 (:0). For channelized OC12/STM4 IQ and IQE interfaces, channel numbering begins with 1 (:1).

The OC-slice range is the range of SONET/SDH slices. For SONET/SDH interfaces, the OC-slice range specifies the bandwidth size required for the interface type you are configuring. For channelized OC1 interfaces, the OC slice can be from 1 through 12. You can configure only one OC slice per channelized OC1 interface.

The interface type is the channelized interface type or clear channel you are creating. For channelized OC12 interfaces, **type** can be **so** or **coc1**.



NOTE: Channelized OC12/STM4 IQ and IQE interfaces in M Series, MX Series, and T Series routers reserve channels 0-3 of each OC12 space for STS3C SONET channels.

When you configure E3 or T3 channels in OC12 spaces on the described PICs, JUNOS Software allocates them starting from channel 4 because channels 0-3 are reserved for four STS3c SONET channels. Channel numbers are allocated sequentially in the following order: 4, 5, 6, 7, 8, 9, 11, 0, 1, 2, 3.

Only after channels 4 through 11 of the OC12 space are exhausted (all 4 through 11 configured) for E3 or T3 channels will JUNOS Software then allocate channel 0-3 space for further E3 or T3 channels; thereby using up the 0-3 space previously reserved for four STS3c SONET channels.

If a subsequent reconfiguration of this OC12 space occurs, where you try to replace channels 4-6 or 7-9 with an OC3 SONET channel; it fails because the channel 0-3 space is already occupied by the last E3 or T3 channels configured. This causes a failure in channel allocation and the Device Control Daemon (DCD) keeps retrying forever to configure the channel allocation on the interface. The only resolution is to reconfigure the last configured E3/T3 channels with OC3 channels, to free channels 0-3.

Example: Configuring T3 Interfaces

Configure a T3 interface using partition 3 and OC slice 3. This configuration creates interface **t3-1/1/0:3**:

```
[edit interfaces coc12-1/1/0]
partition 3 oc-slice 3 interface-type coc1;
[edit interfaces coc1-1/1/0:3]
no-partition interface-type t3;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring OC3 Interfaces

To configure an OC3 interface, include the `partition`, `oc-slice`, and `interface-type` statements at the `[edit interfaces coc12-fpc/pic/port]` hierarchy level, specifying the `so` interface type:

```
[edit interfaces coc12-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type so;
```

The partition number is the sublevel interface partition index. For SONET/SDH interfaces, the partition number does not correlate with bandwidth size. For OC3 interfaces, the partition number can be from 1 through 4.



NOTE: For channelized OC12 interfaces, channel numbering begins with 0 (:0). For channelized OC12 IQ and IQE interfaces, channel numbering begins with 1 (:1).

The OC-slice range is the range of SONET/SDH slices. For SONET/SDH interfaces, the OC-slice range specifies the bandwidth size required for the interface type you are configuring. OC3 interfaces must occupy three consecutive OC slices per interface, in one of the following forms:

- 1–3
- 4–6
- 7–9
- 10–12

By contrast, the T3 and OC1 IQ interfaces each occupy one OC slice per interface.

The interface type is the channelized interface type or data channel you are creating. For channelized OC12 interfaces, the interface type can be `coc1` or `so`.

Example: Configuring OC3 Interfaces

Configure an OC3 interface, using partition 1 and OC slices 4 through 6. This configuration creates interface `so-1/1/0:1`:

```
[edit interfaces coc12-1/1/0]
partition 1 oc-slice 4-6 interface-type so;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring T1 Interfaces

To configure T1 interfaces on a Channelized OC12 IQ or IQE PIC, perform the following tasks:

1. Partition the channelized OC12 interface into channelized OC1 interfaces by including the **partition**, **oc-slice**, and **interface-type** statements at the `[edit interfaces coc12-fpc/pic/port]` hierarchy level, specifying the **coc1** interface type:

```
[edit interfaces coc12-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type coc1;
```

2. If your network equipment uses VT mapping, partition the channelized OC1 interface into T1 interfaces by including the **partition** and **interface-type** statements at the `[edit interfaces coc1-fpc/pic/port]` hierarchy level, specifying the **t1** interface type:

```
[edit interfaces coc1-fpc/pic/port]
partition partition-number interface-type t1;
```

3. If your network equipment uses M13 or C-bit parity, convert the channelized OC1 interface into a channelized T3 interface by including the **no-partition** and **interface-type** statements at the `[edit interfaces coc1-fpc/pic/port:channel]` hierarchy level, specifying the **ct3** interface type. Note that because the **no-partition** statement is included, this configuration does not create another level of channelization, as denoted by the number of colons in the resulting interface.

```
[edit interfaces coc1-fpc/pic/port]
no-partition partition-number interface-type ct3;
```

4. Partition the channelized T3 interface into T1 interfaces by including the **partition** and **interface-type** statements at the `[edit interfaces ct3-fpc/pic/port]` hierarchy level, specifying the **t1** interface type:


```
[edit interfaces ct3-fpc/pic/port]
partition partition-number interface-type t1;
```



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQ interfaces. You can only apply CoS rules to the aggregate bit streams.

Figure 44 on page 430 shows VT-mapped and M13 or C-bit parity-mapped configurations of T1 interfaces.

Figure 44: T1 Interfaces on a Channelized OC12 PIC



Bold entries correspond to actual packet channels.

g003013

Example: Configuring T1 Interfaces

Configure the following T1 interfaces:

```
t1-0/0/0:1:1
t1-0/0/0:1:2
t1-0/0/0:1:3
t1-0/0/0:1:4
t1-0/0/0:1:5
```

VT-Mapped Configuration

```
[edit interfaces coc12-0/0/0]
partition 1 oc-slice 1 interface-type coc1;
```

```
[edit interfaces coc1-0/0/0:1]
partition 1-5 interface-type t1;
```

M13 or C-bit Parity-Mapped Configuration

```
[edit interfaces coc12-0/0/0]
partition 1 oc-slice 1 interface-type coc1;
```

```
[edit interfaces coc1-0/0/0:1]
no-partition interface-type ct3;
```

```
[edit interfaces ct3-0/0/0:1]
partition 1-5 interface-type t1;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring NxDS0 Interfaces

To configure NxDS0 interfaces on a Channelized OC12 IQE PIC, perform the following tasks:

1. Partition the channelized OC12 IQE interface into channelized OC1 interfaces by including the **partition**, **oc-slice**, and **interface-type** statements at the [edit interfaces coc12-fpc/pic/port] hierarchy level, specifying the coc1 interface type:

```
[edit interfaces coc12-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type coc1;
```

2. If your network equipment uses VT mapping, partition the channelized OC1 interface into channelized T1 interfaces by including the **partition** and **interface-type** statements at the [edit interfaces coc1-fpc/pic/port] hierarchy level, specifying the ct1 interface type:

```
[edit interfaces coc1-fpc/pic/port]
partition partition-number interface-type ct1;
```



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQ interfaces. You can only apply CoS rules to the aggregate bit streams.

3. If your network equipment uses M13 or C-bit parity, convert the channelized OC1 interface into a channelized T3 interface by including the **no-partition** and **interface-type** statements at the `[edit interfaces coc1-fpc/pic/port]` hierarchy level, specifying the **ct3** interface type:

```
[edit interfaces coc1-fpc/pic/port]
no-partition partition-number interface-type ct3;
```



NOTE: Because the **no-partition** statement is included, this configuration task does not create another level of channelization, as denoted by the number of colons in the resulting interface.

4. Partition the channelized T3 interface into channelized T1 interfaces by including the **partition** and **interface-type** statements at the `[edit interfaces ct3-fpc/pic/port]` hierarchy level, specifying the **ct1** interface type:


```
[edit interfaces ct3-fpc/pic/port]
partition partition-number interface-type ct1;
```

5. Configure channelized *N*xDS0 IQ interfaces on the channelized T1 IQ interface by including the **partition**, **timeslots**, and **interface-type** statements at the `[edit interfaces ct1-fpc/pic/port:channel:channel]` hierarchy level, specifying the **ds** interface type:

```
[edit interfaces ct1-fpc/pic/port:channel:channel]
partition partition-number timeslots time-slot-range interface-type ds;
```

Figure 45 on page 432 shows VT-mapped and M13 or C-bit parity-mapped configurations of *N*xDS0 IQ interfaces.

Figure 45: Sample Channelization of OC12 IQE PIC



Bold entries correspond to actual packet channels.

g003014

Example: Configuring NxDS0 Interfaces

Configure the following two NxDS0 interfaces with 10 time slots and 4 time slots, respectively:

VT-Mapped Configuration	ds-0/0/0:1:2:1 ds-0/0/0:1:2:2
	[edit interfaces coc12-0/0/0] partition 1 oc-slice 1 interface-type coc1;
	[edit interfaces coc1-0/0/0:1] partition 2 interface-type ct1;
M13 or C-bit Parity-Mapped Configuration	[edit interfaces ct1-0/0/0:1:2] partition 1 timeslots 1-10 interface-type ds; partition 2 timeslots 12-16 interface-type ds;
	[edit interfaces coc12-0/0/0] partition 1 oc-slice 1 interface-type coc1;
	[edit interfaces coc1-0/0/0:1] no-partition interface-type ct3;
	[edit interfaces ct3-0/0/0:1] partition 2 interface-type ct1;
	[edit interfaces ct1-0/0/0:1:2] partition 1 timeslots 1-10 interface-type ds; partition 2 timeslots 12-16 interface-type ds;

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring Fractional T1 Interfaces

By default, all the time slots on a channelized T1 interface are used. To configure a fractional T1 interface on a Channelized OC12 IQE PIC, perform the following tasks:

1. Configure a T1 interface. For more information, see “Configuring T1 Interfaces” on page 410.
2. Configure the number of time slots allocated to the T1 interface by including the `timeslots` statement at the `[edit interfaces t1-fpc/pic/port<:channel> t1-options]` hierarchy level:

```
[edit interfaces t1-fpc/pic/port<:channel> t1-options]
timeslots time-slot-range;
```

For channelized T1 interfaces, the time-slot range is from 1 through 24. You can designate any combination of time slots. To configure ranges, use hyphens. To

configure discontinuous time slots, use commas. Do not include spaces. For more information, see “Configuring Fractional T1 Time Slots” on page 567.

Example: Configuring Fractional T1 Interfaces

Configure a fractional T1 interface that uses time slots 1 through 5 and 10:

```
[edit interfaces coc12-0/0/0]
partition 1 oc-slice 1 interface-type coc1;
[edit interfaces coc1-0/0/0:1]
partition 1 interface-type t1;
[edit interfaces t1-0/0/0:1:1 t1-options]
timeslots 1-5,10;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring Channelized OC12/STM4 IQE Interfaces (SDH Mode)

The Channelized OC12 IQE PIC configured for SDH mode creates a single channelized STM4 interface. You can configure this interface as unpartitioned using the **no-partition** statement at the [edit interfaces cstm4-fpc/pic/port] hierarchy level to create a single SDH VC-4-4C interface, or you can partition it into the following OC slices:

- SDH virtual concatenation 4 (VC-4) and channelized AU-4 interfaces (4 interfaces, any combination)
- E3 interfaces from a channelized AU-4 interface (3 interfaces, any combination)
- Channelized E1 or E1 interfaces from a channelized AU-4 interface (63 interfaces, any combination)
- NxDS0 interfaces from a channelized E1 interface

This section describes how to configure the following channelized OC12 IQE interfaces on a Channelized OC12 IQE PIC configured in SDH mode:

- Configuring Channelized OC12/STM4 IQE PICs for SDH Mode on page 434
- Configuring an Unpartitioned SDH (VC-4-4C) Interface on a Channelized OC12/STM4 IQE PIC on page 435
- Configuring SDH (VC-4) Interfaces on Channelized OC12/STM4 IQE PICs on page 436
- Configuring Channelized AU-4 Interfaces on page 436
- Configuring E3 Interfaces on page 437
- Configuring E1 or Channelized E1 Interfaces on page 438
- Configuring NxDS0 Interfaces on Channelized OC12/STM4 IQE PICs on page 439

Configuring Channelized OC12/STM4 IQE PICs for SDH Mode

The 4-port Channelized OC12 IQE PIC allows SONET/SDH configuration on a per port basis, permitting combinations of SONET and SDH ports on the same PIC. The 1-port Channelized OC12 IQE PIC operates in either SONET or SDH mode only.

To configure a 1-port Channelized OC12 IQE PIC to operate in SDH mode, include the `framing sdh` statement at the `[edit chassis fpc fpc/pic/port]` hierarchy level:

```
[edit chassis]
fpc 0 {
  pic 2 {
    framing sdh;
  }
}
```

This configuration creates interface `cstm4-0/2/0`.

You can also use the above configuration example to configure all 4 ports of a 4-port Channelized OC12 IQE PIC for SDH mode. To configure individual ports to operate in SDH mode, include the `framing sdh` statement at the `[edit chassis fpc fpc/pic/port]` hierarchy level. The following example configures port 2 for SDH mode:

```
[edit chassis]
fpc 0 {
  pic 2 {
    port 2 {
      framing sdh;
    }
  }
}
```

This configuration creates interface `cstm4-0/2/2`.

For more information, see the *JUNOS System Basics Configuration Guide*.

Configuring an Unpartitioned SDH (VC-4-4C) Interface on a Channelized OC12/STM4 IQE PIC

On a Channelized OC12 IQE PIC, you can configure one SDH (VC-4-4C) interface. To configure an SDH (VC-4-4C) interface, include the `no-partition` and `interface-type` statements at the `[edit interfaces cstm4-fpc/pic/port]` hierarchy level:

```
[edit interfaces cstm4-fpc/pic/port]
no-partition interface-type so;
```

This configuration creates interface `so-fpc/pic/port`.

Example: Configuring an Unpartitioned SDH (VC-4-4C) Interface

Configure an unpartitioned SDH (VC-4-4C) interface, using partition 1 and OC slices 4 through 6:

```
[edit interfaces cstm4-0/2/0]
no-partition interface-type so;
```

This configuration creates the interface `so-0/2/0`.

Configuring SDH (VC-4) Interfaces on Channelized OC12/STM4 IQE PICs

To configure an SDH (VC-4) interface on a Channelized OC12 IQE PIC, include the `partition`, `oc-slice`, and `interface-type` statements at the `[edit interfaces cstm4-fpc/pic/port]` hierarchy level, specifying the `so` interface type:

```
[edit interfaces cstm4-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type so;
```

This configuration creates interface `so-fpc/pic/port:channel`.

The partition number is the sublevel interface partition index and is correlated with the channel number. For Channelized OC12 IQE PICs, the OC-slice range can be from 1 through 12.



NOTE: For channelized OC12 IQE interfaces, channel numbering begins with 1 (:1).

The OC-slice range is the range of SONET/SDH slices. For SDH interfaces, the OC-slice range specifies the bandwidth size required for the interface type you are configuring. SDH (VC-4) interfaces must occupy three consecutive OC slices per interface, in one of the following forms:

- 1–3
- 4–6
- 7–9
- 10–12

The interface type is the channelized interface type or data channel you are creating.

Example: Configuring SDH (VC-4) Interfaces

Configure SDH (VC-4) interfaces:

```
[edit interfaces cstm4-0/2/0]
partition 1 oc-slice 1-3 interface-type so;
partition 2 oc-slice 4-6 interface-type so;
partition 3 oc-slice 7-9 interface-type so;
partition 4 oc-slice 10-12 interface-type so;
```

This configuration creates the interfaces `so-0/2/0:1` through `so-0/2/0:4`.

Configuring Channelized AU-4 Interfaces

To configure a channelized AU-4 interface, include the `partition`, `oc-slice`, and `interface-type` statements at the `[edit interfaces cstm4-fpc/pic/port]` hierarchy level, specifying the `cau4` interface type:

```
[edit interfaces cstm4-fpc/pic/port]
```

```
partition partition-number oc-slice oc-slice-range interface-type cau4;
```

This configuration creates interface *cau4-fpc/pic/port:channel*.

The partition number is the sublevel interface partition index. For SDH interfaces, the partition number is not correlated with bandwidth size. A channelized STM-4 interface can have from 1 through 4 partition numbers.



NOTE: For channelized OC12 interfaces, channel numbering begins with 0 (:0). For channelized OC12 interfaces (both IQ and IQE), channel numbering begins with 1 (:1).

The OC-slice range is the range of SONET/SDH slices. For SDH interfaces, the OC-slice range specifies the bandwidth size required for the interface type you are configuring. Channelized AU-4 IQ interfaces must occupy three consecutive OC slices per interface, in one of the following forms:

- 1–3
- 4–6
- 7–9
- 10–12

The interface type is the channelized interface type or data channel you are creating.

Example: Configuring Channelized AU-4 Interfaces

Configure channelized AU-4 interfaces, using partitions 1 through 4:

```
[edit interfaces cstm4-0/2/0]
partition 1 oc-slice 1-3 interface-type cau4;
partition 2 oc-slice 4-6 interface-type cau4;
partition 3 oc-slice 7-9 interface-type cau4;
partition 4 oc-slice 10-12 interface-type cau4;
```

This configuration creates the interfaces *cau4-0/2/0:1* through *cau4-0/2/0:4*.

Configuring E3 Interfaces

To configure E3 interfaces, include the **partition** and **interface-type** statements at the `[edit interfaces cau4-fpc/pic/port]` hierarchy level, specifying the **e3** interface type:

```
[edit interfaces]
cau4-fpc/pic/port {
  partition partition-number interface-type e3;
}
```

This configuration creates the interfaces *e3-fpc/pic/port:channel* and *e3-fpc/pic/port:channel*.



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQE interfaces. You can only apply CoS rules to the aggregate bit streams.



NOTE: Channelized OC12/STM4 IQ and IQE interfaces in M Series, MX Series, and T Series routers reserve channels 0-3 of each OC12 space for STS3C SONET channels.

When you configure E3 or T3 channels in OC12 spaces on the described PICs, JUNOS Software allocates them starting from channel 4 because channels 0-3 are reserved for four STS3c SONET channels. Channel numbers are allocated sequentially in the following order: 4, 5, 6, 7, 8, 9, 11, 0, 1, 2, 3.

Only after channels 4 through 11 of the OC12 space are exhausted (all 4 through 11 configured) for E3 or T3 channels will JUNOS Software then allocate channel 0-3 space for further E3 or T3 channels; thereby using up the 0-3 space previously reserved for four STS3c SONET channels.

If a subsequent reconfiguration of this OC12 space occurs, where you try to replace channels 4-6 or 7-9 with an OC3 SONET channel; it fails because the channel 0-3 space is already occupied by the last E3 or T3 channels configured. This causes a failure in channel allocation and the Device Control Daemon (DCD) keeps retrying forever to configure the channel allocation on the interface. The only resolution is to reconfigure the last configured E3/T3 channels with OC3 channels, to free channels 0-3.

Example: Configuring E3 Interfaces

Configure E3 interfaces, using partition 1:

```
[edit interfaces]
cau4-0/2/0:1 {
  partition 1 interface-type e3;
}
e3-0/2/0:1:1;
```

Configuring E1 or Channelized E1 Interfaces

To configure E1 or channelized E1 interfaces, include the **partition** and **interface-type** statements at the **[edit interfaces cau4-fpc/pic/port]** hierarchy level, specifying the **e1** or **ce1** interface type:

```
[edit interfaces]
cau4-fpc/pic/port {
  partition partition-number interface-type e1;
}
cau4-fpc/pic/port {
  partition partition-number interface-type ce1;
}
```

This configuration creates the interfaces `e1-fpc/pic/port:channel` and `ce1-fpc/pic/port:channel`.



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQE interfaces. You can only apply CoS rules to the aggregate bit streams.

Example: Configuring E1 or Channelized CE1 Interfaces

Configure E1 or channelized CE1 interfaces, using partition 3 and partition 4:

```
[edit interfaces]
cau4-0/2/0:1 {
  partition 3 interface-type e1;
}
cau4-0/2/0:1 {
  partition 4 interface-type ce1;
}
```

This configuration creates interfaces `e1-0/2/0:1:3` and `ce1-0/2/0:1:4`.

Configuring NxDS0 Interfaces on Channelized OC12/STM4 IQE PICs

Configure channelized NxDS0 interfaces on the channelized E1 interface by including the `partition`, `timeslots`, and `interface-type` statements at the `[edit interfaces ce1-fpc/pic/port:channel]` hierarchy level, specifying the `ds` interface type:

```
[edit interfaces ce1-fpc/pic/port:channel:channel]
partition partition-number timeslots time-slot-range interface-type ds;
```

This configuration creates the interface `ds-fpc/pic/port:channel`.

The time-slot range is from 1 through 32. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. You can use a combination of ranges and discontinuous time slots, for example:

```
1,9-18,21
```

Example: Configuring NxDS0 Interfaces

Configure channelized NxDS0 interfaces, using partition 4 and time slots 1 through 10:

```
[edit interfaces]
ce1-0/2/0:1:2:3 {
  partition 4 interface-type ds0 timeslots 1-10;
}
```

This configuration creates interface `ds-0/2/0:1:2:4`.

Configuring Channelized OC12/STM4 IQ Interfaces (SDH Mode)

The Channelized OC12 IQ PIC configured for SDH mode creates a single channelized STM4 interface. You can configure this interface as unpartitioned using the **no-partition** statement at the `[edit interfaces cstm4-fpc/pic/port]` hierarchy level to create a single SDH VC-4-4C interface, or you can partition it into the following OC slices:

- SDH virtual concatenation 4 (VC-4) and channelized AU-4 interfaces (4 interfaces, any combination)
- Channelized T3 or T3 interfaces from a channelized AU-4 interface (3 interfaces, any combination)
- Channelized T1 or T1 interfaces from a channelized T3 interface (28 interfaces, any combination)
- NxDS0 interfaces from a channelized T1 interface

This section describes how to configure the following channelized OC12 IQ interfaces on a Channelized OC12 IQ PIC configured in SDH mode:

- Configuring Channelized OC12/STM4 IQ PICs for SDH Mode on page 440
- Configuring an Unpartitioned SDH (VC-4-4C) Interface on a Channelized OC12/STM4 IQ PIC on page 441
- Configuring SDH (VC-4) Interfaces on Channelized OC12/STM4 IQ PICs on page 441
- Configuring Channelized AU-4 Interfaces on page 442
- Configuring T3 or Channelized T3 Interfaces on page 443
- Configuring T1 or Channelized T1 Interfaces on page 443
- Configuring NxDS0 Interfaces on Channelized OC12/STM4 IQ PICs on page 444

Configuring Channelized OC12/STM4 IQ PICs for SDH Mode

To configure a Channelized OC12 IQ PIC to operate in SDH mode, include the **framing sdh** statement at the `[edit chassis fpc fpc/pic/port]` hierarchy level:

```
[edit chassis]
fpc 0 {
  pic 2 {
    framing sdh;
  }
}
```

This configuration creates interface `cstm4-0/2/0`.

For more information, see the *JUNOS System Basics Configuration Guide*.

Configuring an Unpartitioned SDH (VC-4-4C) Interface on a Channelized OC12/STM4 IQ PIC

On a Channelized OC12 IQ PIC, you can configure one SDH (VC-4-4C) interface. To configure an SDH (VC-4-4C) interface, include the `no-partition` and `interface-type` statements at the `[edit interfaces cstm4-fpc/pic/port]` hierarchy level:

```
[edit interfaces cstm4-fpc/pic/port]
no-partition interface-type so;
```

This configuration creates interface `so-fpc/pic/port`.

Example: Configuring an Unpartitioned SDH (VC-4-4C) Interface

Configure an unpartitioned SDH (VC-4-4C) interface, using partition 1 and OC slices 4 through 6:

```
[edit interfaces cstm4-0/2/0]
no-partition interface-type so;
```

This configuration creates the interface `so-0/2/0`.

Configuring SDH (VC-4) Interfaces on Channelized OC12/STM4 IQ PICs

To configure an SDH (VC-4) interface on a Channelized OC12 IQ PIC, include the `partition`, `oc-slice`, and `interface-type` statements at the `[edit interfaces cstm4-fpc/pic/port]` hierarchy level, specifying the `so` interface type:

```
[edit interfaces cstm4-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type so;
```

This configuration creates interface `so-fpc/pic/port:channel`.

The partition number is the sublevel interface partition index and is correlated with the channel number. For Channelized OC12 IQ PICs, the OC-slice range can be from 1 through 12.



NOTE: For channelized OC12 IQ interfaces, channel numbering begins with 1 (:1).

The OC-slice range is the range of SONET/SDH slices. For SDH interfaces, the OC-slice range specifies the bandwidth size required for the interface type you are configuring. SDH (VC-4) interfaces must occupy three consecutive OC-slices per interface, in one of the following forms:

- 1–3
- 4–6
- 7–9
- 10–12

The interface type is the channelized interface type or data channel you are creating.

Example: Configuring SDH (VC-4) Interfaces

Configure SDH (VC-4) interfaces:

```
[edit interfaces cstm4-0/2/0]
partition 1 oc-slice 1-3 interface-type so;
partition 2 oc-slice 4-6 interface-type so;
partition 3 oc-slice 7-9 interface-type so;
partition 4 oc-slice 10-12 interface-type so;
```

This configuration creates the interfaces `so-0/2/0:1` through `so-0/2/0:4`.

Configuring Channelized AU-4 Interfaces

To configure a channelized AU-4 interface, include the `partition`, `oc-slice`, and `interface-type` statements at the `[edit interfaces cstm4-fpc/pic/port]` hierarchy level, specifying the `cau4` interface type:

```
[edit interfaces cstm4-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type cau4;
```

This configuration creates interface `cau4-fpc/pic/port:channel`.

The partition number is the sublevel interface partition index. For SDH interfaces, the partition number is not correlated with bandwidth size. A channelized STM-4 interface can have from 1 through 4 partition numbers.



NOTE: For channelized OC12 interfaces, channel numbering begins with 0 (:0). For channelized OC12 interfaces (both IQ and IQE), channel numbering begins with 1 (:1).

The OC-slice range is the range of SONET/SDH slices. For SDH interfaces, the OC-slice range specifies the bandwidth size required for the interface type you are configuring. Channelized AU-4 IQ interfaces must occupy three consecutive OC slices per interface, in one of the following forms:

- 1–3
- 4–6
- 7–9
- 10–12

The interface type is the channelized interface type or data channel you are creating.

Example: Configuring Channelized AU-4 Interfaces

Configure channelized AU-4 interfaces, using partitions 1 through 4:

```
[edit interfaces cstm4-0/2/0]
partition 1 oc-slice 1-3 interface-type cau4;
partition 2 oc-slice 4-6 interface-type cau4;
partition 3 oc-slice 7-9 interface-type cau4;
partition 4 oc-slice 10-12 interface-type cau4;
```

This configuration creates the interfaces `cau4-0/2/0:1` through `cau4-0/2/0:4`.

Configuring T3 or Channelized T3 Interfaces

To configure T3 or channelized T3 interfaces, include the `partition` and `interface-type` statements at the `[edit interfaces cau4-fpc/pic/port]` hierarchy level, specifying the `t3` or `ct3` interface type:

```
[edit interfaces]
cau4-fpc/pic/port {
  partition partition-number interface-type t3;
}
cau4-fpc/pic/port {
  partition partition-number interface-type ct3;
}
```

This configuration creates the interfaces `t3-fpc/pic/port:channel` and `t3-fpc/pic/port:channel`.



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQ interfaces. You can only apply CoS rules to the aggregate bit streams.

Example: Configuring T3 or Channelized T3 Interfaces

Configure T3 and channelized T3 interfaces, using partition 1 and partition 2:

```
[edit interfaces]
cau4-0/2/0:1 {
  partition 1 interface-type t3;
}
cau4-0/2/0:1 {
  partition 2 interface-type ct3;
}
t3-0/2/0:1:1 ct3-0/2/0:1:2;
```

Configuring T1 or Channelized T1 Interfaces

To configure T1 or channelized T1 interfaces, include the `partition` and `interface-type` statements at the `[edit interfaces ct3-fpc/pic/port]` hierarchy level, specifying the `t1` or `ct1` interface type:

```
[edit interfaces]
ct3-fpc/pic/port {
  partition partition-number interface-type t1;
```

```

}
ct3-fpc/pic/port {
    partition partition-number interface-type ct1;
}

```

This configuration creates the interfaces *t1-fpc/pic/port:channel* and *ct1-fpc/pic/port:channel*.



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQ interfaces. You can only apply CoS rules to the aggregate bit streams.

Example: Configuring T1 or Channelized T1 Interfaces

Configure T1 or channelized T1 interfaces, using partition 3 and partition 4:

```

[edit interfaces]
ct3-0/2/0:1:2 {
    partition 3 interface-type t1;
}
ct3-0/2/0:1:2 {
    partition 4 interface-type ct1;
}

```

This configuration creates interfaces *t1-0/2/0:1:2:3* and *ct1-0/2/0:1:2:4*.

Configuring NxDS0 Interfaces on Channelized OC12/STM4 IQ PICs

Configure channelized NxDS0 IQ interfaces on the channelized T1 IQ interface by including the *partition*, *timeslots*, and *interface-type* statements at the [edit interfaces *ct1-fpc/pic/port:channel*] hierarchy level, specifying the *ds* interface type:

```

[edit interfaces ct1-fpc/pic/port:channel:channel]
partition partition-number timeslots time-slot-range interface-type ds;

```

This configuration creates the interface *ds-fpc/pic/port:channel*.

The time-slot range is from 1 through 24. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. You can use a combination of ranges and discontinuous time slots:

```
1,9-18,21
```

Example: Configuring NxDS0 Interfaces

Configure channelized NxDS0 interfaces, using partition 4 and time slots 1 through 10:

```

[edit interfaces]
ct1-0/2/0:1:2:3 {
    partition 4 interface-type ds0 timeslots 1-10;
}

```

```
}

```

This configuration creates interface `ds-0/2/0:1:2:3:4`.

Configuring Channelized OC12 Interfaces

On Channelized OC12 PICs, you can configure 12 T3 channels per port. To configure channelized OC12 interface properties, you can include the `sonet-options` and `t3-options` statements at the `[edit interfaces interface-name]` hierarchy level. Some SONET/SDH options are ignored, and some can only be configured for channel 0, though they apply equally to all channels. The `long-buildout` statement under `t3-options` is also ignored.

For T3 channels on a channelized OC12 interface, the `clocking` statement is supported only for channel 0; it is ignored if included in the configuration of channels 1 through 11. The clock source configured for channel 0 applies to all channels on the channelized OC12 interface. The individual T3 channels use a gapped 45-MHz clock as the transmit clock. When you configure the clock source for a channelized interface—`ds-fpc/pic/port :0`, for example—you must also include the `channel-group` statement at the `[edit chassis]` hierarchy level and specify channel group 0. For more information, see “Clock Sources on Channelized Interfaces” on page 390.

For more information, see “Configuring SONET/SDH Interfaces” on page 843 and “Configuring T3 Interfaces” on page 569. For a configuration example, see “Configuring Aggregated SONET/SDH Interfaces” on page 881.

Table 41 on page 445 summarizes the OC12-to-DS3 numbering scheme.

Table 41: OC12-to-DS3 Numbering Scheme

Two-Level STS-1 Number (STS-3,STS-1)	One-Level STS Number	OC12-to-DS3 PIC DS3 Number
1,1	1	0
1,2	2	1
1,3	3	2
2,1	4	3
2,2	5	4
2,3	6	5
3,1	7	6
3,2	8	7
3,3	9	8
4,1	10	9
4,2	11	10

Table 41: OC12-to-DS3 Numbering Scheme (continued)

Two-Level STS-1 Number (STS-3,STS-1)	One-Level STS Number	OC12-to-DS3 PIC DS3 Number
4,3	12	11

Example: Configuring Channelized OC12 Interfaces

The following configuration is sufficient to get the channelized OC12 interface up and running. The OC12 interface can be divided into 12 channels. DS3 channels can use the following encapsulation types:

- PPP, PPP CCC, and PPP TCC
- Frame Relay, Frame Relay CCC, and Frame Relay TCC
- Cisco HDLC, Cisco HDLC CCC, and Cisco HDLC TCC

The channels can also have logical interfaces.

```
[edit interfaces]
t3-fpc/pic/port:0 {
  encapsulation cisco-hdlc;
  t3-options {
    compatibility-mode larscom;
    payload-scrambler;
  }
  unit 0 {
    family inet {
      address 10.11.30.1/30;
    }
    family iso;
  }
}
t3-fpc/pic/port:1 {
  encapsulation ppp;
  t3-options {
    compatibility-mode larscom;
    payload-scrambler;
  }
  unit 0 {
    family inet {
      address 10.11.30.5/30;
    }
    family iso;
  }
}
t3-fpc/pic/port:2 {
  encapsulation frame-relay;
  t3-options {
    compatibility-mode larscom;
    payload-scrambler;
  }
  unit 0 {
```

```

        dlci 100;
        family inet {
            address 10.11.30.9/30;
        }
        family iso;
    }
    unit 1 {
        dlci 101;
        family inet {
            address 10.11.31.9/30;
        }
        family iso;
    }
}
t3-1fpc/pic/port:3 {
    encapsulation cisco-hdlc-ccc;
    t3-options {
        compatibility-mode larscom;
        payload-scrambler;
    }
    unit 0;
}
t3-fpc/pic/port:4 {
    encapsulation ppp-ccc;
    t3-options {
        compatibility-mode larscom;
        payload-scrambler;
    }
    unit 0;
}
t3-fpc/pic/port:5 {
    dce;
    encapsulation frame-relay-ccc;
    t3-options {
        compatibility-mode larscom;
        payload-scrambler;
    }
    unit 0 {
        encapsulation frame-relay-ccc;
        dlci 1000;
    }
    unit 1 {
        encapsulation frame-relay-ccc;
        dlci 1001;
    }
}
}

```

Configuring Link PIC Failover on Channelized OC12/STM4 IQ and IQE Interfaces

For Channelized OC12 IQ or IQE PICs used as linking PICs in redundant LSQ configurations, you can inhibit the router from sending PPP termination-request messages to the remote host if the link PIC fails. To do this, include the **no-termination-request** statement at the [edit interfaces *interface-name* ppp-options] hierarchy level:

```
no-termination-request;
```

The `no-termination-request` statement is supported only with MLPPP and SONET APS configurations and works with PPP, PPP over Frame Relay, and MLPPP interfaces only.

For information about interchassis and intrachassis LSQ failover, see the *JUNOS Services Interfaces Configuration Guide*.

Example: Configuring a Channelized OC12 IQ Interface as an Unpartitioned, Clear Channel

Configuring a SONET/SDH Interface	<p>Configure a channelized OC12 interface as an unpartitioned, clear channel:</p> <pre>[edit interfaces] coc12-5/0/0 { no-partition interface-type so; # so-5/0/0 }</pre>
Configuring Multiple Interface Types	<p>Configure the following interfaces on a Channelized OC12 IQ or IQE PIC:</p> <ol style="list-style-type: none"> An OC3 interface Another OC3 interface A channelized OC1 partitioned into T1 interfaces A channelized OC1 converted into a T3 interface A channelized OC1 partitioned into T1 interfaces and channelized T1s, which are partitioned into NxDS0 interfaces A channelized OC1 converted into a channelized T3, which is partitioned into T1 interfaces A channelized OC1 converted into a channelized T3, which is partitioned into T1 interfaces and a channelized T1, which is partitioned into NxDS0 interfaces A channelized OC1 partitioned into channelized T1s, which are partitioned into NxDS0 interfaces
Configuring the Interface Partitions	<pre>[edit interfaces] coc12-1/1/0 { sonet-options { sonet-options-statements; } partition 1 oc-slice 1-3 interface-type so; # (a) so-1/1/0:1 partition 2 oc-slice 4-6 interface-type so; # (b) so-1/1/0:2 partition 3 oc-slice 7 interface-type coc1; # (c) coc1-1/1/0:3 partition 4 oc-slice 8 interface-type coc1; # (d) coc1-1/1/0:5 partition 5 oc-slice 9 interface-type coc1; # (e) coc1-1/1/0:5 partition 6 oc-slice 10 interface-type coc1; # (f) coc1-1/1/0:6 partition 7 oc-slice 11 interface-type coc1; # (g) coc1-1/1/0:7 partition 8 oc-slice 12 interface-type coc1; # (h) coc1-1/1/0:8 }</pre> <p>(a) <code>so-1/1/0:1 {</code></p>


```

        description "(a) OC-slice 1-3 of coc12-1/1/0. COC12 > OC3.;
        sonet-options {
            sonet-options-statements;
        }
    }

(b) so-1/1/0:2 {
    description "(b) OC-slice 4-6 of coc12-1/1/0. COC12 > OC3.;
    sonet-options {
        sonet-options-statements;
    }
}

(c) coc1-1/1/0:3 {
    description "(c) OC-slice 7 of coc12-1/1/0. COC12 to COC1 VT-mapped to T1s.";
    sonet-options {
        sonet-options-statements;
    }
    partition 1 - 10 interface-type t1; # t1-1/1/0:[1-10]
}
t1-1/1/0:3:1 {
    description "(c) OC-slice 7 of coc12-1/1/0. T1 interface configuration.";
    t1-options {
        t1-options-statements;
    }
}
...

(d) coc1-1/1/0:4 {
    description "(d) OC-slice 8 of coc12-1/1/0. COC12 to COC1 converted to a T3.";
    sonet-options {
        sonet-options-statements;
    }
    no-partition interface-type t3; # t3-1/1/0:4
}
t3-1/1/0:4 {
    description "(d) OC-slice 8 of coc12-1/1/0. T3 interface configuration.";
}

(e) coc1-1/1/0:5 {
    description "(e) OC-slice 9 of coc12-1/1/0. COC12 to COC1 VT-mapped to T1s.";
    sonet-options {
        sonet-options-statements;
    }
    partition 1 - 3 interface-type t1; # t1-1/1/0:5:[1-3]
    partition 4 interface-type ct1; # ct1-1/1/0:5:4
}
t1-1/1/0:5:1 {
    description "(e) OC-slice 9 of coc12-1/1/0. T1 interface configuration.";
    t1-options {
        t1-options-statements;
    }
}
...

```

```

ct1-1/1/0:5:4 {
    description "(e) OC-slice 9 of coc12-1/1/0. CT1 to NxDSOs.;
    t1-options {
        t1-options-statements;
    }
    partition 1 timeslots 0 - 10 interface-type ds0; # ds-1/1/0:5:4:1
    partition 2 timeslots 11- 23 interface-type ds0; # ds-1/1/0:5:4:2
    ...
}

(f) coc1-1/1/0:6 {
    description "(f) OC-slice 10 of coc12-1/1/0. COC12 to COC1 converted to a CT3
        to T1s.";
    sonet-options {
        sonet-options-statements;
    }
    no-partition interface-type ct3; # ct3-1/1/0:6
}
ct3-1/1/0:6 {
    description "(f) COC12 to CT3 M-13 and C-bit parity-mapped to T1s.;
    sonet-options {
        sonet-options-statements;
    }
    partition 1 - 10 interface-type t1; # t1-1/1/0:6:[1-10]
}
t1-1/1/0:6:1 {
    description "(f) T1 interface configuration.";
    t1-options {
        t1-options-statements;
    }
}
...

(g) coc1-1/1/0:7 {
    description "(g) OC-slice 11 of coc12-1/1/0. COC12 to COC1 converted to a CT3
        to T1s and CT1 to NxDSOs.";
    sonet-options {
        sonet-options-statements;
    }
    no-partition interface-type ct3; # ct3-1/1/0:7
}
ct3-1/1/0:7 {
    description "(g) COC12 to CT3 M-13 and C-bit parity-mapped to T1s and CT1.";
    sonet-options {
        sonet-options-statements;
    }
    partition 1 - 10 interface-type t1; # t1-1/1/0:7:[1-10]
    partition 2 interface-type ct1; # ct1-1/1/0:7:11
}
t1-1/1/0:7:1 {
    description "(g) T1 interface configuration.";
    t1-options {
        t1-options-statements;
    }
}

```

```

...
ct1-1/1/0:7:11 {
    description "(g) CT1 to NxDSOs.";
    t1-options {
        t1-options-statements;
    }
    partition 1 timeslots 0 - 10 interface-type ds0; # ds-1/1/0:7:11:1
    partition 2 timeslots 11- 23 interface-type ds0; # ds-1/1/0:7:11:2
    ...
}

(h) coc1-1/1/0:8 {
    description "(h) OC-slice 12 of coc12-1/1/0. COC12 to COC1 VT-mapped to CT1
        to NxDSOs.";
    sonet-options {
        sonet-options-statements;
    }
    partition 1 interface-type t1; # ct1-1/1/0:8:1
}
ct1-1/1/0:8:1 {
    description "(h) CT1 to NxDSOs.";
    t1-options {
        t1-options-statements;
    }
    partition 1 timeslots 0 - 10 interface-type ds0; # ds-1/1/0:8:1:1
    partition 2 timeslots 11- 23 interface-type ds0; # ds-1/1/0:8:1:2
    ...
}

```

For a full configuration example, see the *JUNOS Feature Guide*.

Example: Configuring Channelized OC12 Interfaces with Partitioned Channels

The following configuration is sufficient to get the channelized OC12 interface up and running. The OC12 interface can be divided into 12 channels. DS3 channels can use the following encapsulation types:

- PPP, PPP CCC, and PPP TCC
- Frame Relay, Frame Relay CCC, and Frame Relay TCC
- Cisco HDLC, Cisco HDLC CCC, and Cisco HDLC TCC

The channels can also have logical interfaces.

```

[edit interfaces]
t3-fpc/pic/port:0 {
    encapsulation cisco-hdlc;
    t3-options {
        compatibility-mode larscom;
        payload-scrambler;
    }
    unit 0 {
        family inet {

```

```

        address 10.11.30.1/30;
    }
    family iso;
}
}
t3-fpc/pic/port:1 {
    encapsulation ppp;
    t3-options {
        compatibility-mode larscom;
        payload-scrambler;
    }
    unit 0 {
        family inet {
            address 10.11.30.5/30;
        }
        family iso;
    }
}
t3-fpc/pic/port:2 {
    encapsulation frame-relay;
    t3-options {
        compatibility-mode larscom;
        payload-scrambler;
    }
    unit 0 {
        dlci 100;
        family inet {
            address 10.11.30.9/30;
        }
        family iso;
    }
    unit 1 {
        dlci 101;
        family inet {
            address 10.11.31.9/30;
        }
        family iso;
    }
}
t3-1fpc/pic/port:3 {
    encapsulation cisco-hdlc-ccc;
    t3-options {
        compatibility-mode larscom;
        payload-scrambler;
    }
    unit 0;
}
t3-fpc/pic/port:4 {
    encapsulation ppp-ccc;
    t3-options {
        compatibility-mode larscom;
        payload-scrambler;
    }
    unit 0;
}
t3-fpc/pic/port:5 {

```

```
dce;
encapsulation frame-relay-ccc;
t3-options {
    compatibility-mode larscom;
    payload-scrambler;
}
unit 0 {
    encapsulation frame-relay-ccc;
    dlci 1000;
}
unit 1 {
    encapsulation frame-relay-ccc;
    dlci 1001;
}
}
```


Chapter 20

Configuring Channelized OC3 IQ and IQE Interfaces

This chapter describes how to configure interfaces on Channelized OC3 IQ and IQE PICs, as follows:

- Channelized OC3 IQ and IQE Overview on page 455
- Partitions, OC Slices, Interface Types, and Time Slots on page 456
- Configuring a Clear Channel on Channelized OC3 IQ and IQE PICs on page 457
- Configuring T3 IQ Interfaces on page 457
- Configuring T1 and NxDS0 Interfaces on page 458
- Configuring Fractional T1 IQ Interfaces on page 462
- Configuring Link PIC Failover on Channelized OC3 IQ and IQE Interfaces on page 462

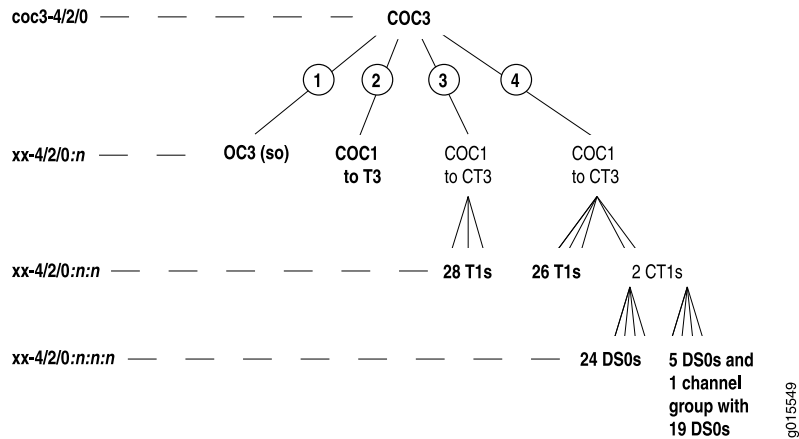
Channelized OC3 IQ and IQE Overview

Channelized intelligent queuing (IQ) and channelized enhanced intelligent queuing (IQE) interfaces allow arbitrary and dynamic channelization of serial links, allowing greater flexibility than regular channelized interfaces.

On each port of a Channelized OC3 IQ or a Channelized OC3 IQE interface, you can configure the following interface types:

- One OC3 SONET interface
- Up to three T3 interfaces
- Up to 84 T1 interfaces
- Up to three E3 interfaces (COC3 IQE PICs in SDH mode)
- Up to 63 E1 interfaces (COC3 IQE PICs in SDH mode)
- Up to 336 NxDS0 interfaces on an M Series router
- Up to 768 NxDS0 interfaces on a T Series router

Figure 46 on page 456 shows an example of how a Channelized OC3 PIC might be partitioned. In the figure, the OC3 SONET interface would be a standalone interface because it would use the entire bandwidth of the PIC. The same applies to each port of the 2-port Channelized OC3 Enhanced IQ (IQE) PIC.

Figure 46: Channelized OC3 IQ Interface Example for Show Interfaces Controller

You can configure the following encapsulation types:

- PPP
- Frame Relay
- Cisco HDLC
- CCC
- TCC
- MPLS—On IQE interfaces.

For more information about interface encapsulation, see “Configuring Interface Encapsulation on Physical Interfaces” on page 106 and “Configuring Interface Encapsulation on Logical Interfaces” on page 160.

To configure channelized interfaces, include the following statements at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
no-partition interface-type type;
partition partition-number oc-slice oc-slice-range interface-type type;
partition partition-number timeslots time-slot-range interface-type type;
```

Partitions, OC Slices, Interface Types, and Time Slots

The partition number is the sublevel interface partition index and is correlated with the channel number. For channelized OC3 interfaces, you can configure up to three OC1 interfaces, so the partition number can be 1, 2, or 3. For channelized T3 interfaces (ct3), you can configure multiple interfaces at once by including a partition range, such as 1-3. This creates three T1 interfaces with channel numbers 1 through 3.



NOTE: For channelized IQ and IQE interfaces, channel numbering begins with 1 (:1). For regular channelized interfaces, channel numbering begins with 0 (:0).

You configure the OC-slice range for SONET/SDH interfaces only. The OC-slice range is correlated with the bandwidth size required for the interface type you are configuring. For example, a channelized OC3 interface (**coc3**) can be divided into three OC1 interfaces, each containing one OC slice. Therefore the OC-slice value must be 1, 2, or 3.

The configurable interface types are dependent on the hierarchy level at which you include the **interface-type** and **partition** or **no-partition** statements. For example, when you include the **no-partition** statement at the [edit interfaces **coc3-fpc/pic/port**] hierarchy level, the only configurable interface type is **so**, because the **no-partition** statement signals that you are creating a clear-channel SONET/SDH interface. When you include the **partition** statement at the [edit interfaces **coc1-fpc/pic/port**] hierarchy level, the configurable interface types are **ct1** or **t1**. If you want to create a T1 interface, include the **t1** option. If you want to further channelize down to the NxDS0 level, include the **ct1** option as an intermediate step before dividing the channelized T1 interface (**ct1**) into NxDS0 interfaces.

You configure time slots for fractional T1 interfaces and NxDS0 interfaces. You can configure ranges by using hyphens. You can configure discontinuous time slots by using commas. Do not include spaces.

Configuring a Clear Channel on Channelized OC3 IQ and IQE PICs

A *clear channel* is an interface that uses the entire bandwidth of the PIC. To configure a clear channel, include the **no-partition** and **interface-type** statements in the configuration.

On Channelized OC3 IQ and IQE PICs, you can configure one OC3 clear-channel interface per port. To configure an OC3 interface, include the **no-partition** and **interface-type** statements at the [edit interfaces **coc3-fpc/pic/port**] hierarchy level:

```
[edit interfaces coc3-fpc/pic/port]
no-partition interface-type so;
```

This configuration creates interface **so-fpc/pic/port**. When you include the **no-partition** statement at the [edit interfaces **coc3-fpc/pic/port**] hierarchy level, the only configurable interface type is **so**, because the **no-partition** statement signals that you are creating a clear-channel SONET/SDH interface.

On a 2-port or 4-port Channelized OC3 IQE PIC, you can configure two to four separate OC3 clear-channel interfaces by additionally specifying the port numbers. Configuration is otherwise the same as previously described on a (1-port) Channelized OC3 IQ PIC.

Configuring T3 IQ Interfaces

To configure a T3 interface on an OC3 PIC, include the **partition**, **oc-slice**, and **interface-type** statements at the [edit interfaces **coc3-fpc/pic/port**] hierarchy level, specifying the **coc1** interface type:

```
[edit interfaces coc3-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type coc1;
```

When you include the **partition** statement at the `[edit interfaces coc3-fpc/pic/port]` hierarchy level, the only configurable interface type is **coc1**. This configuration creates interface `coc1-fpc/pic/port:channel`.



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQ interfaces. You can only apply CoS rules to the aggregate bit streams.

Then, include the **no-partition interface-type** statement at the `[edit interfaces coc1-fpc/pic/port:channel]` hierarchy level, specifying the **t3** interface type:

```
[edit interfaces coc1-fpc/pic/port:channel]
no-partition interface-type t3;
```

This configuration creates interface `t3-fpc/pic/port:channel`.

Example: Configuring T3 Interfaces

Configure a T3 interface using partition 3 and OC slice 3. This configuration creates interface `t3-1/1/0:3`.

```
[edit interfaces coc3-1/1/0]
partition 3 oc-slice 3 interface-type coc1;
[edit interfaces coc1-1/1/0:3]
no-partition interface-type t3;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring T1 and NxDS0 Interfaces

To configure T1 interfaces on a Channelized OC3 IQ or IQE PIC, perform the following tasks:

1. Partition the channelized OC3 interface into channelized OC1 interfaces by including the **partition**, **oc-slice**, and **interface-type** statements at the `[edit interfaces coc3-fpc/pic/port]` hierarchy level, specifying the **coc1** interface type:

```
[edit interfaces coc3-fpc/pic/port]
partition partition-number oc-slice oc-slice-range interface-type coc1;
```

2. If your network equipment uses VT mapping, partition the channelized OC1 interface into T1 interfaces by including the **partition** and **interface-type** statements at the `[edit interfaces coc1-fpc/pic/port:channel]` hierarchy level, specifying the **t1** interface type:

```
[edit interfaces coc1-fpc/pic/port:channel]
partition partition-number interface-type t1;
```

3. If your network equipment uses M13 or C-bit parity, convert the channelized OC1 interface into a channelized T3 interface by including the **no-partition** and

`interface-type` statements at the `[edit interfaces coc1-fpc/pic/port:channel]` hierarchy level, specifying the `ct3` interface type:

```
[edit interfaces coc1-fpc/pic/port:channel]
no-partition partition-number interface-type ct3;
```



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQ interfaces. You can only apply CoS rules to the aggregate bit streams.

Note that because the `no-partition` statement is included, this configuration does not create another level of channelization, as denoted by the number of colons in the resulting interface.

4. To configure T1 interfaces, partition the channelized T3 interface into T1 interfaces by including the `partition` and `interface-type` statements at the `[edit interfaces ct3-fpc/pic/port:channel]` hierarchy level, specifying the `t1` interface type:

```
[edit interfaces ct3-fpc/pic/port:channel]
partition partition-number interface-type t1;
```


5. To configure NxDS0 interfaces, partition the channelized T3 interface into channelized T1 interfaces by including the `partition` and `interface-type` statements at the `[edit interfaces ct3-fpc/pic/port:channel]` hierarchy level and specifying the `ct1` interface type:

```
[edit interfaces ct3-fpc/pic/port:channel]
partition partition-number interface-type ct1;
```



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQ interfaces. You can only apply CoS rules to the aggregate bit streams.

Figure 47 on page 460 shows VT-mapped and M13 or C-bit parity-mapped configurations of T1 IQ interfaces.

Figure 47: T1 Interfaces on a Channelized OC3 PIC


Bold entries correspond to actual packet channels.

g015503

6. Configure channelized *NxDS0* IQ interfaces on the channelized T1 IQ interface by including the **partition**, **timeslots**, and **interface-type** statements at the [edit interfaces *ct1-fpc/pic/port:channel*] hierarchy level, specifying the **ds** interface type:

```
[edit interfaces ct1-fpc/pic/port:channel]
partition partition-number timeslots time-slot-range interface-type ds;
```

Figure 48 on page 460 shows VT-mapped and M13 or C-bit parity-mapped configurations of *NxDS0* IQ interfaces.

Figure 48: Sample Channelization of OC3 IQ or IQE PIC

Bold entries correspond to actual packet channels.

g015504

Example: Configuring T1 and NxDS0 Interfaces

Configure the following T1 interfaces:

```
t1-0/0/0:1:1
t1-0/0/0:1:2
t1-0/0/0:1:3
t1-0/0/0:1:4
t1-0/0/0:1:5
```

VT-Mapped Configuration

```
[edit interfaces coc3-0/0/0]
partition 1 oc-slice 1 interface-type coc1;
[edit interfaces coc1-0/0/0:1]
```

```
partition 1-5 interface-type t1;
```

**M13 or C-bit
Parity-Mapped
Configuration**

```
[edit interfaces coc3-0/0/0]
partition 1 oc-slice 1 interface-type coc1;
[edit interfaces coc1-0/0/0:1]
no-partition interface-type ct3;
[edit interfaces ct3-0/0/0:1]
partition 1-5 interface-type t1;
```

Configure the following two NxDS0 interfaces with 10 time slots and 4 time slots, respectively:

```
ds-0/0/0:1:2:1
ds-0/0/0:1:2:2
```

**VT-Mapped
Configuration**

```
[edit interfaces coc3-0/0/0]
partition 1 oc-slice 1 interface-type coc1;
[edit interfaces coc1-0/0/0:1]
partition 2 interface-type ct1;
[edit interfaces ct1-0/0/0:1:2]
partition 1 timeslots 1-10 interface-type ds;
partition 2 timeslots 12-16 interface-type ds;
```

**M13 or C-bit
Parity-Mapped
Configuration**

```
[edit interfaces coc3-0/0/0]
partition 1 oc-slice 1 interface-type coc1;
[edit interfaces coc1-0/0/0:1]
no-partition interface-type ct3;
[edit interfaces ct3-0/0/0:1]
partition 2 interface-type ct1;
[edit interfaces ct1-0/0/0:1:2]
partition 1 timeslots 1-10 interface-type ds;
partition 2 timeslots 12-16 interface-type ds;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Example: Setting Remote Loopback and Running BERT Tests on NxDS0 Interfaces

For Channelized OC3 IQ and IQE PICs, if you need remote loopback on a far-end NxDS0 interface, and you are running a BERT test from the local NxDS0 interface, you must set remote loopback on the far-end router's associated channelized T1 interface (ct1). To do this, include the `loopback remote` statement at the `[edit interfaces ct1-fpc/picport t1-options]` hierarchy level. For example:

Local router:

```
[edit interfaces]
ct1-0/0/0:2:2 {
  partition 1 timeslots 1-10 interface-type ds;
  ds-0/0/0:2:2:1 {
    ds0-options {
      bert-period 30;
    }
  }
}
```

```
}
```

Remote router:

```
[edit interfaces]
ct1-0/0/0:2:2 {
  partition 1 timeslots 1-10 interface-type ds;
  t1-options {
    loopback remote;
  }
}
```

Configuring Fractional T1 IQ Interfaces

By default, all the time slots on a channelized T1 interface are used. To configure a fractional T1 interface on a Channelized OC3 IQ or IQE PIC, you must perform the following tasks:

1. Configure a T1 interface on the Channelized OC3 IQ or IQE PIC. For more information, see “Configuring T1 and NxDS0 Interfaces” on page 458.
2. Configure the number of time slots allocated to the T1 IQ interface by including the `timeslots` statement at the `[edit interfaces t1-fpc/pic/port<:channel> t1-options]` hierarchy level:

```
[edit interfaces t1-fpc/pic/port<:channel> t1-options]
timeslots time-slot-range;
```

For channelized T1 IQ interfaces, the time-slot range is from 1 through 24. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces. For more information, see “Configuring Fractional T1 Time Slots” on page 567.

Example: Configuring Fractional T1 IQ Interfaces

Configure a fractional T1 interface that uses time slots 1 through 5 and 10:

```
[edit interfaces coc3-0/0/0]
partition 1 oc-slice 1 interface-type coc1;
[edit interfaces coc1-0/0/0:1]
partition 1 interface-type t1;
[edit interfaces t1-0/0/0:1:1 t1-options]
timeslots 1-5,10;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring Link PIC Failover on Channelized OC3 IQ and IQE Interfaces

For Channelized OC3 IQ or IQE PICs used as linking PICs in redundant LSQ configurations, you can inhibit the router from sending PPP termination-request messages to the remote host if the link PIC fails. To do this, include the

`no-termination-request` statement at the [edit interfaces *interface-name* ppp-options] hierarchy level:

```
no-termination-request;
```

The `no-termination-request` statement is supported only with MLPPP and SONET APS configurations and works with PPP, PPP over Frame Relay, and MLPPP interfaces only.

For information about interchassis and intrachassis LSQ failover, see the *JUNOS Services Interfaces Configuration Guide*.

Chapter 21

Configuring Channelized STM1 Interfaces

Each Channelized STM1 PIC and Channelized STM1 Intelligent Queuing (IQ) PIC has one STM1 port.

For the Channelized STM1 IQ or IQE PIC, you can channelize the single port to the NxDS0 level. Each E1 interface has 32 time slots (DS0), in which time slot 0 is reserved.

You can combine one or more of these DS0 time slots (channels) to create a channel group (NxDS0).

This section contains the following topics:

- Configuring Channelized STM1 IQ and IQE Interfaces on page 465
- Configuring Channelized STM1 Interfaces on page 471
- Configuring Link PIC Failover on Channelized STM1 Interfaces on page 475
- Example: Configuring Channelized STM1 Interfaces on page 475

Configuring Channelized STM1 IQ and IQE Interfaces

This section includes the following topics:

- Configuring an STM1 IQ or STM1 IQE Interface on page 465
- Configuring E1 IQ and IQE Interfaces on page 466
- Configuring Fractional E1 IQ and IQE Interfaces on page 467
- Configuring an NxDS0 IQ Interface on page 468
- Example: Configuring Channelized STM1 IQ and IQE Interfaces on page 469

Configuring an STM1 IQ or STM1 IQE Interface

On a one-port Channelized STM1 IQ PIC, or each individual port of the 4-port Channelized STM1 IQE PIC, you can configure one SDH STM1 interface. To configure an SDH STM1 interface, include the **no-partition interface-type** statement at the [edit interfaces cstm1-*fpc/pic/port*] hierarchy level, specifying the **so** interface type:

```
[edit interfaces cstm1-fpc/pic/port]  
no-partition interface-type so;
```

This configuration creates interface *so-fpc/pic/port*.



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQ and IQE interfaces. You can only apply CoS rules to the aggregate bit streams.

Configuring E1 IQ and IQE Interfaces

To configure an E1 interface on a Channelized STM1 IQ or IQE PIC, perform the following tasks:

1. Include the **no-partition** and **interface-type** statements at the **[edit interfaces cstm1-fpc/pic/port]** hierarchy level, specifying the **cau4** interface type. This converts the channelized STM1 interface into a channelized AU-4 interface. The resulting interface name is **cau4-fpc/pic/port**:

```
[edit interfaces cstm1-fpc/pic/port]
no-partition interface-type cau4;
```

2. Partition the channelized AU-4 interface into E1 interfaces by including the **partition** and **interface-type** statements at the **[edit interfaces cau4-fpc/pic/port]** hierarchy level, specifying the **e1** interface type. This configuration creates interface **e1-fpc/pic/port:channel**. The partition number is the sublevel interface partition index and is correlated with the channel number. For channelized E1 interfaces, the partition number can be from 1 through 63. The interface type is the channelized interface type or clear channel you are creating. For channelized AU-4 interfaces, **type** can be **ce1** or **e1**.

```
[edit interfaces cau4-fpc/pic/port]
partition partition-number interface-type e1;
```



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQ or IQE interfaces. You can only apply CoS rules to the aggregate bit streams.



NOTE: For channelized STM1 interfaces, channel numbering begins with 0 (:0). For channelized STM1 IQ and IQE interfaces, channel numbering begins with 1 (:1).

Example: Configuring E1 IQ and IQE Interfaces

Configure the following five E1 interfaces:

```
e1-0/0/0:1
e1-0/0/0:2
e1-0/0/0:3
e1-0/0/0:4
e1-0/0/0:5
```

```
[edit interfaces cstm1-0/0/0]
no-partition interface-type cau4;
[edit interfaces cau4-0/0/0]
partition 1-5 interface-type e1;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring Fractional E1 IQ and IQE Interfaces

By default, all the time slots on a channelized E1 interface are used. To configure a fractional E1 interface on a Channelized STM1 IQ or IQE PIC, perform the following tasks:

1. Include the `no-partition` and `interface-type` statements at the `[edit interfaces cstm1-fpc/pic/port]` hierarchy level, specifying the `cau4` interface type. This converts the channelized STM1 interface into a channelized AU-4 interface. The resulting interface name is `cau4-fpc/pic/port`:

```
[edit interfaces cstm1-fpc/pic/port]
no-partition interface-type cau4;
```

2. Partition the channelized AU-4 interface into E1 interfaces by including the `partition` and `interface-type` statements at the `[edit interfaces cau4-fpc/pic/port]` hierarchy level, specifying the `e1` interface type. The partition number is the sublevel interface partition index and is correlated with the channel number. For channelized E1 interfaces, the partition number can be from 1 through 63. The interface type is the channelized interface type or clear channel you are creating. For channelized AU-4 interfaces, `type` can be `ce1` or `e1`. This configuration creates interface `e1-fpc/pic/port:channel`:

```
[edit interfaces cau4-fpc/pic/port]
partition partition-number interface-type e1;
```

3. Configure the number of time slots allocated to the E1 IQ or IQE interface by including the `timeslots` statement at the `[edit interfaces e1-fpc/pic/port:channel e1-options]` hierarchy level. `NxDS0` time slots configured on either a channelized STM1 IQ or IQE interface or channelized E1 IQ or IQE interface are numbered from 1 to 31 (0 is reserved), while fractional E1 time slots range from 2 to 32 (1 is reserved). To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces.

```
[edit interfaces e1-fpc/pic/port:channel e1-options]
timeslots time-slot-range;
```



NOTE: For channelized STM1 interfaces, channel numbering begins with 0 (:0). For channelized STM1 IQ or IQE interfaces, channel numbering begins with 1 (:1).

For more information about E1 time slots, see “Configuring Fractional E1 Time Slots” on page 548.

Example: Configuring Fractional E1 Interfaces

Configure a fractional E1 interface that uses time slots 2 through 10:

```
[edit interfaces cstm1-0/0/0]
no-partition cau4;
[edit interfaces cau4-0/0/0]
partition 1 interface-type e1;
[edit interfaces e1-0/0/0 e1-options]
timeslots 2-10;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring an NxDS0 IQ Interface

By default, all the time slots on a channelized STM1 interface are used. To configure an NxDS0 IQ interface on a Channelized STM1 IQ or IQE PIC, perform the following tasks:

1. Include the `no-partition` and `interface-type` statements at the `[edit interfaces cstm1-fpc/pic/port]` hierarchy level, specifying the `cau4` interface type. This converts the channelized STM1 interface into a channelized AU-4 interface. The resulting interface name is `cau4-fpc/pic/port`:

```
[edit interfaces cstm1-fpc/pic/port]
no-partition interface-type cau4;
```

2. Partition the channelized AU-4 interface into E1 interfaces by including the `partition` and `interface-type` statements at the `[edit interfaces cau4-fpc/pic/port]` hierarchy level, specifying the `ce1` interface type. This configuration creates interface `ce1-fpc/pic/port:channel`. The partition number is the sublevel interface partition index and is correlated with the channel number. For channelized E1 interfaces, the partition number can be from 1 through 63. The interface type is the channelized interface type or clear channel you are creating. For channelized AU-4 interfaces, `type` can be `ce1` or `e1`:

```
[edit interfaces cau4-fpc/pic/port]
partition partition-number interface-type ce1;
```

3. Configure the number of time slots allocated to the NxDS0 IQ interface by including the `partition`, `timeslots`, and `interface-type` statements at the `[edit interfaces e1-fpc/pic/port:channel]` hierarchy level, specifying the `ds` interface type. For channelized E1 IQ interfaces, the partition number range is from 1 through 31. For E1 IQ interfaces (`e1-fpc/pic/port`), the time-slot range is from 2 through 31. For channelized E1 IQ interfaces (`ce1-fpc/pic/port`), the time-slot range is from 1 through 31. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces:

```
[edit interfaces ce1-fpc/pic/port:channel]
partition partition-number timeslots time-slot-range interface-type ds;
```



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQ and IQE interfaces. You can only apply CoS rules to the aggregate bit streams.



NOTE: For channelized STM1 interfaces, channel numbering begins with 0 (:0). For channelized STM1 IQ and IQE interfaces, channel numbering begins with 1 (:1).

For more information about E1 time slots, see “Configuring Fractional E1 Time Slots” on page 548.

Example: Configuring an NxDS0 IQ Interface

Configure an NxDS0 interface that uses time slots 1 through 10. This configuration creates the ds-0/0/0:1:1 interface.

```
[edit interfaces cstm1-0/0/0]
no-partition interface-type cau4;
[edit interfaces cau4-0/0/0]
partition 1 interface-type ce1;
[edit interfaces ce1-0/0/0:1]
partition 1 timeslots 1-10 interface-type ds;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Example: Configuring Channelized STM1 IQ and IQE Interfaces

Configure STM1, E1, fractional E1, and NxDS0 interfaces:

STM1 Interface	<pre>[edit interfaces] cstm1-0/0/0 { no-partition interface-type so; } so-0/0/0 { unit 0 { family inet { address 10.10.12.1/30; } } }</pre>
E1 Interface	<pre>[edit interfaces] cstm1-1/1/0 { no-partition interface-type cau4; } [edit interfaces] cau4-1/1/0 { partition 1-63 interface-type e1; } [edit interfaces] e1-1/1/0:1 {</pre>

```

        unit 0 {
            family inet {
                address 10.10.10.1/30;
            }
        }
    }
    ...

```

Fractional E1 Interface

```

[edit interfaces]
cstm1-1/0/0 {
    no-partition interface-type cau4;
}
[edit interfaces]
cau4-1/0/0 {
    partition 1-63 interface-type e1;
}
[edit interfaces]
e1-1/1/0:1 {
    e1-options {
        timeslots 2-10;
    }
    unit 0 {
        family inet {
            address 10.10.10.1/30;
        }
    }
}
...

```

DS0 Interface

```

[edit interfaces]
cstm1-2/0/0 {
    no-partition interface-type cau4;
}
[edit interfaces]
cau4-2/0/0 {
    partition 1-10 interface-type ce1;
}
[edit interfaces]
ce1-2/0/0:1 {
    partition 1 interface-type ds timeslots 2-10;
    [edit interfaces]
    ds-2/0/0:1:1 {
        unit 0 {
            family inet {
                address 10.12.12.1/30;
            }
        }
    }
}
...
}

```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring Channelized STM1 Interfaces

To specify the channel number, include it after the colon (:) in the interface name. For example, a Channelized STM1-to-E1 PIC in FPC 1 and slot 1 will have the following physical interface, depending on the media type:

```
e1-1/1/0:x
```

The E1 channel number can be from 0 through 62.

This section contains the following topics:

- Configuring Channelized STM1 Interface Properties on page 471
- Configuring Virtual Tributary Mapping of Channelized STM1 Interfaces on page 472

Configuring Channelized STM1 Interface Properties

To configure the interface properties for Channelized STM1-to-E1 PICs, include the **e1-options** and **sonet-options** statements for both sides of the connection. The following configurations list all the valid statements.

To specify options for each of the E1 channels on the Channelized STM1-to-E1 PIC, include the **e1-options** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
e1-options {
  bert-error-rate;
  bert-period;
  fcs (16 | 32);
  framing (g704 | g704-no-crc4 | unframed);
  idle-cycle-flag (flags | ones);
  loopback (local | remote);
  start-end-flag (filler | shared);
  timeslots time-slot-number;
}
```



NOTE: When a channelized STM1 interface experiences a line transition, the E1 channels configured in unframed mode log a large number of drops (around 24,000) as the channelized STM1 interface clocks resynchronize. This does not occur on framed channels, because the framing resynchronizes clocks very quickly.

To specify options for the SONET/SDH side of the connection, include the **sonet-options** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
sonet-options {
  aps {
    advertise-interval milliseconds;
    authentication-key key;
    force;
```

```

    hold-time milliseconds;
    lockout;
    neighbor address;
    paired-group group-name;
    protect-circuit group-name;
    request;
    revert-time seconds;
    switching-mode (bidirectional | unidirectional);
    working-circuit group-name;
}
bytes {
    e1-quiet value;
    f1 value;
    f2 value;
    s1 value;
    z3 value;
    z4 value;
}
loopback (local | remote);
}

```



NOTE: On channelized STM1 interfaces, you should configure the clock source on one side of the connection to be internal (the default JUNOS configuration) and on the other side of the connection to be external.

For information about Frame Relay DLCI limitations for channelized interfaces, see “Data-Link Connection Identifiers on Channelized Interfaces” on page 388. For more information about Frame Relay DLCIs, see “Configuring a Point-to-Point Frame Relay Connection” on page 380. For information about DLCI sparse mode, see the *JUNOS System Basics Configuration Guide*.

For more information about specific statements, see “Configuring E1 Interfaces” on page 543, “Configuring SONET/SDH Interfaces” on page 843, and “Configuring T1 Interfaces” on page 559. For a configuration example, see “Example: Configuring Channelized STM1 Interfaces” on page 475.

Configuring Virtual Tributary Mapping of Channelized STM1 Interfaces

You can configure virtual tributary mapping to use KLM mode or ITU-T mode. To configure virtual tributary mapping, include the `vtmapping` statement at the [edit chassis fpc slot-number pic pic-number] hierarchy level:

```

[edit chassis fpc slot-number pic pic-number]
  vtmapping (klm | itu-t);

```

By default, virtual tributary mapping uses KLM mode. For more information, see the *JUNOS System Basics Configuration Guide*.

For the Channelized STM1 IQ and IQE PICs, you can configure virtual tributary mapping by including the `vtmapping` statement at the [edit interfaces cau4-fpc/pic/port sonet-options] hierarchy level:


```
[edit interfaces cau4-fpc/pic/port sonet-options]
vtmapping (klm | itu-t);
```

Table 42 on page 473 lists the KLM mappings used by the channelized STM1-to-E1 PIC interfaces. The PIC defaults to KLM numbering with an offset of -1 ; for example, KLM 1 = STM1 PIC 0.

Table 42: Channelized STM1-to-E1 Channel Mapping

Channel Number	KLM Number	Tributary Unit Group 3	Tributary Unit Group 2	Virtual Tributary	ITU-T Number
0	1	1	1	1	1
1	2	1	1	2	22
2	3	1	1	3	43
3	4	1	2	1	4
4	5	1	2	2	25
5	6	1	2	3	46
6	7	1	3	1	7
7	8	1	3	2	28
8	9	1	3	3	49
9	10	1	4	1	10
10	11	1	4	2	31
11	12	1	4	3	52
12	13	1	5	1	13
13	14	1	5	2	34
14	15	1	5	3	55
15	16	1	6	1	16
16	17	1	6	2	37
17	18	1	6	3	58
18	19	1	7	1	19
19	20	1	7	2	40
20	21	1	7	3	61
21	22	2	1	1	2
22	23	2	1	2	23

Table 42: Channelized STM1-to-E1 Channel Mapping *(continued)*

Channel Number	KLM Number	Tributary Unit Group 3	Tributary Unit Group 2	Virtual Tributary	ITU-T Number
23	24	2	1	3	44
24	25	2	2	1	5
25	26	2	2	2	26
26	27	2	2	3	47
27	28	2	3	1	8
28	29	2	3	2	29
29	30	2	3	3	50
30	31	2	4	1	11
31	32	2	4	2	32
32	33	2	4	3	53
33	34	2	5	1	14
34	35	2	5	2	35
35	36	2	5	3	56
36	37	2	6	1	17
37	38	2	6	2	38
38	39	2	6	3	59
39	40	2	7	1	20
40	41	2	7	2	41
41	42	2	7	3	62
42	43	3	1	1	3
43	44	3	1	2	24
44	45	3	1	3	45
45	46	3	2	1	6
46	47	3	2	2	27
47	48	3	2	3	48
48	49	3	3	1	9
49	50	3	3	2	30

Table 42: Channelized STM1-to-E1 Channel Mapping (*continued*)

Channel Number	KLM Number	Tributary Unit Group 3	Tributary Unit Group 2	Virtual Tributary	ITU-T Number
50	51	3	3	3	51
51	52	3	4	1	12
52	53	3	4	2	33
53	54	3	4	3	54
54	55	3	5	1	15
55	56	3	5	2	36
56	57	3	5	3	57
57	58	3	6	1	18
58	59	3	6	2	39
59	60	3	6	3	60
60	61	3	7	1	21
61	62	3	7	2	42
62	63	3	7	3	63

Configuring Link PIC Failover on Channelized STM1 Interfaces

For Channelized STM1 IQ and IQE PICs used as linking PICs in redundant LSQ configurations, you can inhibit the router from sending PPP termination-request messages to the remote host if the link PIC fails. To do this, include the `no-termination-request` statement at the [edit interfaces *interface-name* ppp-options] hierarchy level:

```
no-termination-request;
```

The `no-termination-request` statement is supported only with MLPPP and SONET APS configurations and works with PPP, PPP over Frame Relay, and MLPPP interfaces only.

For information about interchassis and intrachassis LSQ failover, see the *JUNOS Services Interfaces Configuration Guide*.

Example: Configuring Channelized STM1 Interfaces

The following configuration is sufficient to get the Channelized STM1-to-E1 PIC interface up and running. The channelized STM1-to-E1 interface is an STM1 that is divided into 63 E1 interfaces. E1 interfaces can use the following encapsulation types:

- PPP, PPP CCC, and PPP TCC
- Frame Relay, Frame Relay CCC, and Frame Relay TCC
- Cisco HDLC, Cisco HDLC CCC, and Cisco HDLC TCC

The channels can also have logical interfaces. For information about Frame Relay DLCI limitations for channelized interfaces, see “Data-Link Connection Identifiers on Channelized Interfaces” on page 388. For more information about Frame Relay DLCIs, see “Configuring a Point-to-Point Frame Relay Connection” on page 380. For more information about DLCI sparse mode, see the *JUNOS System Basics Configuration Guide*.

You apply all STM1 interface SONET/SDH options to the first E1 interface in the configuration by including the **sonet-options** statement at the [edit interfaces **e1-fpc/pic/port:channel**] hierarchy level:

```
[edit]
interfaces {
  e1-fpc/pic/port:0 {
    encapsulation cisco-hdlc;
    sonet-options {
      no-z0-increment;
    }
    e1-options {
      framing g704;
    }
    unit 0 {
      family inet {
        address 10.11.30.1/30;
      }
    }
  }
  e1-fpc/pic/port:1 {
    encapsulation frame-relay;
    e1-options {
      framing g704;
    }
    unit 1 {
      dlci 16;
      family inet {
        address 10.11.31.9/30;
      }
    }
  }
  e1-fpc/pic/port:2 {
    encapsulation ppp;
    no-keepalives;
    unit 0 {
      family inet {
        address 10.11.31.47/30;
      }
    }
  }
}
[edit]
```

```
chassis {  
  fpc 2 {  
    pic 0 {  
      vtmapping klm;  
    }  
  }  
}
```


Chapter 22

Configuring Channelized T3 Interfaces



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQ interfaces. You can only apply CoS rules to the aggregate bit streams.

This section contains the following topics:

- Configuring Channelized T3 IQ Interfaces on page 479
- Configuring Channelized DS3-to-DS0 Interfaces on page 482
- Configuring Channelized DS3-to-DS1 Interfaces on page 485
- Example: Configuring Channelized T3 IQ Interfaces on page 486
- Examples: Configuring Channelized DS3-to-DS0 Interfaces on page 487
- Examples: Configuring Channelized DS3-to-DS1 Interfaces on page 490

Configuring Channelized T3 IQ Interfaces

This section describes how to configure channelized T3 intelligent queuing (IQ) interfaces, discussing the following topics:

- Configuring T3 IQ Interfaces on page 479
- Configuring T1 IQ Interfaces on page 480
- Configuring Fractional T1 IQ and IQE Interfaces on page 480
- Configuring an NxDS0 IQ Interface on page 481

Configuring T3 IQ Interfaces

To configure a T3 interface, include the **no-partition** and **interface-type** statements at the [edit interfaces *ct3-fpc/pic/port*] hierarchy level:

```
[edit interfaces ct3-fpc/pic/port]
no-partition interface-type t3;
```

This configuration creates interface *t3-fpc/pic/port*.

Configuring T1 IQ Interfaces

On a Channelized DS3 IQ or IQE Physical Interface Card (PIC), you can create up to 112 T1 interfaces. To configure a T1 interface on a Channelized DS3 IQ or IQE PIC, include the **partition** and **interface-type** statements at the `[edit interfaces ct3-fpc/pic/port]` hierarchy level, specifying the **t1** interface type:

```
[edit interfaces ct3-fpc/pic/port]
partition partition-number interface-type t1;
```

This configuration creates interface **t1-fpc/pic/port:channel**.

The partition number is the sublevel interface partition index and is correlated with the channel number. For channelized T3 interfaces, the partition number can be from 1 through 28.



NOTE: For channelized T3 interfaces, channel numbering begins with 0 (:0). For channelized T3 IQ and IQE interfaces, channel numbering begins with 1 (:1).

The interface type is the channelized interface type or clear channel you are creating. For channelized T3 interfaces, **type** can be **ct1** or **t1**.

Example: Configuring T1 IQ and IQE Interfaces

Configure the following five T1 interfaces:

```
t1-0/0/0:1
t1-0/0/0:2
t1-0/0/0:3
t1-0/0/0:4
t1-0/0/0:5
```

```
[edit interfaces ct3-0/0/0]
partition 1-5 interface-type t1;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring Fractional T1 IQ and IQE Interfaces

By default, all the time slots on a channelized T1 interface are used. To configure a fractional T1 interface on a Channelized DS3 IQ or IQE PIC, perform the following tasks:

1. Configure a T1 IQ interface. For more information, see “Configuring T1 IQ Interfaces” on page 480.

This configuration creates interface **t1-fpc/pic/port:channel**.

2. Configure the number of time slots allocated to the T1 IQ interface by including the `timeslots` statement at the `[edit interfaces t1-fpc/pic/port:channel t1-options]` hierarchy level:

```
[edit interfaces t1-fpc/pic/port t1-options]
timeslots time-slot-range;
```

For channelized T1 IQ interfaces, the time-slot range is from 1 through 24. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces. For more information about T1 time slots, see “Configuring Fractional T1 Time Slots” on page 567.

Example: Configuring Fractional T1 IQ Interfaces

Configure a fractional T1 interface that uses time slots 1 through 10:

```
[edit interfaces ct3-0/0/0:1]
partition 1 interface-type t1;
[edit interfaces t1-0/0/0:1:1 t1-options]
timeslots 1-10;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring an NxDS0 IQ Interface

By default, all the time slots on a channelized T3 interface are used. To configure an NxDS0 IQ interface on a Channelized DS3 IQ or IQE PIC, perform the following tasks:

1. Partition the channelized T3 interface into channelized T1 interfaces by including the `partition` and `interface-type` statements at the `[edit interfaces ct3-fpc/pic/port]` hierarchy level, specifying the `ct1` interface type:

```
[edit interfaces ct3-fpc/pic/port]
partition partition-number interface-type ct1;
```

This configuration creates interface `ct1-fpc/pic/port:channel`.

The partition number is the sublevel interface partition index and is correlated with the channel number. For channelized T1 interfaces, the partition number can be from 1 through 28.

The interface type is the channelized interface type or clear channel you are creating. For channelized T3 interfaces, *type* can be `ct1` or `t1`.



NOTE: For channelized T3 interfaces, channel numbering begins with 0 (:0). For channelized T3 IQ interfaces, channel numbering begins with 1 (:1).

2. Configure the number of time slots allocated to the NxDS0 IQ interface by including the `partition`, `timeslots`, and `interface-type` statements at the `[edit`

interfaces *ct1-fpc/pic/port:channel*] hierarchy level, specifying the **ds** interface type:

```
[edit interfaces ct1-fpc/pic/port:channel]
partition partition-number timeslots time-slot-range interface-type ds;
```

For channelized T1 IQ interfaces, the partition number range is from 1 through 28; the time-slot range is from 1 through 24. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces. For more information about T1 time slots, see “Configuring Fractional T1 Time Slots” on page 567.

Example: Configuring an NxDS0 IQ Interface

Configure the following two NxDS0 interfaces with 10 time slots and 4 time slots, respectively:

```
ds-0/0/0:1:1
ds-0/0/0:1:2

[edit interfaces ct3-0/0/0]
partition 1 interface-type ct1;
[edit interfaces ct1-0/0/0:1]
partition 1 timeslots 1-10 interface-type ds;
partition 2 timeslots 12-16 interface-type ds;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring Channelized DS3-to-DS0 Interfaces

For channelized interfaces, you can configure 28 T1 channels per T3 interface. Each T1 link can have up to eight DS0 channel groups, and each channel group can hold any combination of DS0 time slots. To specify the T1 link and DS0 channel group number in the interface name, use colons (:) as separators. For example, a Multichannel DS3 PIC might have the following physical and virtual interfaces:

```
ds-0/0/0:x:y
```

where *x* is a T1 link ranging from 0 through 27 and *y* is a DS0 channel group from 0 through 7. For more information about ranges, see Table 43 on page 483.

You can use any of the values within the range available for *x* and *y*, and you do not have to configure the links sequentially. In addition, the JUNOS Software applies the interface options you configure according to the following rules:

- To configure the T1 options, you must set channel group *y* to 0; the T1 link *x* can be any value:

```
ds-0/0/0:x:0
```

- To configure the T3 options, you must set the T1 link *x* to 0 and channel group *y* to 0:

ds-0/0/0:0:0

- There are no restrictions on configuring the DS0 options.
- If you delete a configuration you previously committed for channel group 0, the options return to default values.

By default, all the time slots are used. To configure the channel groups and time slots for a channelized DS3-to-DS0 interface, include the `channel-group` and `timeslots` statements at the `[edit chassis fpc slot-number pic pic-number ct3 port port-number t1 link-number]` hierarchy level:

```
[edit chassis fpc slot-number pic pic-number ct3 port port-number t1 link-number ]
channel-group group-number;
timeslots time-slot-range;
```



NOTE: If you commit the interface name but do not include the `[edit chassis]` configuration, the channelized DS3-to-DS0 interface behaves like a channelized DS3-to-DS1 interface: none of the DS0 functionality is accessible.

Table 43 on page 483 shows the ranges you can specify for each of the elements in the preceding configuration.

Table 43: Ranges for Channelized DS3-to-DS0 Configuration

Item	Option	Range
FPC slot	slot-number	0 through 7 (see note below)
PIC slot	pic-number	0 through 3
Port	port-number	0 through 1
T1 link	link-number	0 through 27
DS0 channel group	group-number	0 through 7
Time slot	time-slot-range	1 through 24



NOTE: The FPC slot range depends on the router. For a routing matrix, the range is from 0 through 31. For M40, M40e, M160, M320, M120, and other T Series routers, the range is from 0 through 7. For M20 routers, the range is from 0 through 3. For M10 and M10i routers the range is from 0 through 1. For M5 and M7i routers, the only applicable value is 0.

Bandwidth limitations restrict the interface to a maximum of 128 channel groups per T3 port, rather than the theoretical maximum of $8 * 28 = 224$.

There are 24 time slots on a T1 interface. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces. You can use each time slot number on only one channel group within the same T1 link.

To configure channelized DS3-to-DS0 interface properties, you can include the **t3-options**, **t1-options**, and **ds0-options** statements. Only a subset of the T3 options are valid for this configuration, and the **buildout**, **invert-data**, and **line-encoding** statements at the **[edit interfaces *interface-name* t1-options]** hierarchy level are ignored. Likewise, only a subset of the DS0 options are valid for this configuration, and the **bert-algorithm**, **bert-error-rate**, **bert-period**, and **loopback payload** statements at the **[edit interfaces *interface-name* ds0-options]** hierarchy level are ignored. The following configurations list all the valid parameters.



NOTE: The set of options the JUNOS Software applies to the interface depends on how you specify the interface name. For more information, see “Examples: Configuring Channelized DS3-to-DS0 Interfaces” on page 487.

To specify options for the T3 side of the connection, include the **t3-options** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
t3-options {
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  (cbit-parity | no-cbit-parity);
  (long-buildout | no-long-buildout);
  loopback (local | payload | remote);
}
```

The statements at the **t3-options** hierarchy are supported only for channel 0; they are ignored if configured on other channels. To specify options for each of the T1 channels, include the **t1-options** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
t1-options {
  byte-encoding (nx56 | nx64);
  fcs (16 | 32);
  framing (esf | lf);
  idle-cycle-flag (flags | ones);
  invert-data;
  loopback (local | payload | remote);
  start-end-flag (filler | shared);
  timeslots time-slot-number;
}
```

To specify options for each of the DS0 channels, include the **ds0-options** statement at the **[edit interfaces *interface-name*]** hierarchy level:

```
[edit interfaces interface-name]
ds0-options {
```

```

bert-algorithm algorithm;
bert-error-rate rate;
bert-period seconds;
byte-encoding (nx56 | nx64);
fcs (16 | 32);
idle-cycle-flag (flags | ones);
invert-data;
loopback payload;
start-end-flag (filler | shared);
}

```

For more information about specific parameters, see “Configuring E1 Interfaces” on page 543, “Configuring E3 Interfaces” on page 551, “Configuring T1 Interfaces” on page 559, and “Configuring T3 Interfaces” on page 569. For a configuration example, see “Examples: Configuring Channelized DS3-to-DS0 Interfaces” on page 487.

For information about Frame Relay DLCI limitations for channelized interfaces, see “Data-Link Connection Identifiers on Channelized Interfaces” on page 388. For more information about Frame Relay DLCIs, see “Configuring a Point-to-Point Frame Relay Connection” on page 380. For more information about DLCI sparse mode, see the *JUNOS System Basics Configuration Guide*.

Each T1 link can have up to eight DS0 channel groups, and each channel group can hold any combination of DS0 time slots.

Configuring Channelized DS3-to-DS1 Interfaces

You can configure 28 T1 channels per T3 interface, and each interface can have logical interfaces. To specify the channel number, include it after the colon (:) in the interface name. For example, a 4-port T3 PIC in FPC 1 and slot 1 will have the following physical interfaces, depending on the media type:

```

t1-1/1/0:x
t1-1/1/1:x
t1-1/1/2:x
t1-1/1/3:x

```

where x is a channel number ranging from 0 through 27.

To configure channelized DS3-to-DS1 interface properties, you can include both the **t1-options** and **t3-options** statements. Only a subset of the T3 options is valid for this configuration, and the **buildout**, **invert-data**, and **line-encoding** statements at the [edit interfaces *interface-name* **t1-options**] hierarchy level are ignored. Likewise, only a subset of the DS0 options are valid for this configuration, and the **bert-algorithm**, **bert-error-rate**, **bert-period**, and **loopback payload** statements at the [edit interfaces *interface-name* **ds0-options**] hierarchy level are ignored. The following configuration lists all the valid parameters.

To specify options for the T3 side of the connection, include the **t3-options** statement at the [edit interfaces *interface-name*] hierarchy level:

```

[edit interfaces interface-name]
t3-options {

```

```

    bert-algorithm algorithm;
    bert-error-rate rate;
    bert-period seconds;
    (cbit-parity | no-cbit-parity);
    (feac-loop-respond | no-feac-loop-respond);
    loopback (local | payload | remote);
}

```

The statements in the **t3-options** hierarchy are supported only for channel 0; they are ignored if configured on other channels.

To specify options for each of the T1 channels, include the **t1-options** statement at the [edit interfaces *interface-name*] hierarchy level:

```

[edit interfaces interface-name]
t1-options {
    byte-encoding (nx56 | nx64);
    fcs (16 | 32);
    framing (sf | esf);
    idle-cycle-flag (flags | ones);
    loopback (local | payload | remote);
    start-end-flag (filler | shared);
    timeslots time-slot-number;
}

```

For T1 channels on a channelized T3 interface, the **clocking** statement is supported only for channel 0; it is ignored if included in the configuration of channels 1 through 11. The clock source configured for channel 0 applies to all channels on the channelized T3 interface. The individual T1 channels use a gapped 45-MHz clock as the transmit clock. When you configure the clock source for a channelized interface—**ds-fpc/pic/port** :0, for example—you must also include the **channel-group** statement at the [edit chassis] hierarchy level, and specify channel group 0. For more information, see “Clock Sources on Channelized Interfaces” on page 390.

For information about Frame Relay DLCI limitations for channelized interfaces, see “Data-Link Connection Identifiers on Channelized Interfaces” on page 388. For more information about Frame Relay DLCIs, see “Configuring a Point-to-Point Frame Relay Connection” on page 380. For more information about DLCI sparse mode, see the *JUNOS System Basics Configuration Guide*.

For more information about specific parameters, see “Configuring T1 Interfaces” on page 559 and “Configuring T3 Interfaces” on page 569. For a configuration example, see “Examples: Configuring Channelized DS3-to-DS1 Interfaces” on page 490.

Example: Configuring Channelized T3 IQ Interfaces

Configure a channelized T3 interface as an unpartitioned, clear channel.


Configuring a T3 Interface	<pre> [edit interfaces] ct3-5/0/0 { no-partition interface-type t3; } </pre>
---------------------------------------	--

Configuring NxDS0 and T1 Interfaces

Figure 49 on page 487 shows the following interfaces on a Channelized DS3 IQ or IQE PIC:

- A channelized T1, which is partitioned into NxDS0 interfaces
- T1 interfaces

Figure 49: Sample Channelization of DS3 IQ or IQE PIC



Bold entries correspond to actual packet channels.

g003015

```

[edit interfaces]
ct3-1/1/0 {
  description "CT3 to CT1 and CT3 to T1.";
  t3-options {
    loopback remote;
    looptiming;
  }
  partition 1 interface-type ct1; # ct1-1/1/0:1.
  partition 2-28 interface-type t1; # t1-1/1/0:[2-28]
}
ct1-1/1/0:1 {
  description "case (a) CT1s to NxDS0s.";
  t1-options {
    bert-algorithm all-ones-repeating;
    framing sf;
    line-encoding ami;
  }
  partition 1 timeslots 2 - 10 interface-type ds0; # ds-1/1/0:1:1, channel group with
    10 DS0s
  partition 2 timeslots 11- 23 interface-type ds0; # ds-1/1/0:1:2, channel group with
    13 DS0s
  ...
}
  
```

Examples: Configuring Channelized DS3-to-DS0 Interfaces

The following configuration is sufficient to get the channelized DS3-to-DS0 interface up and running. The T3 interface can be divided into 28 channels, each at T1 line rate. DS3 channels can use the following encapsulation types for their logical interfaces:

- PPP, PPP CCC, and PPP TCC
- Frame Relay, Frame Relay CCC, and Frame Relay TCC

■ Cisco HDLC, Cisco HDLC CCC, and Cisco HDLC TCC

For more information, see “Configuring a Point-to-Point Frame Relay Connection” on page 380.



NOTE: All these configuration examples specify channel group 0 in the interface address, which is required for configuring the **t3-options** and **t1-options** statements.

**Configuring Cisco HDLC
Encapsulation on a
Channelized DS3-to-DS0
Interface**

```
[edit interfaces]
ds-2/0/1:20:0 {
  encapsulation cisco-hdlc;
  unit 0 {
    family inet {
      address 10.0.4.40/32 {
        destination 10.0.4.41;
      }
    }
  }
}
[edit chassis]
fpc 2 {
  pic 0 {
    ct3 {
      port 1 {
        t1 20 {
          channel-group 0 timeslots 1-5;
        }
      }
    }
  }
}
```

**Configuring PPP
Encapsulation on a
Channelized DS3-to-DS0
Interface**

```
[edit interfaces]
ds-2/0/1:20:0 {
  encapsulation ppp;
  unit 0 {
    family inet {
      address 10.0.4.40/32 {
        destination 10.0.4.41;
      }
    }
  }
}
[edit chassis]
fpc 2 {
  pic 0 {
    ct3 {
      port 1 {
        t1 20 {
          channel-group 0 timeslots 1-5;
        }
      }
    }
  }
}
```


**Configuring Three Frame
Relay DLCIs on a
Channelized DS3
Interface**

```

    }
  }
}

[edit interfaces]
t1-5/1/3:0 {
  mtu 9192;
  encapsulation frame-relay;
  unit 1 {
    dlc1 101;
    family inet {
      mtu 9000;
      address 10.123.1.2/32 {
        destination 10.123.1.1;
      }
    }
    family iso {
      mtu 9000;
    }
    family mpls {
      mtu 9000;
    }
  }
  unit 2 {
    dlc1 102;
    family inet {
      mtu 9000;
      address 10.123.1.4/32 {
        destination 10.123.1.3;
      }
    }
    family iso {
      mtu 9000;
    }
    family mpls {
      mtu 9000;
    }
  }
  unit 3 {
    dlc1 103;
    family inet {
      mtu 9000;
      address 10.123.1.6/32 {
        destination 10.123.1.5;
      }
    }
    family iso {
      mtu 9000;
    }
    family mpls {
      mtu 9000;
    }
  }
}

```

Configuring Cisco HDLC Encapsulation with Byte-Encoding

```
[edit interfaces ds-0/1/0:5:0]
no-keepalives;
encapsulation cisco-hdlc;
ds0-options {
    byte-encoding nx56;
}
unit 0 {
    family inet {
        address 10.221.2.8/24;
    }
}
```

Configuring Cisco HDLC Encapsulation with Byte-Encoding and Framing

```
[edit interfaces ds-0/1/0:5:0]
no-keepalives;
encapsulation cisco-hdlc;
t1-options {
    byte-encoding nx56;
    framing sf;
}
unit 0 {
    family inet {
        address 10.221.2.8/24;
    }
}
```

Use Time Slots 1 Through 10

```
[edit chassis fpc slot-number pic pic-number ct3 port port-number t1 link-number]
channel-group group-number;
timeslots 1-10;
```

Use Time Slots 1 Through 5, 10, and 24

```
[edit chassis fpc slot-number pic pic-number ct3 port port-number t1 link-number]
channel-group group-number;
timeslots 1-5,10,24;
```

Examples: Configuring Channelized DS3-to-DS1 Interfaces

The following configuration is sufficient to get the channelized DS3-to-DS1 interface up and running. The T3 interface can be divided into 28 channels, each at T1 line rate. DS3 channels can use the following encapsulation types for their logical interfaces:

- PPP, PPP CCC, and PPP TCC
- Frame Relay, Frame Relay CCC, and Frame Relay TCC
- Cisco HDLC, Cisco HDLC CCC, and Cisco HDLC TCC

For more information, see “Configuring a Point-to-Point Frame Relay Connection” on page 380.

Configuring Cisco HDLC Encapsulation on a Channelized DS3 Interface

```
[edit interfaces]
t1-2/0/1:20 {
    encapsulation cisco-hdlc;
    unit 0 {
        family inet {
```

```

        address 10.0.4.40/32 {
            destination 10.0.4.41;
        }
    }
}

```

**Configuring PPP
Encapsulation on a
Channelized DS3
Interface**

```

[edit interfaces]
t1-2/0/1:20 {
    encapsulation ppp;
    unit 0 {
        family inet {
            address 10.0.4.40/32 {
                destination 10.0.4.41;
            }
        }
    }
}

```

**Configuring Five Frame
Relay DLCIs on a
Channelized DS3
Interface**

```

[edit interfaces]
t1-5/1/3:0 {
    mtu 9192;
    encapsulation frame-relay;
    unit 1 {
        dlci 101;
        family inet {
            mtu 9000;
            address 10.123.1.2/32 {
                destination 10.123.1.1;
            }
        }
        family iso {
            mtu 9000;
        }
        family mpls {
            mtu 9000;
        }
    }
    unit 2 {
        dlci 102;
        family inet {
            mtu 9000;
            address 10.123.1.4/32 {
                destination 10.123.1.3;
            }
        }
        family iso {
            mtu 9000;
        }
        family mpls {
            mtu 9000;
        }
    }
    unit 3 {
        dlci 103;
    }
}

```

```

        family inet {
            mtu 9000;
            address 10.123.1.6/32 {
                destination 10.123.1.5;
            }
        }
        family iso {
            mtu 9000;
        }
        family mpls {
            mtu 9000;
        }
    }
    unit 4 {
        dlci 104;
        family inet {
            mtu 9000;
            address 10.123.1.8/32 {
                destination 10.123.1.7;
            }
        }
        family iso {
            mtu 9000;
        }
        family mpls {
            mtu 9000;
        }
    }
    unit 5 {
        dlci 105;
        family inet {
            mtu 9000;
            address 10.123.1.10/32 {
                destination 10.123.1.9;
            }
        }
        family iso {
            mtu 9000;
        }
        family mpls {
            mtu 9000;
        }
    }
}

```

**Configuring Cisco HDLC
Encapsulation with
Byte-Encoding**

```

[edit interfaces t1-1/1/0:1]
no-keepalives;
encapsulation cisco-hdlc;
t1-options {
    byte-encoding nx56;
}
unit 0 {
    family inet {
        address 10.221.2.8/24;
    }
}

```

```
}
```

**Configuring Cisco HDLC
Encapsulation with
Byte-Encoding and
Framing**

```
[edit interfaces t1-1/1/0:1]  
no-keepalives;  
encapsulation cisco-hdlc;  
t1-options {  
    byte-encoding nx56;  
    framing sf;  
}  
unit 0 {  
    family inet {  
        address 10.221.2.8/24;  
    }  
}
```


Chapter 23

Configuring Channelized T1 Interfaces

The Channelized T1 intelligent queuing (IQ) and enhanced intelligent queuing (IQE) PICs have 10 T1 ports that you can channelize to the DS0 level. Each T1 interface has 24 DS0 time slots. You can combine DS0 time slots (channels) to create a channel group (NxDS0).

The Channelized T1 IQ and IQE PICs are supported on the M7i, M10i, M20, M40e, M120, and M320 routers.

This section contains the following topics:

- Configuring Channelized T1 IQ and IQE Interfaces on page 495
- Example: Configuring Channelized T1 IQ and IQE Interfaces on page 499

Configuring Channelized T1 IQ and IQE Interfaces

This section describes how to configure channelized T1 IQ and IQE interfaces, discussing the following topics:

- Configuring T1 IQ and IQE Interfaces on page 495
- Configuring Fractional T1 IQ and IQE Interfaces on page 496
- Configuring NxDS0 IQ and IQE Interfaces on page 497
- Configuring Payload Loopback on page 497
- Configuring Channelized T1 Interface Properties on page 499

Configuring T1 IQ and IQE Interfaces

To configure a T1 interface, include the **no-partition** and **interface-type** statements at the [edit interfaces *ct1-fpc/pic/port*] hierarchy level:

```
[edit interfaces ct1-fpc/pic/port]
no-partition interface-type t1;
```

This configuration creates the interface **t1-fpc/pic/port**.



NOTE: For a T1 (t1-) interface configured on channelized T1 (ct1-) interface on a Channelized T1 IQ or IQE PIC, you can configure the following T1 options, but these options do not take effect for the T1 interface:

- bert-algorithm
- bert-error-rate
- bert-period
- buildout
- framing
- line-encoding
- loopback
- remote-loopback-respond

The T1 interface inherits these option settings from the parent channelized T1 interface.

Configuring Fractional T1 IQ and IQE Interfaces

By default, all the time slots on a channelized T1 interface are used. To configure a fractional T1 interface on a Channelized T1 IQ or IQE PIC, perform the following tasks:

1. Include the **no-partition** statement at the [edit interfaces **ct1-fpc/pic/port**] hierarchy level. This configuration creates the interface **t1-fpc/pic/port**.

```
[edit interfaces ct1-fpc/pic/port]
no-partition interface-type t1;
```

2. Configure the number of time slots allocated to the T1 IQ or IQE interface by including the **timeslots** statement at the [edit interfaces **t1-fpc/pic/port t1-options**] hierarchy level. DS0 time slots configured on the channelized T1 IQ or IQE interface are numbered from 1 to 24. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces.

```
[edit interfaces t1-fpc/pic/port t1-options]
timeslots time-slot-range;
```

For more information about T1 time slots, see “Configuring Fractional T1 Time Slots” on page 567.

Example: Configuring Fractional T1 IQ and IQE Interfaces

Configure a fractional T1 interface that uses time slots 2 through 10:

```
[edit interfaces t1-0/0/0]
no-partition interface-type t1;
[edit interfaces t1-0/0/0 t1-options]
```



```
timeslots 1-10;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring NxDS0 IQ and IQE Interfaces

By default, all the time slots on a channelized T1 interface are used. To configure an NxDS0 IQ or IQE interface on a Channelized T1 IQ or IQE PIC, you must configure the number of time slots allocated to the NxDS0 IQ or IQE interface by including the **partition**, **timeslots**, and **interface-type** statements at the [edit interfaces t1-fpc/pic/port] hierarchy level, specifying the ds interface type:

```
[edit interfaces t1-fpc/pic/port]
partition partition-number timeslots time-slot-range interface-type ds;
```

For channelized T1 IQ or IQE interfaces, the partition number range is from 1 through 24.

For channelized T1 IQ or IQE interfaces (t1-fpc/pic/port), the time-slot range is from 1 through 24. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces. For more information about T1 time slots, see “Configuring Fractional T1 Time Slots” on page 567.

Example: Configuring an NxDS0 IQ or IQE Interface

Configure an NxDS0 interface that uses time slots 2 through 10. This configuration creates the ds-0/0/0:1 interface.

```
[edit interfaces t1-0/0/0:1]
partition 1 timeslots 1-10 interface-type ds;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring Payload Loopback

Clocking and loopback options are configured at the controller level for all IQ-based and IQE-based interfaces. However, for the channelized T1 IQ or IQE interfaces, configure the payload loopback on the T1 interfaces instead of the channelized T1 IQ or IQE interface. To configure the payload option, include the **payload** statement at the [edit interfaces t1-fpc/pic/port t1-options loopback] hierarchy level.

By default, all the time slots on a channelized T1 IQ or IQE interface are used. There can be a maximum of 24 channel groups per channelized T1 IQ or IQE interface. Thus, you can configure a maximum of 240 channel groups per PIC.

To specify the DS0 channel group number in the interface name, include a colon (:) as a separator. For example, a Channelized T1 IQ or IQE PIC might have the following physical and virtual interfaces:

```
ds-0/0/0:x
```

x is a DS0 channel group from 1 through 24 (for more information about ranges, see Table 44 on page 498).

You can use any of the values within the range available for x; you do not have to configure the links sequentially. In addition, the JUNOS Software applies the interface options you configure according to the following rules:

- To configure the **t1-options** statement, you must set channel group x to 0:
`ds-0/0/0:0`
- There are no restrictions on configuring the **ds0-options** statement.
- If you delete a configuration you previously committed for channel group 0, the options return to default values.

To configure the channel groups and time slots for a channelized T1 IQ or IQE interface, include the following statements at the **[edit chassis]** hierarchy level:

```
[edit chassis]
fpc slot-number {
  pic pic-number {
    ct1 {
      t1 link-number {
        channel-group group-number;
        timeslots time-slot-range;
      }
    }
  }
}
```

There are 24 time slots on a T1 interface. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces.

Table 44 on page 498 shows the ranges you can specify.

Table 44: Ranges for Channelized T1 IQ Configuration

Item	Option	Range
FPC slot	<i>slot-number</i>	0 through 7
PIC slot	<i>pic-number</i>	0 through 3
T1 port	<i>port-number</i>	0 through 9
DS0 channel group	<i>partition</i>	1 through 24
Time slot	<i>time-slot-range</i>	1 through 24

The theoretical maximum number of channel groups possible per PIC is $10 * 24 = 240$. This is within the maximum bandwidth available.

Configuring Channelized T1 Interface Properties

To configure channelized T1 IQ or IQE interface properties, include the **t1-options** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
t1-options {
  byte-encoding (nx56 | nx64)
  fcs (16 | 32);
  framing (esf | sf);
  idle-cycle-flag (flags | ones);
  invert-data;
  line-encoding (ami | b8zs);
  loopback (local | payload | remote);
  start-end-flag (filler | shared);
}
```



NOTE: If you configure the **line-encoding** statement with the **ami** option and the **byte-encoding** statement with the **nx64** option, excessive zeros in the payload area may bring the interface down. To prevent this, configure the **byte-encoding** statement with the **nx56** option or include the **invert-data** statement.

To specify options for each of the DS0 channels, include the **ds0-options** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
ds0-options {
  byte-encoding (nx56 | nx64);
  fcs (16 | 32);
  idle-cycle-flag (flags | ones);
  loopback payload;
  start-end-flag (filler | shared);
}
```

Only a subset of the T1 options is valid for the channelized configuration; you specify the time slots using the [edit chassis] configuration described in “Examples: Interface Naming” on page 62. For more information about the T1 and DS0 options, see “Configuring T1 Interfaces” on page 559.

Each T1 interface has 24 time slots (DS0s). You can combine one or more of these DS0 time slots (channels) to create a channel group (NxDS0). There can be a maximum of 24 channel groups per T1 interface.

Example: Configuring Channelized T1 IQ and IQE Interfaces

Configure a channelized T1 interface as an unpartitioned, clear channel.

Configuring a T1 Interface

```
[edit interfaces]
ct1-2/0/0 {
  no-partition interface-type t1; # t1-2/0/0
}
```

Configure a partitioned channel group.

Configuring a Channel Group

```
[edit interfaces]
ct1-0/0/1 {
  partition 1 interface-type ds0 timeslots 1-10;
  partition 2 interface-type ds0 timeslots 11-20;
}
```

The following configuration is sufficient to get the channelized T1 IQ or IQE interface up and running:

Configuring Multiple Interface Types

```
[edit]
interfaces {
  ct1-1/2/3 {
    partition 1 timeslots 10 interface-type ds; # ds-1/2/3:1
    partition 2 timeslots 1-9 interface-type ds; # ds-1/2/3:2
  }
  ds-1/2/3:1 {
    unit 0 {
      family inet {
        address 10.25.1.2/24;
      }
    }
  }
  ds-1/2/3:2 {
    unit 0 {
      family inet {
        address 10.25.2.2/24;
      }
    }
  }
}
[edit]
interfaces {
  ct1-1/2/6 {
    no-partition interface-type t1; # t1-1/2/6
  }
  t1-1/2/6 {
    t1-options {
      timeslots 1-2;
    }
    unit 0 {
      family inet {
        address 10.255.126.2/24;
      }
    }
  }
}
```

Chapter 24

Configuring Channelized E1 Interfaces

Each Channelized E1 PIC, Channelized E1 Intelligent Queuing (IQ) PIC and Channelized E1 Enhanced Intelligent Queuing (IQE) PIC has 10 E1 ports that you can channelize to the *NxDS0* level. Each E1 interface has 32 time slots (DS0), in which time slot 0 is reserved. You can combine one or more of these DS0 time slots (channels) to create a channel group *NxDS0*.

This section contains the following topics:

- Configuring Channelized E1 IQ and IQE Interfaces on page 501
- Configuring Channelized E1 Interfaces on page 503
- Example: Configuring Channelized E1 IQ or IQE Interfaces on page 505
- Example: Configuring Channelized E1 Interfaces on page 506

Configuring Channelized E1 IQ and IQE Interfaces

This section describes how to configure channelized E1 IQ and IQE interfaces, discussing the following topics:

- Configuring E1 IQ and IQE Interfaces on page 501
- Configuring Fractional E1 IQ and IQE Interfaces on page 502
- Configuring *NxDS0* IQ and IQE Interfaces on page 502



NOTE: Class-of-service (CoS) rules cannot be applied to an individual channel configured on channelized IQ and IQE interfaces. You can only apply CoS rules to the aggregate bit streams.

Configuring E1 IQ and IQE Interfaces

To configure an E1 interface, include the **no-partition** and **interface-type** statements at the [edit interfaces *ce1-fpc/pic/port*] hierarchy level:

```
[edit interfaces ce1-fpc/pic/port]
no-partition interface-type e1;
```

This configuration creates interface **e1-fpc/pic/port**.

Configuring Fractional E1 IQ and IQE Interfaces

By default, all the time slots on a channelized E1 interface are used. To configure a fractional E1 interface on a Channelized E1 IQ PIC, perform the following tasks:

1. Include the `no-partition` statement at the `[edit interfaces ce1-fpc/pic/port]` hierarchy level:

```
[edit interfaces ce1-fpc/pic/port]
no-partition interface-type e1;
```

This configuration creates interface `e1-fpc/pic/port`.

2. Configure the number of time slots allocated to the E1 IQ or IQE interface by including the `timeslots` statement at the `[edit interfaces e1-fpc/pic/port e1-options]` hierarchy level:

```
[edit interfaces e1-fpc/pic/port e1-options]
timeslots time-slot-range;
```

NxDS0 time slots configured on either a channelized STM1 IQ or IQE interface or a channelized E1 IQ or IQE interface are numbered from 1 to 31 (0 is reserved), while fractional E1 time slots are numbered from 2 to 32 (1 is reserved).

To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces.

For more information about E1 time slots, see “Configuring Fractional E1 Time Slots” on page 548.

Example: Configuring Fractional E1 IQ and IQE Interfaces

Configure a fractional E1 interface that uses time slots 2 through 10:

```
[edit interfaces ce1-0/0/0]
no-partition interface-type e1;
[edit interfaces e1-0/0/0 e1-options]
timeslots 2-10;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring NxDS0 IQ and IQE Interfaces

By default, all the time slots on a channelized E1 interface are used. To configure an NxDS0 IQ interface on a Channelized E1 IQ or IQE PIC, you must configure the number of time slots allocated to the NxDS0 IQ or IQE interface by including the `partition`, `timeslots`, and `interface-type` statements at the `[edit interfaces ce1-fpc/pic/port]` hierarchy level, specifying the `ds` interface type:

```
[edit interfaces ce1-fpc/pic/port]
partition partition-number timeslots time-slot-range interface-type ds;
```

For channelized E1 IQ and IQE interfaces, the partition number range is from 1 through 31.

For E1 IQ and IQE interfaces (*e1-fpc/pic/port*), the time-slot range is from 2 through 31. For channelized E1 IQ and IQE interfaces (*ce1-fpc/pic/port*), the time-slot range is from 1 through 31. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces. For more information about E1 time slots, see “Configuring Fractional E1 Time Slots” on page 548.

Example: Configuring an NxDS0 IQ or IQE Interface

Configure an NxDS0 interface that uses time slots 2 through 10. This configuration creates the *ds-0/0/0:1:1* interface.

```
[edit interfaces ce1-0/0/0:1]
partition 1 timeslots 2-10 interface-type ds;
```

For a full configuration example, see the *JUNOS Feature Guide*.

Configuring Channelized E1 Interfaces

By default, all the time slots on a channelized E1 interface are used. There can be a maximum of 24 channel groups per channelized E1 interface. Thus, you can configure a maximum of 240 channel groups per PIC.

To specify the DS0 channel group number in the interface name, include a colon (:) as a separator. For example, a Channelized E1 PIC might have the following physical and virtual interfaces:

ds-0/0/0:x

where *x* is a DS0 channel group from 0 through 23 (for more information about ranges, see Table 45 on page 504).

You can use any of the values within the range available for *x*; you do not have to configure the links sequentially. In addition, the JUNOS Software applies the interface options you configure according to the following rules:

- To configure the **e1-options** statement, you must set channel group *x* to 0:
ds-0/0/0:0
- There are no restrictions on configuring the **ds0-options** statement.
- If you delete a configuration you previously committed for channel group 0, the options return to default values.

To configure the channel groups and time slots for a channelized E1 interface, include the following statements at the [edit chassis] hierarchy level:

```
[edit chassis]
fpc slot-number {
  pic pic-number {
```

```

ce1 {
  e1 link-number {
    channel-group group-number;
    timeslots time-slot-range;
  }
}

```



NOTE: If you commit the interface name but do not include the [edit chassis] configuration, the Channelized E1 PIC behaves like a standard E1 PIC, and none of the DS0 functionality is accessible.

There are 32 time slots on an E1 interface; however, time slot 0 is reserved. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces.

Table 45 on page 504 shows the ranges you can specify.

Table 45: Ranges for Channelized E1 Configuration

Item	Option	Range
FPC slot	<i>slot-number</i>	0 through 7 (see note below)
PIC slot	<i>pic-number</i>	0 through 3
E1 link	<i>link-number</i>	0 through 9
DS0 channel group	<i>group-number</i>	0 through 23
Time slot	<i>time-slot-range</i>	0 through 31 (with time slot 0 reserved) (see note below)

The theoretical maximum number of channel groups possible per PIC is $10 * 24 = 240$. This is within the maximum bandwidth available.



NOTE: NxDS0 time slots configured on either a channelized STM1 IQ or IQE interface or channelized E1 IQ or IQE interface are numbered from 1 to 31 (0 is reserved), while fractional E1 time slots range from 2 to 32 (1 is reserved).

The FPC slot range depends on the router. For a routing matrix, the range is from 0 through 31. For M40, M40e, M160, M320, M120, and other T Series routers, the range is from 0 through 7. For M20 routers, the range is from 0 through 3. For M10 and M10i routers, the range is from 0 through 1. For M5 and M7i routers, the only applicable value is 0.

Configuring Channelized E1 Interface Properties

To configure channelized E1 interface properties, include the **e1-options** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
e1-options {
  fcs (16 | 32);
  framing (g704 | g704-no-crc4 | unframed);
  idle-cycle-flag (flags | ones);
  loopback (local | remote);
  start-end-flag (filler | shared);
}
```

To specify options for each of the DS0 channels, include the **ds0-options** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
ds0-options {
  byte-encoding (nx56 | nx64);
  fcs (16 | 32);
  idle-cycle-flag (flags | ones);
  loopback payload;
  start-end-flag (filler | shared);
}
```

For DS0 channels on a channelized E1 interface, the **clocking** statement is supported only for channel 0; it is ignored if included in the configuration of channels 1 through 11. The clock source configured for channel 0 applies to all channels on the channelized E1 interface. The individual DS0 channels use a gapped 45-MHz clock as the transmit clock. When you configure the clock source for a channelized interface—**ds-fpc/pic/port:0**, for example—you must also include the **channel-group** statement at the [edit chassis] hierarchy level, and specify channel group 0. For more information, see “Clock Sources on Channelized Interfaces” on page 390.

Only a subset of the E1 options is valid for the channelized configuration; you specify the time slots using the [edit chassis] configuration described in “Examples: Interface Naming” on page 62. For more information about the E1 and DS0 options, see “Configuring E1 Interfaces” on page 543 and “Configuring T1 Interfaces” on page 559.

Each E1 interface has 32 time slots (DS0s), in which time slot 0 is reserved. You can combine one or more of these DS0 time slots (channels) to create a channel group (NxDS0). There can be a maximum of 24 channel groups per E1 interface.

Example: Configuring Channelized E1 IQ or IQE Interfaces

Configure a channelized E1 interface as an unpartitioned, clear channel:

Configuring an E1 Interface	<pre>[edit interfaces] ce1-2/0/0 { no-partition interface-type e1; # e1-2/0/0 }</pre>
------------------------------------	---

The following configuration is sufficient to get the channelized E1 IQ or IQE interface up and running:

**Configuring Multiple
Interface Types**

```
[edit]
interfaces {
  ce1-1/2/3 {
    partition 1 timeslots 10 interface-type ds; # ds-1/2/3:1
    partition 2 timeslots 1-9 interface-type ds; # ds-1/2/3:2
  }
  ds-1/2/3:1 {
    unit 0 {
      family inet {
        address 10.25.1.2/24;
      }
    }
  }
  ds-1/2/3:2 {
    unit 0 {
      family inet {
        address 10.25.2.2/24;
      }
    }
  }
}
[edit]
interfaces {
  ce1-1/2/6 {
    no-partition interface-type e1; # e1-1/2/6
  }
  e1-1/2/6 {
    e1-options {
      timeslots 1-2;
    }
    unit 0 {
      family inet {
        address 10.255.126.2/24;
      }
    }
  }
}
```

Example: Configuring Channelized E1 Interfaces

The following configuration is sufficient to get the channelized E1 interface up and running:

**Configuring an E1
Interface, E1 Options,
and DSO Options**

```
[edit chassis]
fpc 0 {
  pic 1 {
    ce1 {
      e1 0 {
        channel-group 0 timeslots 1;
        channel-group 1 timeslots 2;
        channel-group 5 timeslots 5-7;
      }
    }
  }
}
```

```

        e1 4 {
            channel-group 10 timeslots 11,17,28-31;
        }
    }
}
[edit interfaces ds-0/1/0:0]
e1-options {
    fcs 32;
    framing g704-non-grc;
    loopback remote;
}
[edit interfaces ds-0/1/4:10]
ds0-options {
    byte-encoding nx56;
    start-end-flag filler;
}

```

The above configuration results in the following interfaces:

```

ds-0/1/0:1, with time slot 1 allocated
ds-0/1/0:5, with time slots 5 through 7 allocated
ds-0/1/4:10, with time slots 11, 17, and 28 through 31 allocated

```

The remaining ports (other than 0 and 4) remain as regular E1 interfaces (and follow the e1-0/1/x naming convention).

```

[edit chassis]
fpc 0 {
    pic 1 {
        ce1 {
            e1 0 {
                channel-group 1 timeslots 1;
                channel-group 5 timeslots 5-7;
            }
            e1 4 {
                channel-group 10 timeslots 11,17, 28-31;
            }
        }
    }
}

```

Use Time Slots 1 Through 10	<code>[edit chassis fpc slot-number pic pic-number ce1 e1 link-number]</code> channel-group <i>group-number</i> ; timeslots 1-10;
Use Time Slots 1 Through 5, 10, and 24	<code>[edit chassis fpc slot-number pic pic-number ce1 e1 link-number]</code> channel-group <i>group-number</i> ; timeslots 1-5,10,24;

Chapter 25

Configuring Channelized E1 PRI and T1 PRI Interfaces

This section contains the following topics:

- Channelized E1 PRI and T1 PRI Overview on page 509
- Configuring a Clear Channel on a Dual-Port Channelized T1-E1 PIM on page 510
- Configuring a Channelized T1/E1 Interface to Drop and Insert Time Slots on page 510
- Configuring Primary Rate Interfaces on page 512
- Allocating B-Channels for Dialout on page 513
- Configuring PRI Interfaces on page 513
- Example: Configuring a Channelized T1 Interface as Primary Rate Interface on page 514

Channelized E1 PRI and T1 PRI Overview

J Series Services Routers equipped with a Dual-Port Channelized T1/E1 PIM support Integrated Services Digital Network (ISDN) Primary Rate Interfaces (PRIs). ISDN PRI, referred to as S2M in Europe, is the “primary” extended ISDN network interface. It offers a larger capacity of digital channels utilizing a variety of improved mediums, and is used by large organizations with intensive communication needs. In contrast, the ISDN Basic Rate Interface (BRI), known as SO in Europe, provides a limited number of channels, transmitting over copper wire, and is used by smaller organizations or individuals with less intensive communication needs. For more information about configuring ISDN BRI interfaces, see “Configuring ISDN Physical Interface Properties” on page 821.

Unlike channelized PICs on the M Series and T Series routers, the interface type on the Dual-Port Channelized T1/E1 PIM is configurable. A single interface can operate as either a channelized T1 or channelized E1 interface (or clear channel) or as an ISDN PRI. The ISDN PRI channels can operate on the same interface as T1 or E1 channels. The PIM also supports a “drop-and-insert” feature, allowing you to insert channels from one port on the PIM into the other port on the PIM.

These ISDN channels are delivered to the user in one of two predefined configurations:

- ISDN BRI is configured by specifying properties for a physical (**br-**) interface and a logical (**dln**) interface.

- For ISDN PRI, you configure:
 1. Either a channelized E1 (**ce1-pim/0/port**) or channelized T1 (**ct1-pim/0/port**) interface.
 2. Time slots within a **ce1-pim/0/port** interface or **ct1-pim/0/port** interface.
 3. A bearer (B) channel **bc-pim/0/port:channel** interface for each time slot that you want to function as an ISDN PRI B-channel. The B-channel is used for data, video, voice, and multimedia. You can create up to 30 B-channels on a channelized E1 interface, and 23 B-channels on a channelized T1 interface.
 4. One delta (D) channel, used between switching equipment in the ISDN network and the ISDN equipment at your site for signaling. For channelized E1, the D-channel must be time slot 16. For channelized T1, the D-channel must be time slot 24.



NOTE: Time slots can also be shared with **ds-pim/0/port** time slots within the same channelized interface.

Channelized E1 and T1 PIMs on J Series routers provide support for ISDN PRI connectivity for dial-in and callback, and for use as primary or backup network connections.

Configuring a Clear Channel on a Dual-Port Channelized T1-E1 PIM

A *clear channel* is an interface that uses the entire bandwidth of the port on a PIM. To configure a clear channel, include the **no-partition** and **interface-type** statements in the configuration. On a Dual-Port Channelized T1-E1 PIM, you can configure two clear-channel interfaces.

To configure an E1 interface, include the **no-partition** and **interface-type** statements at the [edit interfaces **ce1-pim/0/port**] hierarchy level:

```
[edit interfaces ce1-pim/0/port]
no-partition interface-type e1;
```

This configuration creates interface **e1-pim/0/port**.

To configure a T1 interface, include the **no-partition** and **interface-type** statements at the [edit interfaces **ct1-pim/0/port**] hierarchy level:

```
[edit interfaces ct1-pim/0/port]
no-partition interface-type t1;
```

This configuration creates interface **t1-pim/0/port**.

Configuring a Channelized T1/E1 Interface to Drop and Insert Time Slots

On channelized T1/E1 interfaces configured for channelized operation, you can insert channels (time slots) from one port (for example, channels carrying voice) directly

into the other port on the PIM, to replace channels coming through the Routing Engine. This feature, known as drop and insert, allows you to integrate voice and data on a single T1 or E1 link by removing the DS0 time slots of one T1 or E1 port and replacing them by inserting the time slots of another T1 or E1 port. It is not necessary to use the same time slots on both interfaces, but the time slots count must be same. The channels that are not configured for the drop-and-insert feature are used for normal traffic.

You can configure:

- 30 channelized E1 time slots, with the 16th time slot operating as the signaling channel
- 23 channelized T1 time slots, with the 24th time slot operating as the signaling channel

The signaling channel, or D-channel, must be part of the channels that are being switched through the drop-and-insert functionality. The JUNOS Software does not support switching of voice and data between ports by default.

Both ports involved in the drop-and-insert configuration must use the same clock source—either the router's internal clock or an external clock.

The following clock source settings are valid:

- When port 0 is set to use the internal clock, port 1 must also be set to use it, and vice versa.
- When port 0 is set to use its external clock, port 1 must be set to run on the same clock—the external clock for port 0.
- When port 1 is set to use its external clock, port 0 must be set to run on the same clock—the external clock for port 1.

For more details about valid clock combinations, see the *J-series Services Router Basic LAN and WAN Access Configuration Guide*.

To configure drop-and-insert time slots on a channelized T1 interface, include the **partition** statement at the `[edit interfaces ct1-pim/0/port]` hierarchy level with the **timeslots** statement and **interface-type** statements specified:

```
[edit interfaces]
ct1-pim/0/port {
  partition 1 timeslots 1-10 interface-type ds;
  partition 2 timeslots 11-14 interface-type ds;
  partition 3 timeslots 15-32 interface-type ds;
}
```

This configuration creates interfaces `ds-pim/0/port:1`, `ds-pim/0/port:2`, and `ds-pim/0/port:3`.

Use the same configuration to create drop-and insert time slots on a channelized E1 interface by including the **partition** statement and options at the `[edit interfaces ce1-pim/0/port]` hierarchy level.

Configuring Primary Rate Interfaces

Primary rate interfaces are a combination of B-channels with one controlling D-channel for the group. Configure B-channel interfaces for each time slot that you want to function as an ISDN PRI interface. The B-channel is used for data, video, voice, and multimedia. You can create:

- 23 B-channels on a channelized T1 interface
- 30 B-channels on a channelized E1 interface

To configure B-channels on a channelized T1 interface, include the **partition** statement at the `[edit interfaces ct1-pim/0/port]` hierarchy level with the **timeslots** statement and **interface-type bc** specified:

```
[edit interfaces]
ct1-pim/0/port {
    partition 1-23 timeslots 1-23 interface-type bc;
}
```

This configuration creates interfaces `bc-pim/0/port:1` through `bc-pim/0/port:1`, and `ds-pim/0/port:3`.

Use the same configuration to create B-channels on a channelized E1 interface by including the **partition** statement and options at the `[edit interfaces ce1-pim/0/port]` hierarchy level.

One D-channel is used between switching equipment in the ISDN network and the ISDN equipment at your site for signaling. For channelized E1, the D-channel must be time slot 16. For channelized T1, the D-channel must be time slot 24.

To configure a D-channel on a channelized T1 interface, include the **partition** statement at the `[edit interfaces ct1-pim/0/port]` hierarchy level with the **timeslots** statement and **interface-type dc** specified:

```
[edit interfaces]
ct1-pim/0/port {
    partition 24 timeslots 24 interface-type dc;
}
```

This configuration creates interfaces `dc-pim/0/port`.

Use the same configuration to create B-channels on a channelized E1 interface by including the **partition** statement and options at the `[edit interfaces ce1-pim/0/port]` hierarchy level.

```
[edit interfaces]
ce1-pim/0/port {
    partition 16 timeslots 16 interface-type dc;
}
```

To view PRI or ISDN options information about interface, use the following operational mode commands supporting BRI interfaces:

- `show interfaces interface-name detail`
- `show interface dln`
- `show isdn calls`
- `show isdn history`
- `show isdn q921 statistics`
- `show isdn q931 statistics`
- `show isdn status`



NOTE: You must configure a D-channel and B-channels to complete your ISDN PRI line configuration.



NOTE: You can configure dso-options on the B-channel, but you cannot configure parameters for a D-channel. However, when interface statistics are displayed, both B-channel and D-channel interfaces have statistical values.

Allocating B-Channels for Dialout

You can configure the system to allocate B-channels for dialout from lowest or highest numbered B-channel (ascending or descending order). By configuring this feature, you reduce chances of “glare” on PRI lines carrying a mix of incoming and outgoing calls.

To configure the B-channel allocation, include the `idsn-options` and `bchannel-allocation` statements at the `[edit interfaces ct1-pim/0/port | ce1-pim/0/port]` hierarchy level:

```
[edit interfaces]
(ct1-pim/0/port | ce1-pim/0/port) {
  isdn-options {
    (bchannel-allocation (ascending | descending);
  }
}
```

Configuring PRI Interfaces

When you create a PRI from a channelized E1 or channelized T1 interface, you can select all the slots for the PRI, or just a few of them, leaving the rest as **ds-** interfaces.

To configure a PRI from a channelized T1 interface, include the `partition` statement at the `[edit interfaces ct1-pim/0/port]` hierarchy level with the `timeslots` statement and `interface-type bc` specified:

```
[edit interfaces]
ct1-pim/0/port {
  partition 1 timeslots 1-10 interface-type ds;
```

```
    partition 2 timeslots 11-24 interface-type pr;
}
```

This configuration creates interfaces `ds-pim/0/port:1` through `pr-pim/0/port:2`.

Use the same configuration to create interfaces on a channelized E1 interface by including the `partition` statement and options at the `[edit interfaces ce1-pim/0/port]` hierarchy level.

To configure channelized E1 interface properties, include the `e1-options` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]
e1-options {
  fcs (16 | 32);
  framing (g704 | g704-no-crc4 | unframed);
  idle-cycle-flag (flags | ones);
  loopback (local | remote);
  start-end-flag (filler | shared);
}
```

To specify options for each of the DS0 channels, include the `ds0-options` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]
ds0-options {
  byte-encoding (nx56 | nx64);
  fcs (16 | 32);
  idle-cycle-flag (flags | ones);
  loopback payload;
  start-end-flag (filler | shared);
}
```

Example: Configuring a Channelized T1 Interface as Primary Rate Interface

Configure a channelized T1 interface to operate fully as a PRI:

```
[edit interfaces]
ct1-2/0/0 {
  partition 1-23 timeslots 1-23 interface-type bc;
  partition 24 timeslots 24 interface-type dc;
  t1-options {
    line-encoding b8zs;
    framing esf;
  }
  traceoptions {
    flag q931;
    flag q921;
    file {
      pri_trace_log;
    }
  }
  dialer-options {
    pool 1 priority 25;
  }
}
```

```

    }
    isdn-options {
        switch-type att5e;
        bchannel-allocation descending;
        incoming-called-number 384101;
        incoming-called-number 384102;
        incoming-called-number 384103;
    }
}

[edit interfaces]
d10 {
    unit 0 {
        dialer-options {
            pool 1;
            dial-string 384010;
            incoming-map {
                accept-all;
            }
        }
        family inet {
            filter {
                dialer int-packet;
            }
            address 13.1.1.2/24;
        }
    }
}

[edit firewall]
family inet {
    dialer-filter int-packet {
        term term1 {
            from {
                destination address {
                    13.1.1.1/24;
                }
                protocol icmp;
                then note;
            }
        }
        term term2 {
            then ignore;
        }
    }
}

```


Part 8

Configuring Circuit Emulation PICs

- Circuit Emulation PICs Overview on page 519
- Configuring SAToP Support on Circuit Emulation PICs on page 523
- Configuring ATM Support on Circuit Emulation PICs on page 529

Chapter 26

Circuit Emulation PICs Overview

This chapter includes the following sections:

- Mobile Backhaul and Circuit Emulation Overview on page 519
- Mobile Backhaul Application Overview on page 519
- Circuit Emulation PIC Types on page 520
- Circuit Emulation PICs Clocking Features on page 520
- T1 and E1 Options Exceptions on Circuit Emulation PICs on page 521
- Displaying Information About Circuit Emulation PICs on page 522

Mobile Backhaul and Circuit Emulation Overview

This section provides an overview of the Circuit Emulation PICs, describes their main features, and indicates which routers support them.

Juniper Networks mobile backhaul (IP/MPLS) solutions provide the following benefits:

- Flexibility to support converged networks that accommodate both IP and legacy services (leveraging proven circuit emulation techniques).
- Scalability to support emerging data intensive technologies.
- Cost-effectiveness to compensate for rising levels of backhaul traffic.


M7i, M10i, M40e, M120, and M320 routers with 12-port T1/E1 CE and 4-port Channelized OC3/STM1 CE circuit emulation (CE) interfaces offer IP/MPLS-based mobile backhaul solutions that enable operators to combine diverse transport technologies onto a single transport architecture, to reduce operating costs while enhancing user features and increasing profits. This architecture accommodates the backhaul of legacy services, emerging IP-based services, location-based services, mobile gaming and mobile TV, and new emerging technologies such as LTE and WiMAX.

Mobile Backhaul Application Overview

This section provides an application example (see Figure 50 on page 520) based on the mobile backhaul reference model where Customer Edge 1 (CE1) is a base station controller (BSC), Provider Edge 1 (PE1) is a cell site router, PE2 is an M Series (aggregation) router, and CE2 is a BSC and Radio Network Controller (RNC). The Internet Engineering Task Force (RFC 3895) describes pseudowire as “a mechanism

that emulates the essential attributes of a telecommunications service (such as a T1 leased line or Frame Relay) over a PSN” (Packet Switching Network).

Figure 50: Mobile Backhaul Application



Circuit Emulation PIC Types

The following Circuit Emulation PICs are specifically designed for mobile backhaul applications:

- Four-Port Channelized OC3/STM1 Circuit Emulation PIC on page 520
- Twelve-Port T1/E1 Circuit Emulation PIC on page 520

Four-Port Channelized OC3/STM1 Circuit Emulation PIC

The four-port Channelized OC3/STM1 CE PIC allows each of its four ports to be independently configured to either SONET or SDH framing mode, and supports mixed SAToP and ATM interfaces on any port. In SONET mode, each OC3 port can be channelized down to three coc1 channels, and then each coc1 can channel down to 28 T1 channels. In SDH mode, each STM1 port can be channelized down to four cau4 channels, and then each cau4 can channel down to 63 E1 channels. The T1/E1 channels support time-division multiplexing (TDM) interfaces using the Structure-Agnostic time-division multiplexing over Packet (SAToP) protocol [RFC 4553] encapsulation, and support T1/E1 and SONET clocking features. Mixing T1 and E1 channels is not supported on individual ports.

Twelve-Port T1/E1 Circuit Emulation PIC

The twelve-port Channelized T1/E1 Circuit Emulation PIC supports time-division multiplexing (TDM) interfaces using the Structure-Agnostic time-division multiplexing over Packet (SAToP) protocol [RFC 4553] encapsulation, and supports T1/E1 and SONET clocking features. The 12-port Channelized T1/E1 Circuit Emulation PIC can be configured to work as either 12 T1s or 12 E1s. Mixing T1s and E1s is not supported.

Circuit Emulation PICs Clocking Features

All Circuit Emulation PICs support the following clocking features:

- External clocking—Also known as “loop timing”. Clock is distributed via the TDM interfaces.
- Internal clocking with external synchronization—Also known as “external timing” or “external synchronization”.

- Internal clocking with PIC-level line synchronization—Synchronizing the PIC's internal clock with a clock recovered from a TDM interface local to the PIC.

This feature set is useful for aggregation in mobile backhaul applications.



NOTE: The PRS of the clock recovered from one interface may not be the same as that of another TDM interface. There is a limitation on the number of timing domains that can be supported in practice.

T1 and E1 Options Exceptions on Circuit Emulation PICs

This section contains the following topics:

- T1 and E1 Options Exceptions on 12-Port T1/E1 Circuit Emulation PICs on page 521
- T1 and E1 Options Exceptions on 4-Port Channelized OC3/STM1 Circuit Emulation PICs on page 522

T1 and E1 Options Exceptions on 12-Port T1/E1 Circuit Emulation PICs

Twelve-port T1/E1 Circuit Emulation PICs support T1 and E1 options with the following exceptions:

- `bert-algorithm`, `bert-error-rate`, and `bert-period` options are supported for CT1 or CE1 configurations only.
- `framing` is supported for CT1 or CE1 configurations only. It is not applicable in SAToP configurations.
- `buildout` is supported in CT1 configurations only.
- `line-encoding` is supported in CT1 configurations only.
- `loopback local` and `loopback remote` are supported in CE1 and CT1 configurations only.
- `loopback payload` is not supported. It is not applicable in SAToP configurations.
- `idle-cycle-flag` is not supported. It is not applicable in SAToP or ATM configurations.
- `start-end-flag` is not supported. It is not applicable in SAToP or ATM configurations.
- `invert-data` is not supported. It is not applicable in SAToP configurations.
- `fcs32` is not supported. `fcs` is not applicable in SAToP or ATM configurations.
- `timeslots` is not supported. It is not applicable in SAToP configurations.
- `byte-encoding nx56` is not supported. It is not applicable in SAToP or ATM configurations.
- `crc-major-alarm-threshold` and `crc-minor-alarm-threshold` are not supported.
- `remote-loopback-respond` is not supported. It is not applicable in SAToP configurations.

T1 and E1 Options Exceptions on 4-Port Channelized OC3/STM1 Circuit Emulation PICs

Four-port Channelized OC3/STM1 Circuit Emulation PICs support T1 and E1 options with the following exceptions:

- **bert-algorithm**, **bert-error-rate**, and **bert-period** options are supported for CT1 or CE1 configurations only.
- **framing** is supported for CT1 or CE1 configurations only. It is not applicable in SAToP configurations.
- **buildout** is supported in CT1 configurations only.
- **line-encoding** is supported in CT1 configurations only.
- **loopback local** and **loopback remote** are supported in CE1 and CT1 configurations only. By default, no loopback is configured.
- **loopback payload** is not supported. It is not applicable in SAToP configurations.
- **idle-cycle-flag** is not supported. It is not applicable in SAToP configurations.
- **start-end-flag** is not supported. It is not applicable in SAToP configurations.
- **invert-data** is not supported. It is not applicable in SAToP configurations.
- **fcs16** is not supported in E1 and T1 configurations only.
- **fcs32** is not supported in E1 and T1 configurations only. It is not applicable in SAToP configurations.
- **timeslots** is not supported. It is not applicable in SAToP or ATM configurations.
- **byte-encoding** is not supported in T1 configurations only. It is not applicable in SAToP configurations. **nx56** byte encoding is not supported.
- **crc-major-alarm-threshold** and **crc-minor-alarm-threshold** are T1 options supported in SAToP configurations only.
- **remote-loopback-respond** is not supported. It is not applicable in SAToP configurations.

Displaying Information About Circuit Emulation PICs

Use the CLI **show chassis hardware** command to display information about the PIC configuration.

- For a T1 CE PIC configuration, the output designation is **T1 CE**.
- For an E1 CE PIC configuration, the output designation is **E1 CE**.
- For a COC3 CE PIC configuration, the output designation is **COC1 CE**.
- For a CSTM1 CE PIC configuration, the output designation is **CSTM1 CE**.

Chapter 27

Configuring SAToP Support on Circuit Emulation PICs

This chapter includes information on configuring SAToP support on Circuit Emulation PICs and it includes the following sections:

- Configuring SAToP on 4-port Channelized OC3/STM1 CE PICs on page 523
- Configuring SAToP Emulation on T1/E1 Interfaces on CE PICs on page 525

Configuring SAToP on 4-port Channelized OC3/STM1 CE PICs

This configuration example applies to the mobile backhaul application shown in Figure 50 on page 520.

Configuring SONET/SDH Framing Mode at the PIC Level

To set the framing mode at the PIC level, for all four ports on the PIC, include the framing statement at the [edit chassis fpc fpc-slot pic pic-slot] hierarchy level:

```
[edit chassis fpc fpc-slot pic pic-slot]  
framing (sonet | sdh); # SONET for COC3 or SDH for CSTM1
```

After a PIC is brought online, interfaces are created for the PIC's available ports according to the PIC type and the framing option used:

- If you include the framing **sonet** statement (for a COC3 CE PIC), four COC3 interfaces are created.
- If you include the framing **sdh** statement (for a CSTM1 CE PIC), four CSTM1 interfaces are created.
- If you do not specify framing at the PIC level, then the default framing is SONET for all four ports.



NOTE: If you set the framing option incorrectly for the PIC type, the commit operation fails.

Configuring SONET/SDH Framing Mode at the Port Level

Each port's framing mode can be configured individually, as either COC3 (SONET) or STM1 (SDH). Ports not configured for framing retain the PIC framing configuration, which is SONET by default; if you have not specified framing at the PIC level. To set the framing mode for individual ports, include the **framing** statement at the [edit chassis fpc fpc-slot pic pic-slot port port-number] hierarchy level:

```
[edit chassis fpc fpc-slot pic pic-slot port port-number]
framing (sonet | sdh); # SONET for COC3 or SDH for CSTM1
```

Configuring the framing mode at the port level overwrites the previous PIC level framing mode configuration for the specified port. Subsequently configuring the PIC level framing mode overwrites the port level framing configuration. For example, if you would like three STM1s and one COC3; then it would be practical to first configure the PIC for SDH framing and then configure one port for SONET framing.

Configuring COC3 Ports Down to T1 Channels

On any port configured for SONET framing (numbered 0 through 3), you can configure three COC1 channels (numbered 1 through 3). On each COC1 channel, you can configure 28 T1 channels (numbered 1 through 28). To configure COC3 channelization down to COC1 and then down to T1 channels, include the **partition** statement at the [edit interfaces (coc1 | coc3)-fpc/pic/port] hierarchy level as in the following example:

```
[edit interfaces]
coc3-1/0/0 {
  partition 1 oc-slice 1 interface-type coc1;
}
coc1-1/0/0:1 {
  partition 1 interface-type t1;
}
```

After you partition the T1 channels, configure the SAToP options on them in the same way as for T1 interfaces. See “Setting the SAToP Options” on page 526.

Configuring CSTM1 Ports Down to E1 Channels

On any port configured for SDH framing (numbered 0 through 3), you can configure one CAU4 channel. On each CAU4 channel, you can configure 63 T1 channels (numbered 1 through 63). To configure CSTM1 channelization down to CAU4 and then down to E1 channels, include statements for the various interface types at the [edit interfaces] hierarchy level as in the following example:

```
[edit interfaces]
cstm1-1/0/1 {
  no-partition interface-type cau4;
}
cau4-1/0/1 {
  partition 1 interface-type e1;
}
e1-1/0/1:1 {
```

```
encapsulation satop;
unit 0;
}
```

After you configure the E1 channels, configure SAToP options on them in the same way as for E1 interfaces. See “Setting the SAToP Options” on page 526.

Configuring SAToP Emulation on T1/E1 Interfaces on CE PICs

This configuration example applies to the mobile backhaul application shown in Figure 50 on page 520.

- Setting the Emulation Mode on page 525
- Configuring SAToP Emulation on T1/E1 Interfaces on page 525

Setting the Emulation Mode

To set the framing emulation mode, include the **framing** statement at the [edit chassis fpc fpc-slot pic pic-slot] hierarchy level:

```
[edit chassis fpc fpc-slot pic pic-slot]
framing (t1 | e1);
```

After a PIC is brought online, interfaces are created for the PIC’s available ports according to the PIC type and the framing option used:

- If you include the **framing t1** statement (for a T1 CE PIC), 12 CT1 interfaces are created.
- If you include the **framing e1** statement (for an E1 CE PIC), 12 CE1 interfaces are created.



NOTE: If you set the **framing** option incorrectly for the PIC type, the commit operation fails.

CE PICs with SONET and SDH ports require prior channelization down to T1 or E1 before configuring. Only T1/E1 channels support SAToP encapsulation or SAToP options.

Configuring SAToP Emulation on T1/E1 Interfaces

The following topics are covered in this section:

- Setting the Encapsulation Mode on page 526
- T1/E1 Loopback Support on page 526
- T1 FDL Support on page 526
- Setting the SAToP Options on page 526
- Pseudowire Interface Configuration on page 527

Setting the Encapsulation Mode

T1/E1 channels on CE PICs can be configured with SAToP encapsulation at the PE router, as follows:

```
[edit interfaces (t1|e1)-fpc/pic/port]
encapsulation satop;
unit logical-unit-number;
```

You do not need to configure any cross-connect circuit family because it is automatically created for the above encapsulation.

T1/E1 Loopback Support

Use the CLI to configure remote and local loopback as T1 (CT1) or E1 (CE1). By default, no loopback is configured. See “Configuring T1 Loopback Capability” on page 565 and “Configuring E1 Loopback Capability” on page 547.

T1 FDL Support

If T1 is used for SAToP, the T1 facility data-link (FDL) loop is NOT supported on the CT1 interface device because SAToP does not analyze T1 framing bits.

Setting the SAToP Options

You can configure the following SAToP options:

- **groups**—Specify groups.
- **excessive-packet-loss-rate**—Set packet loss options.
- **idle-pattern**—An 8-bit hexadecimal pattern to replace TDM data in a lost packet (from 0 to 255).
- **jitter-buffer-auto-adjust**—Automatically adjust the jitter buffer.
- **jitter-buffer-latency**—Number of milliseconds delay in the jitter buffer (from 1 to 1000 milliseconds).
- **jitter-buffer-packets**—Number of packets in the jitter buffer (from 1 to 64 packets).
- **payload-size**—Configure the payload size, in bytes (from 32 to 1024 bytes).
- **sample-period**—Number of milliseconds over which excessive packet loss rate is calculated.
- **threshold**—Percentile designating the threshold of excessive packet loss rate (from 1 to 100 percent).

The following example shows the SAToP configuration options:

```
[edit interfaces (t1|e1)-fpc/pic/port]
satop-options {
  excessive-packet-loss-rate {
    groups group-names;
```

```

        sample-period milliseconds;
        threshold percentile;
    )
    idle-pattern pattern;
    jitter-buffer-auto-adjust;
    jitter-buffer-latency milliseconds;
    jitter-buffer-packets packets;
    payload-size bytes;
}

```

Pseudowire Interface Configuration

Configuration for the TDM pseudowire at the PE uses the existing Layer 2 circuit infrastructure:

```

[edit protocols l2circuit]
neighbor address {
    interface t1-fpc/pic/port.0 {
        virtual-circuit-id 1;
    }
}

```

After the CE-bound interfaces (for both PEs) are configured with proper encapsulation, payload size, and other parameters, the two PEs try to establish a pseudowire with pseudowire emulation edge-to-edge (PWE3) signaling extensions. The following pseudowire interface configurations are disabled or ignored for TDM pseudowires:

- ignore-encapsulation
- mtu

The supported pseudowire types are:

- 0x0011 Structure-Agnostic E1 over Packet
- 0x0012 Structure-Agnostic T1 (DS1) over Packet

When the local interface parameters match the received parameters, and the pseudowire type and control world bit are equal, the pseudowire is established.

For detailed information about configuring TDM pseudowire, see the *JUNOS VPNs Configuration Guide*.

For detailed information about PICs, see the *PIC Guide* for your router.

Chapter 28

Configuring ATM Support on Circuit Emulation PICs

This chapter includes information on configuring ATM support on Circuit Emulation PICs and it includes the following sections:

- Overview of ATM Support on Circuit Emulation PICs on page 529
- Configuring the 12-Port Channelized T1/E1 CE PIC Operating Mode on page 530
- Configuring the 4-Port Channelized COC3/STM1 CE PIC Operating Mode on page 532
- Configuring ATM Pseudowires on page 534
- ATM OAM on page 536
- Scaling on page 537
- Congestion Control on page 537
- QoS/Shaping on page 537
- Configuring the PIC Type on page 537
- Configuring Layer 2 Circuit and Layer 2 VPN Pseudowires on page 537
- Supported Interface Configurations on page 538
- ATM Limitations on page 539

Overview of ATM Support on Circuit Emulation PICs

M7i, M10i, and M40e routers with 4-port COC3/CSTM1 CE PIC and 12-port T1/E1 CE PIC support ATM over MPLS (RFC 4717) and packet encapsulations (RFC 2684). CE PIC ATM configuration and behavior is consistent with existing ATM2 PICs.

The following protocols are supported:

- ATM over MPLS (RFC 4717)
- ATM via dynamic labels (LDP, RSVP-TE)

ATM OAM support:

- Generation and monitoring of F4 and F5 OAM cells
- Generation and monitoring of end-to-end cells of type AIS and RDI

- Monitor and terminate loopback cells
- Supports OAM on each VP and VC simultaneously

The following protocols are not supported:

- QoS or COS queues. All VCs are unspecified bit rate (UBR).
- NxDS0 grooming.

The following ATM2 encapsulations are not supported:

- `atm-cisco-nlpid`—Cisco-compatible ATM NLPID encapsulation
- `atm-mlppp-llc`—ATM MLPPP over AAL5/LLC
- `atm-nlpid`—ATM NLPID encapsulation
- `atm-ppp-llc`—ATM PPP over AAL5/LLC
- `atm-ppp-vc-mux`—ATM PPP over raw AAL5
- `atm-snap`—ATM LLC/SNAP encapsulation
- `atm-tcc-snap`—ATM LLC/SNAP for translational cross-connect
- `atm-tcc-vc-mux`—ATM VC for translational cross-connect
- `vlan-vci-ccc`—CCC for VLAN Q-in-Q and ATM VPI/VCI interworking
- `atm-vc-mux`—ATM VC multiplexing
- `ether-over-atm-llc`—Ethernet over ATM (LLC/SNAP) encapsulation
- `ether-vpls-over-atm-llc`—Ethernet VPLS over ATM (bridging) encapsulation

Configuring the 12-Port Channelized T1/E1 CE PIC Operating Mode

This section contains the following topics:

- T1/E1 Mode Selection on page 530
- 12-Port Channelized T1/E1 CE PIC Configuration Statements on page 531

T1/E1 Mode Selection

All ATM interfaces are either T1 or E1 channels within the COC3/CSTM1 hierarchy. Each COC3 can be partitioned as three COC1 slices, each of which in turn can be partitioned further into 28 ATM interfaces the size of a T1. Each CSMT1 can be partitioned as 1 CAU4 which can be further partitioned as 84 ATM interfaces the size of an E1. All ports operate by default as T1s.

To configure the PIC mode, use the following `set` command, selecting either the T1 or E1 option:

```
set chassis fpc fpc-slot pic pic-slot framing (t1 | e1)
```

After the PIC is brought online, 12 ct1 interfaces or 12 ce1 interfaces are created, depending on the T1 or E1 mode selection of the PIC.

Figure 51 on page 531 and Figure 52 on page 531 illustrate the possible interfaces that can be created on the 12-port T1/E1 CE PIC:

Figure 51: 12-Port T1/E1 CE PIC Possible Interfaces (T1 Size)





Figure 52: 12-Port T1/E1 CE PIC Possible Interfaces (E1 Size)



12-Port Channelized T1/E1 CE PIC Configuration Statements

Setting the T1/E1 Mode at the PIC Level

To set the T1/E1 mode at the PIC level, enter the following command:

```
set chassis fpc fpc-slot pic pic-slot framing t1|e1
```

Or specify the following:

```
chassis {
  fpc fpc-slot {
    pic pic-slot {
      framing <t1 | e1>;
    }
  }
}
```

After the PIC is brought online, 12 ct1 interfaces or 12 ce1 interfaces are created.

If the mode is not manually configured, then the PIC defaults to T1.

Creating an ATM Interface on a CT1 or CE1

To create an ATM interface on a CT1, enter the following command:

```
set interfaces ct1-x/y/z no-partition interface-type at
```

Or specify the following:

```
interfaces {
  ct1-x/y/z {
    no-partition {
      interface-type <at>;
    }
  }
}
```

To create an ATM interface on a CE1:

```
set interfaces ce1-x/y/z no-partition interface-type at
```

Or specify the following:

```
interfaces {
  ce1-x/y/z {
    no-partition {
      interface-type <at>;
    }
  }
}
```

The interface `at-x/y/z` is created.

You can use the `show chassis hardware` command to display a list of the installed PICs.

Configuring the 4-Port Channelized COC3/STM1 CE PIC Operating Mode

This section contains the following topics:

- T1/E1 Mode Selection on page 532
- Configuring a Port for SONET or SDH Mode on a 4-Port Channelized COC3/STM1 CE PIC on page 533
- Configuring an ATM Interface on a COC1 on page 534


T1/E1 Mode Selection

All ATM interfaces are either T1 or E1 channels within the COC3/CSTM1 hierarchy. Each COC3 can be partitioned as three COC1 slices, each of which in turn can be partitioned further into 28 ATM interfaces and the size of each interface created is that of a T1. Each CSMT1 can be portioned as 1 CAU4, which can be further partitioned as 84 E1 sized ATM interfaces.


To configure the T1/E1 mode selection:

1. To create `coc3-x/y/z` or `cstm1-x/y/z` interfaces, `chassisd` will look for configuration at the `[edit chassis fpc x pic y port z framing (<sonet> | <sdh>)]` hierarchy level. If the `sdh` option is specified, `chassisd` will create a `cstm1-x/y/z` interface. Otherwise, `chassisd` will create `coc3-x/y/z` interfaces.
2. Only interface `coc1` can be created from `coc3`, and `t1` can be created from `coc1`.
3. Only interface `cau4` can be created from `cstm1`, and `e1` can be created from `cau4`.

Figure 53 on page 533 and Figure 54 on page 533 illustrate the possible interfaces that can be created on the 4-port Channelized COC3/STM1 CE PIC:

Figure 53: 4-Port Channelized COC3/STM1 CE PIC Possible Interfaces (T1 Size)

g017388

Figure 54: 4-Port Channelized COC3/STM1 CE PIC Possible Interfaces (E1 Size)

g017389

Subrate T1 is not supported.

ATM NxDS0 grooming is not supported.

External and internal loopback of T1/E1 (on ct1/ce1 physical interfaces) can be configured using the `sonet-options` statement. By default, no loopback is configured.

Configuring a Port for SONET or SDH Mode on a 4-Port Channelized COC3/STM1 CE PIC

Each port of the 4-port Channelized COC3/STM1 CE PIC can be independently configured for either SONET or SDH mode. To configure a port for either SONET or SDH mode, enter the `[framing <sonet> | <sdh>]` statement at the `chassis fpc number pic number port number` hierarchy level.

The following example shows how to configure FPC 1, PIC 1, and port 0 for SONET mode and port 1 for SDH mode:

```

set chassis fpc 1 pic 1 port 0 framing sonet
set chassis fpc 1 pic 1 port 1 framing sdh
  
```

Or specify the following:

```

[edit]
fpc 1 {
  pic 1 {
    port 0 {
      framing sonet;
    }
    port 1 {
      framing sdh;
    }
  }
}
  
```

Configuring an ATM Interface on a COC1

To create an ATM interface on a COC1, enter the following command:

```
set interfaces coc1-x/y/z:1 partition 2 interface-type at
[edit]
fpc 1 {
  pic 1 {
    port 0 {
      framing sonet;
    }
    port 1 {
      framing sonet;
    }
  }
}
```

Or specify the following:

```
interfaces {
  coc1-x/y/z:1 {
    partition number {
      interface-type <at>;
    }
  }
}
```

The interfaces `at-x/y/z:1:2` are created.

To create an ATM interface on CAU4:

```
set interfaces cau4-x/y/z partition 2 interface-type at
```

Or specify the following:

```
interfaces {
  cau4-x/y/z {
    partition number {
      interface-type <at>;
    }
  }
}
```

You can use the `show chassis hardware` command to display a list of the installed PICs.

Configuring ATM Pseudowires

ATM pseudowires are described in RFC 4717. Pseudowire encapsulation is selected by configuring for a cell-relay pseudowire:

```
[edit interfaces at-fpc/pic/port:unit n]
encapsulation atm-ccc-cell-relay;
```

```
atm-l2circuit-mode cell;
```

Or for an AAL5 pseudowire:

```
encapsulation atm-ccc-vc-mux;
atm-l2circuit-mode aal5;
```



NOTE: encapsulation <atm-ccc-cell-relay> can be set at either the physical interface or logical interface level. atm-ccc-vc-mux can only be set at the logical interface level.

Cell Relay Mode (*atm-l2circuit-mode cell*)

In cell relay mode, one or more cells are bundled together to form a packet that is sent across the PSN tunnel. N-to-one mode is used to encapsulate cell bundles. In this mode, 52 bytes of each cell are transported across the PSN (the HEC field of the ATM header is omitted). The optional one-to-one mode is not supported.

By default, each ATM cell is encapsulated into a pseudowire packet (per RFC 4717) and sent over the pseudowire (*cell-bundle-size* = 1). The pseudowire may be configured to aggregate a user-configured number of cells into a packet to increase network utilization efficiency.

```
[edit interfaces at-fpc/pic/port]
atm-options {
    cell-bundle-size cells;
}
```

Where *cells* is the number of cells each pseudowire packet should contain.

Configuring VP or Port Promiscuous Mode

By default, all incoming cells are mapped from a single VC to an ATM pseudowire. For ATM physical interfaces configured with *atm-l2circuit-mode cell*, you can configure port or VP promiscuous mode.

In VP promiscuous mode, all cells with the same VPI are forwarded on a single pseudowire:

```
[edit interfaces at-fpc/pic/port]
atm-options {
    pic-type atm-ce;
    promiscuous-mode {
        vpi number;
    }
}
unit 0 {
    vpi 0;
}
```

In port promiscuous mode, all cells received on a T1 or E1 ATM port are forwarded across a single pseudowire:

```
[edit interfaces at-fpc/pic/port]
encapsulation <atm-ccc-cell-relay>;
atm-options {
  pic-type atm-ce;
  promiscuous-mode
}
unit 0 {
  allow-any-vci
}
```

Use the `show interface at-x/y/z:n` command to view cell relay statistics.

Configuring AAL5 SDU Mode (*atm-l2circuit-mode aal5*)

In AAL5 SDU mode, the ATM logical interface (VC) expects all data to be either AAL5 encapsulated packets or OAM cells. AAL5 packets are deencapsulated (AAL5 trailer is stripped off), prepended with an ATM pseudowire control word (RFC 4717) and forwarded on the pseudowire.

OAM cells that are received while an AAL5 packet is being reassembled are forwarded on the pseudowire immediately (they are reordered ahead of the packet being reassembled).

Use the `show interface at-x/y/z:n` command to view AAL5 statistics.

ATM OAM

CE PICs provide ATM support for the following OAM-FM cell types:

- F4 AIS (end-to-end)
- F4 RDI (end-to-end)
- F4 loopback (end-to-end)
- F5 loopback
- F5 AIS
- F5 RDI

For information on OAM configuration (for loopback OAM cells) and behavior (for AIS and RDI) see *Network Interfaces Configuration Guide*.

VP Pseudowires (CCC Encapsulation)

In the case of ATM VP pseudowires (all VCs in a VP are transported over a single *N-to-one* mode pseudowire), all F4 and F5 OAM cells are forwarded through the pseudowire.

Port Pseudowires (CCC Encapsulation)

Like VP pseudowires, port pseudowires all F4 and F5 OAM cells are forwarded through the pseudowire.

VC Pseudowires (CCC Encapsulation)

In the case of VC pseudowires, F5 OAM cells are forwarded through the pseudowire while F4 OAM cells are terminated at the RE.

Scaling

The 12-port Channelized T1/E1 CE PIC supports a maximum of 1000 VCs.

The 4-port Channelized COC3/STM1 CE PIC supports a maximum of 2000 VCs.

Congestion Control

ATM encapsulations provide congestion control via EPD thresholds on a per logical interface basis. For CE PICs, the EPD number specifies the number of packets (or frames, or cell bundles).

```
[edit interfaces at-fpc/pic/port unit unit-number]
epd-threshold <packets> plp1 <packets>;
```

QoS/Shaping

Unspecified bit rate (UBR) is supported.

Configuring the PIC Type

To configure Circuit Emulation PICs, you must specify the `atm-options` statement `pic-type` option as `atm-ce`, as follows:

```
[edit interfaces at-fpc/pic/port]
atm-options {
  pic-type <atm-ce>;
}
```

Configuring Layer 2 Circuit and Layer 2 VPN Pseudowires

ATM Layer 2 circuit and Layer 2 VPN pseudowires are configured using the same syntax described for ATM2 PICs:

```
protocols {
  # MPLS and routing config omitted for brevity.
  l2circuit {
    neighbor 10.255.245.1 { # loopback addr on remote router
```

```

        interface at-0/0/0.0 { # Pinot PIC ATM interface configured for CCC
            virtual-circuit-id 100;
        }
    }
}

```

Supported Interface Configurations

With 12-port Channelized T1/E1 CE PICs, ATM supports T1 and E1 interfaces:

```

[edit interfaces at-fpc/pic/port ]
t1-options
e1-options

```

A sample configuration follows:

```

at-0/2/1:3 {
    atm-options {
        pic-type atm-ce;
    }
    e1-options {
        framing g704;
    }
    t1-options {
        framing sf;
    }
}

```



NOTE: In the sample configuration above, both T1 and E1 framing are set. Depending on which CE PIC you are using (T1 or E1), only the appropriate options are functional.

The following CLI output showing the available T1 interface options:

```

[edit interfaces at-0/2/1:3]
user@host# set t1-options ?

```

```

Possible completions:
+ apply-groups          Groups from which to inherit configuration data
+ apply-groups-except   Don't inherit configuration data from these groups
bert-algorithm          Set BERT algorithm
byte-encoding           Byte encoding
crc-major-alarm-threshold  CRC Major alarm threshold value
crc-minor-alarm-threshold  CRC Minor alarm threshold value
framing                 Framing mode
invert-data             Invert data
line-encoding           Line encoding
loopback                Loopback mode

```

The following CLI output showing the available E1 interface options:

```

[edit interfaces at-0/2/1:3]
user@host# set e1-options ?

```

Possible completions:	
+ apply-groups	Groups from which to inherit configuration data
+ apply-groups-except	Don't inherit configuration data from these groups
bert-error-rate	Bit error rate (10^{-n} for $n > 0$, and zero for $n = 0$)
(0..7)	
bert-period	Length of BERT test (1..86400 seconds)
framing	Framing mode
loopback	Loopback mode

ATM Limitations

The following limitations apply to ATM support on Circuit Emulation PICs:

- Packet MTU—Packet MTU is limited to 2048 bytes.
- Trunk Mode ATM Pseudowires—CE PICs do not support trunk mode ATM pseudowires:
- OAM-FM Segment—Segment F4 flows are not supported. Only end-to-end F4 flows are supported.
- IP and Ethernet Encapsulations—IP and Ethernet encapsulations are not supported.
- F5 OAM—OAM termination is not supported.

Part 9

Configuring E1, E3, T1, and T3 Interfaces

- Configuring E1 Interfaces on page 543
- Configuring E3 Interfaces on page 551
- Configuring T1 Interfaces on page 559
- Configuring T3 Interfaces on page 569

Chapter 29

Configuring E1 Interfaces

This chapter contains the following sections:

- E1 Interfaces Overview on page 543
- Configuring E1 Physical Interface Properties on page 544
- Configuring E1 BERT Properties on page 544
- Configuring the E1 Frame Checksum on page 545
- Configuring E1 Framing on page 546
- Configuring the E1 Idle Cycle Flag on page 546
- Configuring E1 Data Inversion on page 546
- Configuring E1 Loopback Capability on page 547
- Configuring E1 Start and End Flags on page 548
- Configuring Fractional E1 Time Slots on page 548

E1 Interfaces Overview

E1 is a standard WAN digital communication format designed to operate over copper facilities at a rate of 2.048 Mbps. Widely used outside North America, it is a basic time-division multiplexing scheme used to carry digital circuits. The following standards apply to E1 interfaces:

- ITU-T Recommendation G.703, *Physical/electrical characteristics of hierarchical digital interfaces*, describes data rates and multiplexing schemes for the E series.
- ITU-T Recommendation G.751, *General Aspects of Digital Transmission Systems: Terminal Equipment*, describes framing methods.
- ITU-T Recommendation G.775, *Loss of Signal (LOS) and Alarm Indication Signal (AIS) Defect Detection and Clearance Criteria*, describes alarm reporting methods.



NOTE: The Juniper Networks E1 Physical Interface Card (PIC) does not support Channel Associated Signaling (CAS).

Configuring E1 Physical Interface Properties

To configure E1-specific physical interface properties, include the **e1-options** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
e1-options {
  bert-error-rate rate;
  bert-period seconds;
  fcs (16 | 32);
  framing (g704 | g704-no-crc4 | unframed);
  idle-cycle-flag (flags | ones);
  invert-data;
  loopback (local | remote);
  start-end-flag (filler | shared);
  timeslots time-slot-range;
}
```

Configuring E1 BERT Properties

This section discusses BERT properties for the E1 interface specifically. For general information about the JUNOS implementation of the BERT procedure, see “Interface Diagnostics” on page 134.

You can configure an E1 interface or a CE1 or E1 partition on a channelized PIC to execute a bit error rate test (BERT) when the interface receives a request to run this test. You specify the duration of the test and the error rate to include in the bit stream by including the **bert-period** and **bert-error-rate** statements at the [edit interfaces *interface-name* **e1-options**] hierarchy level:

```
[edit interfaces interface-name e1-options]
bert-error-rate rate;
bert-period seconds;
```

By default, the BERT period is 10 seconds. You can configure the BERT period to last from 1 through 239 seconds on some PICs and from 1 through 240 seconds on other PICs. Standard CE1, standard E1, E1 IQ, and E1 IQE interfaces, and PICs partitioned to CE1 and E1 channels, support an extended BERT period range, up to 86,400 seconds (24 hours), and have a default BERT period value of 240 seconds.



NOTE: When configuring E1 and CE1 interfaces on 10-port Channelized E1/T1 IQE PICs, the **bert-period** statement must be included at the [edit interfaces *ce1-fpc/pic/port*] hierarchy level.

rate is the bit error rate. This can be an integer from 0 through 7, which corresponds to a bit error rate from 10^{-0} (0, which corresponds to no errors) to 10^{-7} (1 error per 10 million bits). The default is 0.

Individual concatenated E1 interfaces do not support the **bert-algorithm** configuration statement. For individual concatenated E1 interfaces, the **bert-algorithm** statement at the [edit interfaces *interface-name* e1-options] hierarchy level is ignored. The algorithm for the E1 BERT procedure is **pseudo-2e15-o151** (pattern is $2^{15}-1$, as defined in the CCITT/ITU O.151 standard).

For channelized E1 intelligent queuing (IQ and IQE) interfaces, you can configure the BERT algorithm by including the **bert-algorithm** statement at the [edit interfaces *ce1-fpc/pic/port* e1-options] or [edit interfaces *e1-fpc/pic/port* e1-options] hierarchy level:

```
[edit interfaces ce1-fpc/pic/port e1-options]
bert-algorithm algorithm;
[edit interfaces e1-fpc/pic/port e1-options]
bert-algorithm algorithm;
```

For a list of supported algorithms, enter a ? after the **bert-algorithm** statement; for example:

```
[edit interfaces ce1-0/0/0 e1-options]
user@host# set bert-algorithm ?
Possible completions:
pseudo-2e11-o152 Pattern is 2^11 -1 (per O.152 standard)
pseudo-2e15-o151 Pattern is 2^15 - 1 (per O.152 standard)
pseudo-2e20-o151 Pattern is 2^20 - 1 (per O.151 standard)
pseudo-2e20-o153 Pattern is 2^20 - 1 (per O.153 standard)
```

Configuring the E1 Frame Checksum

By default, the E1 interface supports a 16-bit checksum. You can configure a 32-bit checksum, which provides more reliable packet verification. However, some older equipment might not support 32-bit checksums.

To configure a 32-bit checksum, include the **fcs 32** statement at the [edit interfaces *interface-name* e1-options] hierarchy level:

```
[edit interfaces interface-name e1-options]
fcs 32;
```

To return to the default 16-bit frame checksum, delete the **fcs 32** statement from the configuration:

```
[edit]
user@host# delete interfaces e1-fpc/pic/port e1-options fcs 32
```

To explicitly configure a 16-bit checksum, include the **fcs 16** statement at the [edit interfaces *interface-name* e1-options] hierarchy level:

```
[edit interfaces interface-name e1-options]
fcs 16;
```

Configuring E1 Framing

By default, E1 interfaces use the G704 framing mode. You can configure the alternative unframed mode if needed.

To have the interface use the unframed mode, include the **framing** statement at the [edit interfaces *interface-name* e1-options] hierarchy level, specifying the **unframed** option:

```
[edit interfaces interface-name e1-options]  
framing unframed;
```

To explicitly configure G704 framing, include the **framing** statement at the [edit interfaces *interface-name* e1-options] hierarchy level, specifying the **g704** option:

```
[edit interfaces interface-name e1-options]  
framing g704;
```

By default, G704 framing uses CRC4. To explicitly configure an interface's G704 framing to not use CRC4, include the **framing** statement at the [edit interfaces *interface-name* e1-options] hierarchy level, specifying the **g704-no-crc4** option:

```
[edit interfaces interface-name e1-options]  
framing g704-no-crc4;
```

Configuring the E1 Idle Cycle Flag

By default, an E1 interface transmits the value 0x7E in the idle cycles. To have the interface transmit the value 0xFF (all ones) instead, include the **idle-cycle-flag** statement at the [edit interfaces *interface-name* e1-options] hierarchy level, specifying the **ones** option:

```
[edit interfaces interface-name e1-options]  
idle-cycle-flag ones;
```

To explicitly configure the default value of 0x7E, include the **idle-cycle-flag** statement with the **flags** option:

```
[edit interfaces interface-name e1-options]  
idle-cycle-flag flags;
```

Configuring E1 Data Inversion

By default, data inversion is disabled. To enable data inversion at the HDLC level, include the **invert-data** statement at the [edit interfaces *interface-name* e1-options] hierarchy level:

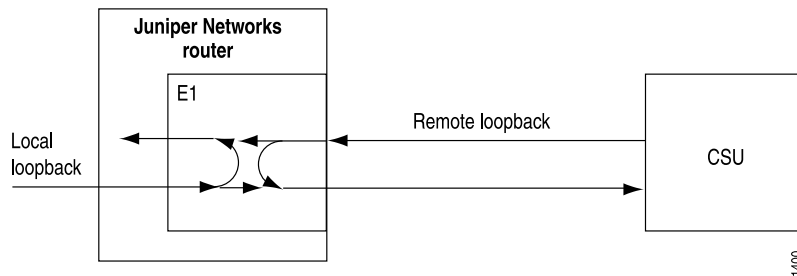
```
[edit interfaces interface-name e1-options]  
invert-data;
```

When you enable data inversion, all data bits in the data stream are transmitted inverted; that is, zeroes are transmitted as ones and ones as zeroes. Data inversion is normally used only in AMI mode to guarantee ones density in the transmitted stream.

Configuring E1 Loopback Capability

You can configure loopback capability between the local E1 interface and the remote channel service unit (CSU), as shown in Figure 55 on page 547. You can configure the loopback to be local or remote. With local loopback, the E1 interface can transmit packets to the CSU, but receives its own transmission back again and ignores data from the CSU. With remote loopback, packets sent from the CSU are received by the E1 interface, forwarded if there is a valid route, and immediately retransmitted to the CSU.

Figure 55: Remote and Local E1 Loopback



To configure loopback capability on an E1 interface, include the `loopback` statement at the [edit interfaces *interface-name* e1-options] hierarchy level:

```
[edit interfaces interface-name e1-options]
loopback (local | remote);
```

Packets can be looped on either the local router or the remote CSU.

To exchange BERT patterns between a local router and a remote router, include the `loopback remote` statement in the interface configuration at the remote end of the link. From the local router, you issue the `test interface` command.

For more information about configuring BERT, see “Interface Diagnostics” on page 134. For more information about using operational mode commands to test interfaces, see the *JUNOS System Basics and Services Command Reference*.

To turn off the loopback capability, remove the `loopback` statement from the configuration:

```
[edit]
user@host# delete interfaces e1-fpc/pic/port e1-options loopback
```

You can determine whether there is an internal problem or an external problem by checking the error counters in the output of the `show interface interface-name extensive` command:

```
user@host> show interfaces interface-name extensive
```

Example: Configuring E1 Loopback Capability

To determine whether a problem is internal or external, loop packets on both the local and the remote router. To do this, include the `no-keepalives` and `encapsulation cisco-hdlc` statements at the `[edit interfaces interface-name]` hierarchy level and the `loopback local` statement at the `[edit interfaces interface-name e1-options]` hierarchy level.

With this configuration, the link stays up, so you can loop ping packets to a remote router. The `loopback local` statement causes the interface to loop within the PIC just before the data reaches the transceiver.

```
[edit interfaces]
e1-1/0/0 {
  no-keepalives;
  encapsulation cisco-hdlc;
  e1-options {
    loopback local;
  }
  unit 0 {
    family inet {
      address 10.100.100.1/24;
    }
  }
}
```

Configuring E1 Start and End Flags

By default, start and end flags are shared.

To configure an E1 interface to wait two idle cycles between the start and end flags, include the `start-end-flag` statement with the `filler` option at the `[edit interfaces interface-name e1-options]` hierarchy level:

```
[edit interfaces interface-name e1-options]
start-end-flag filler;
```

To revert to the default behavior, sharing the transmission of start and end flags, include the `start-end-flag` statement with the `shared` option at the `[edit interfaces interface-name e1-options]` hierarchy level:

```
[edit interfaces interface-name e1-options]
start-end-flag shared;
```

Configuring Fractional E1 Time Slots

By default, all the time slots on an E1 interface are used. To configure the number of time slots allocated to a fractional E1 interface, include the `timeslots` statement at the `[edit interfaces interface-name e1-options]` hierarchy level:

```
[edit interfaces interface-name e1-options]
timeslots time-slot-range;
```

There are 32 time slots on an E1 interface. Time slot 0 is always reserved for framing and cannot be used to configure a fractional E1 interface.

Time slot numbering constraints vary for different E1 PICs, as follows:

- For 4-port E1 PICs, the configurable time slot range is 1 through 31 (time slot 0 is reserved for framing).
- For 10-port Channelized E1 and 10-port Channelized E1 Intelligent Queuing (IQ) PICs, the configurable time slot range is 2 through 32 (time slots 0 and 1 are reserved for framing).
- For Enhanced Intelligent Queuing (IQE) PICs, the configurable time slot range is 2 through 32.
- NxDS0 time slots configured on either a channelized STM1 IQ interface or a channelized E1IQ interface are numbered from 1 to 31 (0 is reserved), while fractional E1 time slots are numbered from 2 to 32 (0 and 1 are reserved).
- For fractional E1 interfaces only, if you connect a 4-port E1 PIC to a device that uses time slot numbering from 2 through 32, you must subtract 1 from the configured number of time slots. To do this, include the **timeslots** statement at the [edit interfaces *interface-name* e1-options] hierarchy level, and offset 1 from the specified slot number.



NOTE: When configuring fractional E1 time slots, you also must include the **framing g704** statement at the [edit interfaces *e1-fpc/pic/port* e1-options] hierarchy level.

To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces.

Example: Configuring Fractional E1 Time Slots

In this example, time slots are offset by 1 to compensate for the fractional E1 interface being connected to a device that uses time slot numbering from 0 through 31.

Use Time Slots 3 Through 5, 10, and 24	[edit interfaces <i>interface-name</i> e1-options] # Fractional E1 interface timeslots 4-6,11,25;
Use Time Slots 1 Through 10	[edit interfaces <i>interface-name</i> e1-options] timeslots 1-10;
Use Time Slots 1 Through 5, 10, and 24	[edit interfaces <i>interface-name</i> e1-options] timeslots 1-5,10,24;

Chapter 30

Configuring E3 Interfaces

This chapter contains the following sections:

- E3 Interfaces Overview on page 551
- Configuring E3 Physical Interface Properties on page 552
- Configuring E3 BERT Properties on page 552
- Configuring the E3 CSU Compatibility Mode on page 553
- Configuring the E3 Frame Checksum on page 554
- Configuring the E3 Idle Cycle Flag on page 555
- Configuring E3 Data Inversion on page 555
- Configuring E3 Loopback Capability on page 555
- Configuring E3 HDLC Payload Scrambling on page 557
- Configuring the E3 Start and End Flags on page 557
- Configuring E3 IQ and IQE Unframed Mode on page 558

E3 Interfaces Overview

E3 is a high-speed WAN digital communication technique designed to operate over copper facilities at a rate of 34.368 Mbps. Widely used outside North America, it is the time-division multiplexing scheme used to carry 16 E1 circuits. The following standards apply to E3 interfaces:

- ITU-T Recommendation G.703, *Physical/electrical characteristics of hierarchical digital interfaces*, describes data rates and multiplexing schemes for the E series.
- ITU-T Recommendation G.751, *General Aspects of Digital Transmission Systems: Terminal Equipment*, describes framing methods.
- ITU-T Recommendation G.775, *Loss of Signal (LOS) and Alarm Indication Signal (AIS) Defect Detection and Clearance Criteria*, describes alarm reporting methods.

The JUNOS Software supports the E3 Physical Interface Card (PIC) and the E3 Intelligent Queuing (IQ and IQE) PICs. The E3 IQ and E3 IQE PICs supports transmission scheduling on logical interfaces. For more information, see the *JUNOS Class of Service Configuration Guide*.



NOTE: In unframed mode, the E3 IQ and E3 IQE PICs do not detect yellow or loss-of-frame alarms.

Configuring E3 Physical Interface Properties

To configure E3-specific physical interface properties, include the **e3-options** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
e3-options {
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  compatibility-mode (digital-link | kentrox | larscom) <subrate value>;
  fcs (16 | 32);
  idle-cycle-flag value;
  invert-data;
  loopback (local | remote);
  (payload-scrambler | no-payload-scrambler);
  start-end-flag value;
  (unframed | no-unframed);
}
```

Configuring E3 BERT Properties

This section discusses BERT properties for the E3 interface specifically. For general information about the JUNOS implementation of the BERT procedure, see “Interface Diagnostics” on page 134.

You can configure an E3 interface to execute a bit error rate test (BERT) when the interface receives a request to run this test. You specify the duration of the test, the pattern to send in the bit stream, and the error rate to include in the bit stream by including the **bert-period**, **bert-algorithm**, and **bert-error-rate** statements at the [edit interfaces *interface-name* e3-options] hierarchy level:

```
[edit interfaces interface-name e3-options]
bert-algorithm algorithm;
bert-error-rate rate;
bert-period seconds;
```

By default, the BERT period is 10 seconds. You can configure the BERT period to last from 1 through 239 seconds on some PICs and from 1 through 240 seconds on other PICs.

rate is the bit error rate. This can be an integer from 0 through 7, which corresponds to a bit error rate from 10^{-0} (0, which corresponds to no errors) to 10^{-7} (1 error per 10 million bits).

algorithm is the pattern to send in the bit stream. On E3 interfaces, you can also select the pattern to send in the bit stream by including the **bert-algorithm** statement at the [edit interfaces *interface-name* *interface-options*] hierarchy level:


```
[edit interfaces interface-name interface-options]
bert-algorithm algorithm;
```

For a list of supported algorithms, enter a ? after the **bert-algorithm** statement; for example:

```
[edit interfaces e3-0/0/0 e3-options]
user@host# set bert-algorithm ?
Possible completions:
pseudo-2e11-o152 Pattern is 2^11 - 1 (per O.152 standard)
pseudo-2e15-o151 Pattern is 2^15 - 1 (per O.152 standard)
pseudo-2e20-o151 Pattern is 2^20 - 1 (per O.151 standard)
pseudo-2e20-o153 Pattern is 2^20 - 1 (per O.153 standard)
```

For specific hierarchy information, see individual interface types. For information about running the BERT procedure, see the *JUNOS System Basics and Services Command Reference*.

Configuring the E3 CSU Compatibility Mode

Subrating an E3 interface reduces the maximum allowable peak rate by limiting the High-level Data Link Control (HDLC)-encapsulated payload. Subrate modes configure the PIC to connect with channel service units (CSUs) that use proprietary methods of multiplexing.

On M Series and T Series routers, you can configure E3 interfaces to be compatible with a Digital Link, Kentrox, or Larscom CSU. On J Series Services Routers, you can configure E3 interfaces to be compatible with a Digital Link or Kentrox CSU.



NOTE: To subrate an E3 interface to be compatible with a Kentrox CSU, you must have an IQ or IQE-based PIC. Non-IQ or IQE PICs allow a commit of the configuration, but the interfaces remain at the full E3 rate for the Kentrox compatibility mode.

To configure an E3 interface so that it is compatible with the CSU at the remote end of the line, include the **compatibility-mode** statement at the `[edit interfaces interface-name e3-options]` hierarchy level:

```
[edit interfaces interface-name e3-options]
compatibility-mode (digital-link | kentrox | larscom) <subrate value>;
```

The subrate of an E3 interface must exactly match that of the remote CSU. To specify the subrate, include the **subrate** statement in the configuration:

- For Kentrox CSUs, specify the subrate as a number from 1 through 48 that exactly matches the value configured on the CSU. Each increment of the subrate value corresponds to a rate increment of about 0.5 Mbps.
- For Digital Link CSUs, you can specify the subrate value to match the data rate configured on the CSU in the format *xkb* or *x.xMb*. You can configure the subrate values shown in Table 46 on page 554.
- Larscom CSUs do not support the E3 subrate.

Table 46: Subrate Values for E3 Digital Link Compatibility Mode

358 Kbps	7.2 Mbps	14.0 Mbps	20.8 Mbps	27.6 Mbps
716 Kbps	7.5 Mbps	14.3 Mbps	21.1 Mbps	27.9 Mbps
1.1 Mbps	7.9 Mbps	14.7 Mbps	21.5 Mbps	28.3 Mbps
1.4 Mbps	8.2 Mbps	15.0 Mbps	21.8 Mbps	28.6 Mbps
1.8 Mbps	8.6 Mbps	15.4 Mbps	22.2 Mbps	29.0 Mbps
2.1 Mbps	9.0 Mbps	15.8 Mbps	22.6 Mbps	29.4 Mbps
2.5 Mbps	9.3 Mbps	16.1 Mbps	22.9 Mbps	29.7 Mbps
2.9 Mbps	9.7 Mbps	16.5 Mbps	23.3 Mbps	30.1 Mbps
3.2 Mbps	10.0 Mbps	16.8 Mbps	23.6 Mbps	30.4 Mbps
3.6 Mbps	10.4 Mbps	17.2 Mbps	24.0 Mbps	30.8 Mbps
3.9 Mbps	10.7 Mbps	17.5 Mbps	24.3 Mbps	31.1 Mbps
4.3 Mbps	11.1 Mbps	17.9 Mbps	24.7 Mbps	31.5 Mbps
4.7 Mbps	11.5 Mbps	18.3 Mbps	25.1 Mbps	31.9 Mbps
5.0 Mbps	11.8 Mbps	18.6 Mbps	25.4 Mbps	32.2 Mbps
5.4 Mbps	12.2 Mbps	19.0 Mbps	25.8 Mbps	32.6 Mbps
5.7 Mbps	12.5 Mbps	19.3 Mbps	26.1 Mbps	32.9 Mbps
6.1 Mbps	12.9 Mbps	19.7 Mbps	26.5 Mbps	33.3 Mbps
6.4 Mbps	13.2 Mbps	20.0 Mbps	26.9 Mbps	33.7 Mbps
6.8 Mbps	13.6 Mbps	20.4 Mbps	27.2 Mbps	

For information about subrating a T3 interface, see “Configuring the T3 CSU Compatibility Mode” on page 572.

Configuring the E3 Frame Checksum

You can configure a 32-bit checksum, which provides more reliable packet verification. However, some older equipment might not support 32-bit checksums.

On a channelized OC12 interface, the **fcs** statement is not supported. To configure FCS on each E3 channel, you must include the **e3-options fcs** statement in the configuration for each channel.

To configure a 32-bit checksum, include the **fcs** statement at the **[edit interfaces interface-name e3-options]** hierarchy level:

```
[edit interfaces interface-name e3-options]
fcs 32;
```

To return to the default 16-bit frame checksum, delete the `fcs 32` statement from the configuration:

```
[edit]
user@host# delete interfaces e3-fpc/pic/port e3-options fcs 32
```

To explicitly configure a 16-bit checksum, include the `fcs` statement at the `[edit interfaces interface-name e3-options]` hierarchy level:

```
[edit interfaces interface-name e3-options]
fcs 16;
```

Configuring the E3 Idle Cycle Flag

By default, an E3 interface transmits the value 0x7E in the idle cycles. To have the interface transmit the value 0xFF (all ones) instead, include the `idle-cycle-flag` statement at the `[edit interfaces interface-name e3-options]` hierarchy level, specifying the `ones` option:

```
[edit interfaces interface-name e3-options]
idle-cycle-flag ones;
```

To explicitly configure the default value of 0x7E, include the `idle-cycle-flag` statement with the `flags` option:

```
[edit interfaces interface-name e3-options]
idle-cycle-flag flags;
```

Configuring E3 Data Inversion

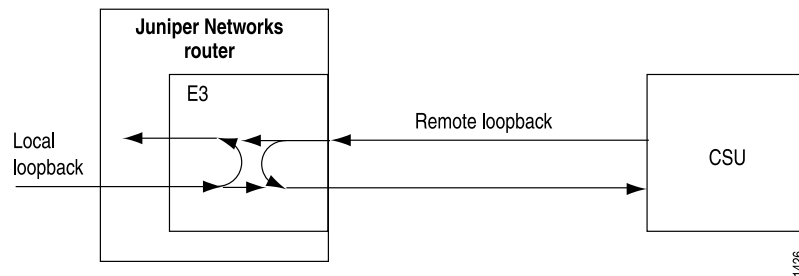
By default, data inversion is disabled. To enable data inversion at the HDLC level, include the `invert-data` statement at the `[edit interfaces interface-name e3-options]` hierarchy level:

```
[edit interfaces interface-name e3-options]
invert-data;
```

When you enable data inversion, unused data bits in the data stream are transmitted inverted; that is, zeroes are transmitted as ones and ones as zeroes. Enable inversion to be compatible with another vendor's E3 interface.

Configuring E3 Loopback Capability

You can configure loopback capability between the local E3 interface and the remote CSU. You can configure the loopback to be local or remote. With local loopback, the E3 interface can transmit packets to the CSU, but receives its own transmission back again and ignores data from the CSU. With remote loopback, packets sent from the CSU are received by the E3 interface, forwarded if there is a valid route, and immediately retransmitted to the CSU (see Figure 56 on page 556).

Figure 56: Remote and Local E3 Loopback

To configure loopback capability on an E3 interface, include the **loopback** statement at the [edit interfaces *interface-name* e3-options] hierarchy level:

```
[edit interfaces interface-name e3-options]
  loopback (local | remote);
```

Packets can be looped on either the local router or the remote CSU.

To exchange BERT patterns between a local router and a remote router, include the **loopback remote** statement in the interface configuration at the remote end of the link. From the local router, you issue the **test interface** command.

For more information about configuring BERT, see “Interface Diagnostics” on page 134. For more information about using operational mode commands to test interfaces, see the *JUNOS System Basics and Services Command Reference*.

To turn off the loopback capability, remove the **loopback** statement from the configuration:

```
[edit]
user@host# delete interfaces e3-fpc/pic/port e3-options loopback
```

You can determine whether there is an internal problem or an external problem by checking the error counters in the output of the **show interface *interface-name* extensive** command:

```
user@host> show interfaces interface-name extensive
```

Example: Configuring E3 Loopback Capability

To determine whether a problem is internal or external, loop packets on both the local and the remote router. To do this, include the **no-keepalives** and **encapsulation cisco-hdlc** statements at the [edit interfaces *interface-name*] hierarchy level and the **loopback local** statement at the [edit interfaces *interface-name* e3-options] hierarchy level. With this configuration, the link stays up, so you can loop ping packets to a remote router. The **loopback local** statement causes the interface to loop within the PIC just before the data reaches the transceiver.

```
[edit interfaces]
e3-1/0/0 {
  no-keepalives;
  encapsulation cisco-hdlc;
```

```

    e3-options {
        loopback local;
    }
    unit 0 {
        family inet {
            address 10.100.100.1/24;
        }
    }
}

```

Configuring E3 HDLC Payload Scrambling

E3 HDLC payload scrambling, which is disabled by default, provides better link stability. Both sides of a connection must either use or not use scrambling.

To configure scrambling on the interface, you can include the `payload-scrambler` statement at the `[edit interfaces interface-name e3-options]` hierarchy level:

```

[edit interfaces interface-name e3-options]
payload-scrambler;

```

To explicitly disable HDLC payload scrambling, include the `no-payload-scrambler` statement at the `[edit interfaces interface-name e3-options]` hierarchy level:

```

[edit interfaces interface-name e3-options]
no-payload-scrambler;

```

To disable payload scrambling again (return to the default), delete the `payload-scrambler` statement from the configuration:

```

[edit]
user@host# delete interfaces e3-fpc/pic/port e3-options payload-scrambler

```

Configuring the E3 Start and End Flags

By default, an E3 interface shares the transmission of the start and end flags

To configure an E3 interface to wait two idle cycles between the start and end flags, include the `start-end-flag` statement with the `filler` option at the `[edit interfaces interface-name e3-options]` hierarchy level:

```

[edit interfaces interface-name e3-options]
start-end-flag filler;

```

To revert to the default behavior, sharing the transmission of start and end flags, include the `start-end-flag` statement with the `shared` option at the `[edit interfaces interface-name e3-options]` hierarchy level:

```

[edit interfaces interface-name e3-options]
start-end-flag shared;

```

Configuring E3 IQ and IQE Unframed Mode

For E3 IQ and IQE interfaces only, you can enable or disable unframed mode. In unframed mode, the E3 IQ and IQE interfaces do not detect yellow (ylw) or loss-of-frame (lof) alarms.

By default, unframed mode is disabled. To enable unframed mode, include the **unframed** statement at the [edit interfaces *interface-name* e3-options] hierarchy level:

```
[edit interfaces interface-name e3-options]  
unframed;
```

To explicitly configure the default of framed mode, include the **no-unframed** statement:

```
[edit interfaces interface-name e3-options]  
no-unframed;
```

Chapter 31

Configuring T1 Interfaces

This chapter includes an overview of T1 interfaces and configuration information as follows:

- T1 Interfaces Overview on page 559
- Configuring T1 Physical Interface Properties on page 560
- Configuring T1 BERT Properties on page 560
- Configuring the T1 Buildout on page 561
- Configuring T1 Byte Encoding on page 561
- Configuring T1 CRC Error Major Alarm Thresholds on page 562
- Configuring T1 CRC Error Minor Alarm Thresholds on page 562
- Configuring T1 Data Inversion on page 563
- Configuring the T1 Frame Checksum on page 563
- Configuring the T1 Remote Loopback Response on page 564
- Configuring T1 Framing on page 564
- Configuring T1 Line Encoding on page 564
- Configuring T1 Loopback Capability on page 565
- Configuring the T1 Idle Cycle Flag on page 566
- Configuring T1 Start and End Flags on page 567
- Configuring Fractional T1 Time Slots on page 567

T1 Interfaces Overview

T1 is the basic physical layer protocol used by the Digital Signal level 1 (DS1) multiplexing method in North America. A T1 interface operates at a bit rate of 1.544 Mbps and can support 24 DS0 channels. Supported DS1 standards include:

- ANSI T1.107, T1.102
- GR 499-core, GR 253-core
- AT&T Pub 54014
- ITU G.751, G.703

Configuring T1 Physical Interface Properties

To configure T1-specific physical interface properties, include the **t1-options** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
t1-options {
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  buildout value;
  byte-encoding (nx56 | nx64);
  crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);
  crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);
  fcs (16 | 32);
  framing (esf | sf);
  idle-cycle-flag (flags | ones);
  invert-data;
  line-encoding (ami | b8zs);
  loopback (local | payload | remote);
  remote-loopback-respond;
  start-end-flag (filler | shared);
  timeslots time-slot-range;
}
```

Configuring T1 BERT Properties

This section discusses BERT properties for the T1 interface specifically. For general information about the JUNOS implementation of the BERT procedure, see “Interface Diagnostics” on page 134.

You can configure a T1 interface or partitioned CT1 or T1 channel to execute a bit error rate test (BERT) when the interface receives a request to run this test. You specify the duration of the test and the error rate to include in the bit stream by including the **bert-period** and **bert-error-rate** statements at the [edit interfaces *interface-name* t1-options] hierarchy level:

```
[edit interfaces interface-name t1-options]
bert-algorithm algorithm;
bert-error-rate rate;
bert-period seconds;
```

seconds is the duration of the BERT procedure. The test can last from 1 through 239 seconds; the default is 10 seconds. Standard CT1, standard T1, T1 IQ, and T1 IQE interfaces, and PICs partitioned to CT1 and T1 channels, support an extended BERT period range, up to 86,400 seconds (24 hours), and have a default BERT period value of 240 seconds.



NOTE: When configuring T1 and CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the **bert-period** statement must be included at the [edit interfaces *ct1-fpc/pic/port*] hierarchy level.

rate is the bit error rate. This can be an integer from 0 through 7, which corresponds to a bit error rate from 10^{-0} (1 error per bit) to 10^{-7} (1 error per 10 million bits).

algorithm is the pattern to send in the bit stream. On T1 interfaces, you can also select the pattern to send in the bit stream by including the **bert-algorithm** statement at the [edit interfaces *interface-name interface-options*] hierarchy level:

```
[edit interfaces interface-name interface-options]
bert-algorithm algorithm;
```

For a list of supported algorithms, enter a ? after the **bert-algorithm** statement; for example:

```
[edit interfaces t1-0/0/0 t1-options]
user@host# set bert-algorithm ?
Possible completions:
pseudo-2e11-o152 Pattern is 2^11 - 1 (per 0.152 standard)
pseudo-2e15-o151 Pattern is 2^15 - 1 (per 0.152 standard)
pseudo-2e20-o151 Pattern is 2^20 - 1 (per 0.151 standard)
pseudo-2e20-o153 Pattern is 2^20 - 1 (per 0.153 standard)
```

For specific hierarchy information, see individual interface types. For information about running the BERT procedure, see the *JUNOS System Basics and Services Command Reference*.

Configuring the T1 Buildout

A T1 interface has five possible setting ranges for the T1 line buildout: 0-132, 133-265, 266-398, 399-531, or 532-655 feet. By default, the T1 interface uses the shortest setting (0-132).

To have the interface drive a line at one of the longer distance ranges, include the **buildout** statement with the appropriate value at the [edit interfaces *interface-name t1-options*] hierarchy level:

```
[edit interfaces interface-name t1-options]
buildout value;
```

Configuring T1 Byte Encoding

By default, T1 interfaces use a byte encoding of 8 bits per byte (nx64). You can configure an alternative byte encoding of 7 bits per byte (nx56).

To have the interface use 7 bits per byte encoding, include the **byte-encoding** statement at the [edit interfaces *interface-name t1-options*] hierarchy level, specifying the nx56 option:

```
[edit interfaces interface-name t1-options]
byte-encoding nx56;
```

To explicitly configure nx64 byte encoding, include the **byte-encoding** statement at the [edit interfaces *interface-name* t1-options] hierarchy level, specifying the nx64 option:

```
[edit interfaces interface-name t1-options]
byte-encoding nx64;
```

Configuring T1 CRC Error Major Alarm Thresholds

JUNOS Software collects CRC errors from PICs every second. On Channelized OC3 IQ and IQE PICs, Channelized OC12 IQ and IQE PICs, and Channelized T3 IQ PICs, you can configure major error thresholds for T1 CRC errors.

When the threshold is exceeded for 1 second, a defect condition is declared. If the defect condition continues for the monitoring period, an alarm condition is declared. You can display the CRC error threshold configuration, CRC errors count, and the alarm condition using the **show interfaces extensive** command.

To configure a CRC major error threshold, include the **crc-major-alarm-threshold** statement at the [edit interfaces *interface-name* t1-options] hierarchy level, specifying the errors per bits as 1e-3, 5e-4, 1e-4, 5e-5 or 1e-5:

```
[edit interfaces interface-name t1-options]
crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);
```

To configure a T1 CRC error major alarm for five errors in 10^{-4} bits, include the **crc-major-alarm-threshold** statement at the [edit interfaces *interface-name* t1-options] hierarchy level, specifying the 5e-4 option:

```
[edit interfaces interface-name t1-options]
crc-major-alarm-threshold 5e-4;
```

All settings except 1e-5 use a 10-second monitoring period. The 1e-5 value uses a 50-second monitoring period.

Configuring T1 CRC Error Minor Alarm Thresholds

JUNOS Software collects CRC errors from PICs every second. On Channelized OC3 IQ and IQE PICs, Channelized OC12 IQ and IQE PICs, and Channelized T3 IQ PICs, you can configure minor error thresholds for T1 CRC errors.

When the threshold is exceeded for 1 second, a defect condition is declared. If the defect condition continues for the monitoring period, an alarm condition is declared. You can display the CRC error threshold configuration, CRC errors count, and the alarm condition using the **show interfaces extensive** command.

To configure a CRC minor error threshold, include the **crc-minor-alarm-threshold** statement at the [edit interfaces *interface-name* t1-options] hierarchy level, specifying the errors per bits as 1e-3, 5e-4, 1e-4, 5e-5, 1e-5, 5e-6, or 1e-6:

```
[edit interfaces interface-name t1-options]
crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);
```

To configure a T1 CRC error minor alarm for five errors in 10^{-4} bits, include the `crc-minor-alarm-threshold` statement at the `[edit interfaces interface-name t1-options]` hierarchy level, specifying the `5e-4` option:

```
[edit interfaces interface-name t1-options]
crc-minor-alarm-threshold 5e-4;
```

The 10-second monitoring period is used for values `1e-3`, `5e-4`, `1e-4`, and `5e-5`. The `1e-5` value uses a 50-second monitoring period. The `5e-6` value uses a 100-second monitoring period. The `1e-6` value uses a 500-second monitoring period.

Configuring T1 Data Inversion

By default, data inversion is disabled. To enable data inversion at the HDLC level, include the `invert-data` statement at the `[edit interfaces interface-name t1-options]` hierarchy level:

```
[edit interfaces interface-name t1-options]
invert-data;
```

When you enable data inversion, all data bits in the data stream are transmitted inverted; that is, zeroes are transmitted as ones and ones as zeroes. Data inversion is normally used only in AMI mode to guarantee ones density in the transmitted stream.

Configuring the T1 Frame Checksum

By default, T1 interfaces use a 16-bit frame checksum. You can configure a 32-bit checksum, which provides more reliable packet verification. However, some older equipment might not support 32-bit checksums.

To configure a 32-bit checksum, include the `fcs 32` statement at the `[edit interfaces interface-name t1-options]` hierarchy level:

```
[edit interfaces interface-name t1-options]
fcs 32;
```

To return to the default 16-bit frame checksum, delete the `fcs 32` statement from the configuration:

```
[edit]
user@host# delete interfaces t1-fpc/pic/port t1-options fcs 32
```

To explicitly configure a 16-bit checksum, include the `fcs 16` statement at the `[edit interfaces interface-name t1-options]` hierarchy level:

```
[edit interfaces interface-name t1-options]
fcs 16;
```

Configuring the T1 Remote Loopback Response

The T1 facilities data-link loop request signal is used to communicate various network information in the form of in-service monitoring and diagnostics. Extended superframe, through the facilities data link (FDL), supports nonintrusive signaling and control, thereby offering clear-channel communication. Remote loopback requests can be over the FDL or inband. To configure the router to respond to remote loopback requests, include the `remote-loopback-respond` statement at the `[edit interfaces interface-name t1-options]` hierarchy level:

```
[edit interfaces interface-name t1-options]
remote-loopback-respond;
```

By default, the router does not respond to remote loopback requests.

Configuring T1 Framing

By default, T1 interfaces use extended superframe framing format. You can configure SF (superframe) as an alternative.

To have the interface use the SF framing format, include the `framing` statement at the `[edit interfaces interface-name t1-options]` hierarchy level, specifying the `sf` option:

```
[edit interfaces interface-name t1-options]
framing sf;
```

To explicitly configure ESF framing, include the `framing` statement at the `[edit interfaces interface-name t1-options]` hierarchy level, specifying the `esf` option:

```
[edit interfaces interface-name t1-options]
framing esf;
```

Configuring T1 Line Encoding

By default, T1 interfaces use B8ZS line encoding. You can configure AMI line encoding if necessary.

To have the interface use AMI line encoding, include the `line-encoding` statement at the `[edit interfaces interface-name t1-options]` hierarchy level, specifying the `ami` option:

```
[edit interfaces interface-name t1-options]
line-encoding ami;
```

To explicitly configure B8ZS line encoding, include the `line-encoding` statement at the `[edit interfaces interface-name t1-options]` hierarchy level, specifying the `b8zs` option:

```
[edit interfaces interface-name t1-options]
line-encoding b8zs;
```


For M Series and T Series routers, you must set the line encoding parameter for paired ports to the same value. Ports 0 and 1 must share the same value, and likewise

ports 2 and 3 must share the same value, but ports 0 and 1 can have a different value from that of ports 2 and 3.

Configuring T1 Loopback Capability

You can configure loopback capability between the local T1 interface and the remote channel service unit (CSU), as shown in Figure 57 on page 565. You can configure the loopback to be local or remote. With local loopback, the T1 interface can transmit packets to the CSU, but receives its own transmission back again and ignores data from the CSU. With remote loopback, packets sent from the CSU are received by the T1 interface, forwarded if there is a valid route, and immediately retransmitted to the CSU.

Figure 57: Remote and Local T1 Loopback



To configure loopback capability on a T1 interface, include the **loopback** statement at the [edit interfaces *interface-name* t1-options] hierarchy level:

```
[edit interfaces interface-name t1-options]
loopback (local | payload | remote);
```

Packets can be looped on either the local router or the remote CSU. Local and remote loopback loop back both data and clocking information.

To exchange BERT patterns between a local router and a remote router, include the **loopback remote** statement in the interface configuration at the remote end of the link. From the local router, issue the **test interface** command.

For more information about configuring BERT, see “Interface Diagnostics” on page 134. For more information about using operational mode commands to test interfaces, see the *JUNOS System Basics and Services Command Reference*.

For channelized T3, T1, and NxDS0 intelligent queuing (IQ) interfaces only, you can include the **loopback payload** statement in the configuration to loop back data only (without clocking information) on the remote router’s PIC. In payload loopback, overhead is recalculated. For T3 IQ interfaces, you can include the **loopback payload** statement at the [edit interfaces ct3-fpc/pic/port] and [edit interfaces t3-fpc/pic/port:channel] hierarchy levels. For T1 interfaces, you can include the **loopback payload** statement in the configuration at the [edit interfaces t1-fpc/pic/port:channel] hierarchy level; it is ignored if included at the [edit interfaces ct1-fpc/pic/port] hierarchy level. For NxDS0 interfaces, payload and remote loopback

are the same. If you configure one, the other is ignored. NxDS0 IQ interfaces do not support local loopback.

To determine whether a problem is internal or external, you can loop packets on both the local and the remote router. To do this, include the `no-keepalives` and `encapsulation cisco-hdlc` statements at the `[edit interfaces interface-name]` hierarchy level and the `loopback local` statement at the `[edit interfaces interface-name t1-options]` hierarchy level, as shown in the following example:

```
[edit interfaces]
t1-1/0/0 {
  no-keepalives;
  encapsulation cisco-hdlc;
  t1-options {
    loopback local;
  }
  unit 0 {
    family inet {
      address 10.100.100.1/24;
    }
  }
}
```

With this configuration, the link stays up, so you can loop ping packets to a remote router. The `loopback local` statement causes the interface to loop within the PIC just before the data reaches the transceiver.

To turn off the loopback capability, remove the `loopback` statement from the configuration:

```
[edit]
user@host# delete interfaces t1-fpc/pic/port t1-options loopback
```

You can determine whether there is an internal problem or an external problem by checking the error counters in the output of the `show interface interface-name extensive` command, for example:

```
user@host> show interfaces t1-fpc/pic/port extensive
```

Configuring the T1 Idle Cycle Flag

By default, a T1 interface transmits the value 0x7E in the idle cycles. To have the interface transmit the value 0xFF (all ones) instead, include the `idle-cycle-flag` statement at the `[edit interfaces interface-name t1-options]` hierarchy level, specifying the `ones` option:

```
[edit interfaces interface-name t1-options]
idle-cycle-flag ones;
```

To explicitly configure the default value of 0x7E, include the `idle-cycle-flag` statement with the `flags` option:

```
[edit interfaces interface-name t1-options]
idle-cycle-flag flags;
```

Configuring T1 Start and End Flags

By default, a T1 interface shares the transmission of the start and end flags.

To configure a T1 interface to wait two idle cycles between the start and end flags, include the `start-end-flag` statement with the `filler` option at the `[edit interfaces interface-name t1-options]` hierarchy level:

```
[edit interfaces interface-name t1-options]
start-end-flag filler;
```

To revert to the default behavior, sharing the transmission of start and end flags, include the `start-end-flag` statement with the `shared` option at the `[edit interfaces interface-name t1-options]` hierarchy level:

```
[edit interfaces interface-name t1-options]
start-end-flag shared;
```

Configuring Fractional T1 Time Slots

By default, all the time slots on a T1 interface are used. To configure the number of time slots allocated to a fractional T1 interface, include the `timeslots` statement at the `[edit interfaces interface-name t1-options]` hierarchy level:

```
[edit interfaces interface-name t1-options]
timeslots time-slot-range;
```

For T1 interfaces, the time-slot range is from 1 through 24. There are 24 time slots on a T1 interface. You can designate any combination of time slots. To configure ranges, use hyphens. To configure discontinuous time slots, use commas. Do not include spaces.

Example: Configuring Fractional T1 Time Slots

Use Time Slots 1 Through 10	<code>[edit interfaces interface-name t1-options]</code> <code>timeslots 1-10;</code>
Use Time Slots 1 Through 5, 10, and 24	<code>[edit interfaces interface-name t1-options]</code> <code>timeslots 1-5,10,24;</code>
Use the First Four Odd-Numbered Time Slots	<code>[edit interfaces interface-name t1-options]</code> <code>timeslots 1,3,5,7;</code>

Chapter 32

Configuring T3 Interfaces

This chapter includes an overview of T3 interfaces and configuration information as follows:

- T3 Interfaces Overview on page 569
- Configuring T3 Physical Interface Properties on page 570
- Configuring T3 BERT Properties on page 570
- Disabling T3 C-Bit Parity Mode on page 571
- Configuring the T3 CSU Compatibility Mode on page 572
- Configuring the T3 Frame Checksum on page 574
- Configuring the T3 FEAC Response on page 575
- Configuring the T3 Idle Cycle Flag on page 575
- Configuring the T3 Line Buildout on page 575
- Configuring the Channelized T3 Loop Timing on page 576
- Configuring T3 Loopback Capability on page 576
- Configuring T3 HDLC Payload Scrambling on page 578
- Configuring T3 Start and End Flags on page 579
- Examples: Configuring T3 Interfaces on page 579

T3 Interfaces Overview

T3 is the physical layer protocol used by the Digital Signal level 3 (DS3) multiplexing method in North America. A T3 interface operates at a bit rate of 44.736 Mbps. The JUNOS Software supports payload scrambling and subrate operation on each physical T3 interface. One encapsulation format—Point-to-Point Protocol (PPP), Frame Relay, or High-level Data Link Control (HDLC)—must be configured for the interface. DS3 standards supported include:

- ANSI T1.107, T1.102
- GR 499-core, GR 253-core
- Bellcore TR-TSY-000009
- AT&T Pub 5404
- ITU G.751, G.703, G823

Configuring T3 Physical Interface Properties

To configure T3-specific physical interface properties, include the **t3-options** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
t3-options {
  bert-algorithm algorithm;
  bert-error-rate rate;
  bert-period seconds;
  (cbit-parity | no-cbit-parity);
  compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate value>;
  fcs (16 | 32);
  (feac-loop-respond | no-feac-loop-respond);
  idle-cycle-flag value;
  (long-buildout | no-long-buildout);
  (loop-timing | no-loop-timing);
  loopback (local | payload | remote);
  (payload-scrambler | no-payload-scrambler);
  start-end-flag value;
}
```

Configuring T3 BERT Properties

This section discusses BERT properties for the T3 interface specifically. For general information about the JUNOS implementation of the BERT procedure, see “Interface Diagnostics” on page 134.

You can configure a T3 interface to execute a bit error rate test (BERT) when the interface receives a request to run this test. You specify the duration of the test, the pattern to send in the bit stream, and the error rate to include in the bit stream by including the **bert-period**, **bert-algorithm**, and **bert-error-rate** statements at the [edit interfaces *interface-name* t3-options] hierarchy level:

```
[edit interfaces interface-name t3-options]
bert-algorithm algorithm;
bert-error-rate rate;
bert-period seconds;
```

By default, the BERT period is 10 seconds. You can configure the BERT period to last from 1 through 239 seconds on some PICs and from 1 through 240 seconds on other PICs.

rate is the bit error rate. This can be an integer from 0 through 7, which corresponds to a bit error rate from 10^{-0} (1 error per bit) to 10^{-7} (1 error per 10 million bits).

algorithm is the pattern to send in the bit stream. The default algorithm for the DS3 BERT procedure is **pseudo-2e15-o151** (pattern is $2^{15}-1$, as defined in the CCITT/ITU O.151 standard).

On T3 interfaces, you can also select the pattern to send in the bit stream by including the **bert-algorithm** statement at the [edit interfaces *interface-name* *interface-options*] hierarchy level:

```
[edit interfaces interface-name interface-options]
bert-algorithm algorithm;
```

For a list of supported algorithms, enter a ? after the **bert-algorithm** statement; for example:

```
[edit interfaces t3-0/0/0 t3-options]
user@host# set bert-algorithm ?
Possible completions:
all-ones-repeating Repeating one bits
all-zeros-repeating Repeating zero bits
alternating-double-ones-zeros Alternating pairs of ones and zeros
alternating-ones-zeros Alternating ones and zeros
pseudo-2e10 Pattern is 2^10 - 1
...
```

For specific hierarchy information, see individual interface types. For information about running the BERT procedure, see the *JUNOS System Basics and Services Command Reference*.

Disabling T3 C-Bit Parity Mode

C-bit parity mode controls the type of framing that is present on the transmitted T3 signal. When C-bit parity mode is enabled, the C-bit positions are used for the FEBE, FEAC, terminal data link, path parity, and mode indicator bits, as defined in ANSI T1.107a-1989. When C-bit parity mode is disabled, the basic T3 framing mode (M13) is used.

By default, C-bit parity mode is enabled. To disable C-bit parity mode and use M13 framing for your T3 link, include the **no-cbit-parity** statement at the [edit interfaces *interface-name* *t3-options*] hierarchy level:

```
[edit interfaces interface-name t3-options]
no-cbit-parity;
```



NOTE: For ATM and ATM2 IQ2 and IQ2-E interfaces, M23 framing is used when the **no-cbit-parity** statement is included. For all other interfaces, M13 framing is used when the **no-cbit-parity** statement is included.

To return to the default, enabling C-bit parity mode, delete the **no-cbit-parity** statement from the configuration:

```
[edit]
user@host# delete interfaces t3-fpc/pic/port t3-options no-cbit-parity
```

To explicitly enable C-bit parity mode, include the **cbit-parity** statement at the [edit interfaces *interface-name* *t3-options*] hierarchy level:

```
[edit interfaces interface-name t3-options]
cbit-parity;
```

Configuring the T3 CSU Compatibility Mode

Subrating a T3 interface reduces the maximum allowable peak rate by limiting the HDLC-encapsulated payload. Subrate modes configure the PIC to connect with channel service units (CSUs) that use proprietary methods of multiplexing.

You can configure T3 interfaces to be compatible with a Digital Link, Kentrox, or Larscom CSUs. For T3 intelligent queuing (IQ) channels only, you can also configure Adtran or Verilink CSU compatibility.



NOTE: To subrate an E3 interface to be compatible with a Kentrox CSU, you must have an IQ or IQE based PIC. Non-IQ or IQE PICs allow a commit of the configuration, but the interfaces remain at the full E3 rate for the Kentrox compatibility mode.

To configure a T3 interface so that it is compatible with the CSU at the remote end of the line, include the **compatibility** statement at the [edit interfaces *interface-name* t3-options] hierarchy level:

```
[edit interfaces interface-name t3-options]
compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate value>;
```

The subrate of a T3 interface must exactly match that of the remote CSU. To specify the subrate, include the **subrate** statement in the configuration:

- For Adtran CSUs, specify the subrate as a number from 1 through 588 that exactly matches the value configured on the CSU. A subrate value of 588 corresponds to 44.2 Mbps, or 100 percent of the HDLC-encapsulated payload. A subrate value of 1 corresponds to $44.2 / 588$, which is 75.17 Kbps, or 0.17 percent of the HDLC-encapsulated payload.
- For Digital Link CSUs, specify the subrate as the data rate you configured on the CSU in the format xKb or x.xMb. For Digital Link CSUs, you can specify the subrate value to match the data rate configured on the CSU in the format xkb or x.xMb. You can configure the subrate values shown in Table 47 on page 573.
- For Kentrox CSUs, specify the subrate as a number from 1 through 69 that exactly matches the value configured on the CSU. A subrate value of 69 corresponds to 34.995097 Mbps, or 79.17 percent of the HDLC-encapsulated payload (44.2 Mbps). A subrate value of 1 corresponds to 999.958 Kbps, which is 2.26 percent of the HDLC-encapsulated payload. Each increment of the subrate value corresponds to a rate increment of about 0.5 Mbps.
- For Larscom CSUs, specify the subrate as a number from 1 through 14 that exactly matches the value configured on the CSU. A subrate value of 14 corresponds to 44.2 Mbps, or 100 percent of the HDLC-encapsulated payload. A subrate value of 1 corresponds to $44.2 / 14$, which is 3.16 Mbps, 7.15 percent of the HDLC-encapsulated payload.
- For Verilink CSUs, specify the subrate as a number from 1 through 28 that exactly matches the value configured on the CSU. To calculate the maximum allowable

peak rate, multiply the configured subrate by 1.578 Mbps. For example, a subrate value of 28 corresponds to 28×1.578 Mbps, which is 44.2 Mbps, 100 percent of the HDLC-encapsulated payload. A subrate value of 1 corresponds to 1.578 Mbps, 3.57 percent of the HDLC-encapsulated payload. A subrate value of 20 corresponds to 20×1.578 Mbps, which is 31.56 Mbps, 71.42 percent of the HDLC-encapsulated payload.



NOTE: Verilink configuration is not functional if an IQ interface is paired with an IQE interface.

Table 47: Subrate Values for T3 Digital Link Compatibility Mode

301 Kbps	9.3 Mbps	18.3 Mbps	27.4 Mbps	36.4 Mbps
601 Kbps	9.6 Mbps	18.6 Mbps	27.7 Mbps	36.7 Mbps
902 Kbps	9.9 Mbps	18.9 Mbps	28.0 Mbps	37.0 Mbps
1.2 Mbps	10.2 Mbps	19.2 Mbps	28.3 Mbps	37.3 Mbps
1.5 Mbps	10.5 Mbps	19.5 Mbps	28.6 Mbps	37.6 Mbps
1.8 Mbps	10.8 Mbps	19.8 Mbps	28.9 Mbps	37.9 Mbps
2.1 Mbps	11.1 Mbps	20.1 Mbps	29.2 Mbps	38.2 Mbps
2.4 Mbps	11.4 Mbps	20.5 Mbps	29.5 Mbps	38.5 Mbps
2.7 Mbps	11.7 Mbps	20.8 Mbps	29.8 Mbps	38.8 Mbps
3.0 Mbps	12.0 Mbps	21.1 Mbps	30.1 Mbps	39.1 Mbps
3.3 Mbps	12.3 Mbps	21.4 Mbps	30.4 Mbps	39.4 Mbps
3.6 Mbps	12.6 Mbps	21.7 Mbps	30.7 Mbps	39.7 Mbps
3.9 Mbps	12.9 Mbps	22.0 Mbps	31.0 Mbps	40.0 Mbps
4.2 Mbps	13.2 Mbps	22.3 Mbps	31.3 Mbps	40.3 Mbps
4.5 Mbps	13.5 Mbps	22.6 Mbps	31.6 Mbps	40.6 Mbps
4.8 Mbps	13.8 Mbps	22.9 Mbps	31.9 Mbps	40.9 Mbps
5.1 Mbps	14.1 Mbps	23.2 Mbps	32.2 Mbps	41.2 Mbps
5.4 Mbps	14.4 Mbps	23.5 Mbps	32.5 Mbps	41.5 Mbps
5.7 Mbps	14.7 Mbps	23.8 Mbps	32.8 Mbps	41.8 Mbps
6.0 Mbps	15.0 Mbps	24.1 Mbps	33.1 Mbps	42.1 Mbps
6.3 Mbps	15.3 Mbps	24.4 Mbps	33.4 Mbps	42.4 Mbps

Table 47: Subrate Values for T3 Digital Link Compatibility Mode (*continued*)

6.6 Mbps	15.6 Mbps	24.7 Mbps	33.7 Mbps	42.7 Mbps
6.9 Mbps	15.9 Mbps	25.0 Mbps	34.0 Mbps	43.0 Mbps
7.2 Mbps	16.2 Mbps	25.3 Mbps	34.3 Mbps	43.3 Mbps
7.5 Mbps	16.5 Mbps	25.6 Mbps	34.6 Mbps	43.6 Mbps
7.8 Mbps	16.8 Mbps	25.9 Mbps	34.9 Mbps	43.9 Mbps
8.1 Mbps	17.1 Mbps	26.2 Mbps	35.2 Mbps	44.2 Mbps
8.4 Mbps	17.4 Mbps	26.5 Mbps	35.5 Mbps	
8.7 Mbps	17.7 Mbps	26.8 Mbps	35.8 Mbps	
9.0 Mbps	18.0 Mbps	27.1 Mbps	36.1 Mbps	

For information about subrating an E3 interface, see “Configuring the E3 CSU Compatibility Mode” on page 553.

Configuring the T3 Frame Checksum

By default, T3 interfaces use a 16-bit frame checksum. You can configure a 32-bit checksum, which provides more reliable packet verification. However, some older equipment might not support 32-bit checksums.

On a channelized OC12 interface, the **fcs** statement is not supported. To configure FCS on each DS3 channel, you must include the **t3-options fcs** statement in the configuration for each channel.

To configure a 32-bit checksum, include the **fcs** statement at the [edit interfaces *interface-name* t3-options] hierarchy level:

```
[edit interfaces interface-name t3-options]
fcs 32;
```

To return to the default 16-bit frame checksum, delete the **fcs 32** statement from the configuration:

```
[edit]
user@host# delete interfaces t3-fpc/pic/port t3-options fcs 32
```

To explicitly configure a 16-bit checksum, include the **fcs** statement at the [edit interfaces *interface-name* t3-options] hierarchy level:

```
[edit interfaces interface-name t3-options]
fcs 16;
```

Configuring the T3 FEAC Response

The T3 far-end alarm and control (FEAC) signal is used to send alarm or status information from the far-end terminal back to the near-end terminal and to initiate T3 loopbacks at the far-end terminal from the near-end terminal.

By default, the router does not respond to FEAC requests. To allow the remote CSU to place the local router into loopback, you must configure the router to respond to the CSU's FEAC request by including the **feac-loop-respond** statement at the [edit interfaces *interface-name* t3-options] hierarchy level:

```
[edit interfaces interface-name t3-options]
feac-loop-respond;
```

If you configure remote or local loopback with the T3 **loopback** statement, the router does not respond to FEAC requests from the CSU even if you include the **feac-loop-respond** statement in the configuration. For the router to respond, you must delete the **loopback** statement from the configuration.

To explicitly configure the router not to respond to FEAC requests, include the **no-feac-loop** statement at the [edit interfaces *interface-name* t3-options] hierarchy level:

```
[edit interfaces interface-name t3-options]
no-feac-loop-respond;
```

Configuring the T3 Idle Cycle Flag

By default, a T3 interface transmits the value 0x7E in the idle cycles. To have the interface transmit the value 0xFF (all ones) instead, include the **idle-cycle-flag** statement at the [edit interfaces *interface-name* t3-options] hierarchy level, specifying the **ones** option:

```
[edit interfaces interface-name t3-options]
idle-cycle-flag ones;
```

To explicitly configure the default value of 0x7E, include the **idle-cycle-flag** statement with the **flags** option:

```
[edit interfaces interface-name t3-options]
idle-cycle-flag flags;
```

Configuring the T3 Line Buildout

A T3 interface has two settings for the T3 line buildout: a short setting, which is less than 255 feet (about 68 meters), and a long setting, which is greater than 255 feet and less than 450 feet (about 137 meters). By default, the interface uses the short setting.

The **long-buildout** and **no-long-buildout** statements apply only to copper-cable-based T3 interfaces. You cannot configure a line buildout for a DS3 channel on a channelized

OC12 interface, which runs over fiber-optic cable. If you configure this statement on a channelized OC12 interface, it is ignored.

To have the interface drive a line that is longer than 255 feet and shorter than 450 feet, include the **long-buildout** statement at the [edit interfaces *interface-name* t3-options] hierarchy level:

```
[edit interfaces interface-name t3-options]
long-buildout;
```

To explicitly configure the default short line buildout, include the **no-long-buildout** statement at the [edit interfaces *interface-name* t3-options] hierarchy level:

```
[edit interfaces interface-name t3-options]
no-long-buildout;
```

Configuring the Channelized T3 Loop Timing

By default, internal clocking (line timing) is used on channelized IQ and IQE interfaces. To configure SONET/SDH or DS3-level external clocking, include the **loop-timing** statement:

```
loop-timing;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *ct3-fpc/pic/port* t3-options]
- [edit interfaces *stm1-fpc/pic/port* sonet-options]

To explicitly configure the default line timing, include the **no-loop-timing** statement in the configuration:

```
no-loop-timing;
```

The **loop-timing** and **no-loop-timing** statements apply only to E1 and T1 interfaces you configure on channelized IQ and IQE PICs. If you attempt to include these statements on any other interface type, they are ignored.

For all channelized IQ and IQE PICs, the **clocking** statement is supported on all channels. To configure clocking on individual interfaces, include the **clocking** statement at the [edit interfaces *type-fpc/pic/port:channel*] hierarchy level. If you do not include the **clocking** statement, the individual interfaces use internal clocking by default.

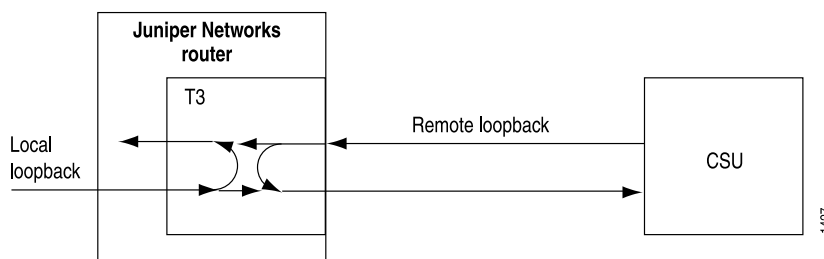
For more information, see “Configuring the Clock Source” on page 128 and “Clock Sources on Channelized Interfaces” on page 390.

Configuring T3 Loopback Capability

You can configure loopback capability between the local T3 interface and the remote CSU, as shown in Figure 58 on page 577. You can configure the loopback to be local or remote. With local loopback, the T3 interface can transmit packets to the CSU, but receives its own transmission back again and ignores data from the CSU. With

remote loopback, packets sent from the CSU are received by the T3 interface, forwarded if there is a valid route, and immediately retransmitted to the CSU.

Figure 58: Remote and Local T3 Loopback



To configure loopback capability on a T3 interface, include the `loopback` statement at the `[edit interfaces interface-name t3-options]` hierarchy level:

```
[edit interfaces interface-name t3-options]
  loopback (local | payload | remote);
```

Packets can be looped on either the local router or the remote CSU. Local and remote loopback loop back both data and clocking information.

To exchange BERT patterns between a local router and a remote router, include the `loopback remote` statement in the interface configuration at the remote end of the link. From the local router, you issue the `test interface` command.

For more information about configuring BERT, see “Interface Diagnostics” on page 134. For more information about using operational mode commands to test interfaces, see the *JUNOS System Basics and Services Command Reference*.

For channelized T3, T1, and NxDS0 IQ interfaces only, you can include the `loopback payload` statement in the configuration to loop back data only (without clocking information) on the remote router’s PIC. In payload loopback, overhead is recalculated. For T3 IQ interfaces, you can include the `loopback payload` statement at the `[edit interfaces ct3-fpc/pic/port]` and `[edit interfaces t3-fpc/pic/port:channel]` hierarchy levels. For T1 interfaces, you can include the `loopback payload` statement in the configuration at the `[edit interfaces t1-fpc/pic/port:channel]` hierarchy level; it is ignored if included at the `[edit interfaces ct1-fpc/pic/port]` hierarchy level. For NxDS0 interfaces, payload and remote loopback are the same. If you configure one, the other is ignored. NxDS0 IQ interfaces do not support local loopback.

To determine whether a problem is internal or external, you can loop packets on both the local and the remote router. To do this, include the `no-keepalives` and `encapsulation cisco-hdlc` statements at the `[edit interfaces interface-name]` hierarchy level and the `loopback local` statement at the `[edit interfaces interface-name t3-options]` hierarchy level, as shown in the following example:

```
[edit interfaces]
  t3-1/0/0 {
    no-keepalives;
    encapsulation cisco-hdlc;
    t3-options {
```

```

        loopback local;
    }
    unit 0 {
        family inet {
            address 10.100.100.1/24;
        }
    }
}

```

With this configuration, the link stays up, so you can loop ping packets to a remote router. The `loopback local` statement causes the interface to loop within the PIC just before the data reaches the transceiver.

To turn off the loopback capability, remove the `loopback` statement from the configuration:

```

[edit]
user@host# delete interfaces t3-fpc/pic/port t3-options loopback

```

You can determine whether there is an internal problem or an external problem by checking the error counters in the output of the `show interface interface-name extensive` command, for example:

```

user@host> show interfaces t3-fpc/pic/port extensive

```

For channel 0 on channelized interfaces only, you can include the `loopback` statement at the `[edit interfaces interface-name interface-type-options]` hierarchy level. The loopback setting configured for channel 0 applies to all channels on the channelized interface. The `loopback` statement is ignored if you include it at this hierarchy level in the configuration of other channels. To configure loopbacks on individual channels, you must include the `channel-type-options loopback` statement in the configuration for each channel. This allows each channel to be put in loopback mode independently.

For example, for DS3 channels on a channelized OC12 interface, the `sonet-options loopback` statement is supported only for channel 0; it is ignored if included in the configuration for channels 1 through 11. The SONET loopback configured for channel 0 applies to all 12 channels equally. To configure loopbacks on the individual DS3 channels, you must include the `t3-options loopback` statement in the configuration for each channel. This allows each DS3 channel can be put in loopback mode independently.

Configuring T3 HDLC Payload Scrambling

T3 HDLC payload scrambling, which is disabled by default, provides better link stability. Both sides of a connection must either use or not use scrambling.

On a channelized OC12 interface, the SONET `payload-scrambler` statement is ignored. To configure scrambling on the DS3 channels on the interface, you can include the `t3-options payload-scrambler` statement at the `[edit interfaces interface-name t3-options]` hierarchy level for each DS3 channel.

If you enable HDLC payload scrambling on a T3 interface, you must also configure the interface to be compatible with the channel service unit (CSU) at the remote end

of the line before you commit the interface configuration. For information about subrating a T3 interface, see “Configuring the T3 CSU Compatibility Mode” on page 572.

```
[edit interfaces interface-name t3-options]
compatibility-mode (adtran | digital-link | kentrox | larscom | verilink) <subrate value>;
payload-scrambler;
```

To explicitly disable HDLC payload scrambling, include the `no-payload-scrambler` statement at the `[edit interfaces interface-name t3-options]` hierarchy level:

```
[edit interfaces interface-name t3-options]
no-payload-scrambler;
```

To disable payload scrambling again (return to the default), delete the `payload-scrambler` statement from the configuration:

```
[edit]
user@host# delete interfaces t3-fpc/pic/port t3-options payload-scrambler
```

Configuring T3 Start and End Flags

By default, a T3 interface shares the transmission of the start and end flags.

To configure a T3 interface to wait two idle cycles between the start and end flags, include the `start-end-flag` statement with the `filler` option at the `[edit interfaces interface-name t3-options]` hierarchy level:

```
[edit interfaces interface-name t3-options]
start-end-flag filler;
```

To revert to the default behavior, sharing the transmission of start and end flags, include the `start-end-flag` statement with the `shared` option at the `[edit interfaces interface-name t3-options]` hierarchy level:

```
[edit interfaces interface-name t3-options]
start-end-flag shared;
```

Examples: Configuring T3 Interfaces

T3 interfaces can use PPP, Cisco HDLC, or Frame Relay encapsulation.

PPP Encapsulation on a DS3 PIC

```
[edit]
interfaces {
  t3-0/0/0 {
    encapsulation ppp;
    t3-options {
      no-long-buildout;
      compatibility-mode larscom;
      payload-scrambler;
    }
    unit 0 {
      family inet {
        address 10.0.0.1/32 {
```

```

        destination 10.0.0.2;
    }
}
family iso;
}
}
}

```

**Cisco HDLC
Encapsulation on a DS3
PIC**

```

[edit]
interfaces {
  t3-0/0/1 {
    encapsulation cisco-hdlc;
    t3-options {
      no-long-buildout;
      compatibility-mode larscom;
      payload-scrambler;
    }
    unit 0 {
      family inet {
        address 10.0.0.1/32 {
          destination 10.0.0.2;
        }
      }
      family iso;
    }
  }
}

```

Configure Frame Relay encapsulation on two routers, where one router is a DTE device and the other is a DCE device:

On DTE Router

```

[edit]
interfaces {
  t3-1/0/1 {
    encapsulation frame-relay;
    t3-options {
      no-long-buildout;
      compatibility-mode larscom;
      payload-scrambler;
    }
    unit 1 {
      dlci 1;
      family inet {
        address 10.0.0.1/32 {
          destination 10.0.0.2;
        }
      }
      family iso;
    }
  }
  unit 2 {
    dlci 2;
    family inet {
      address 10.0.0.3/32 {
        destination 10.0.0.4;
      }
    }
  }
}

```

```

    }
    family iso;
  }
}

```

On DCE Router

```

[edit]
interfaces {
  t3-1/1/1 {
    dce;
    encapsulation frame-relay;
    t3-options {
      no-long-buildout;
      compatibility-mode larscom;
      payload-scrambler;
    }
    unit 1 {
      dlci 1;
      family inet {
        address 10.0.0.2/32 {
          destination 10.0.0.1;
        }
      }
      family iso;
    }
    unit 2 {
      dlci 2;
      family inet {
        address 10.0.0.4/32 {
          destination 10.0.0.3;
        }
      }
      family iso;
    }
  }
}

```


Part 10

Configuring Ethernet Interfaces

- Configuring Ethernet Interfaces on page 585
- Configuring 802.1Q VLANs on page 599
- Configuring Aggregated Ethernet Interfaces on page 623
- Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 641
- Configuring Layer 2 Bridging Interfaces on page 663
- Configuring TCC and Layer 2.5 Switching on page 665
- Configuring Static ARP Table Entries on page 669
- Configuring Unrestricted Proxy ARP on page 671
- Configuring MAC Address Validation on Static Ethernet Interfaces on page 675
- Enabling Passive Monitoring on Ethernet Interfaces on page 677
- Configuring IEEE 802.1ag OAM Connectivity-Fault Management on page 679
- Configuring ITU-T Y.1731 Ethernet Service OAM on page 711
- Configuring IEEE 802.1x Port-Based Network Access Control on page 741
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 745
- Configuring VRRP and VRRP for IPv6 on page 753
- Configuring Gigabit Ethernet Accounting and Policing on page 755
- Configuring Gigabit Ethernet Autonegotiation on page 767
- Configuring Gigabit Ethernet OTN Options on page 773
- Configuring the Management Ethernet Interface on page 775
- Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength on page 779
- Configuring 10-Gigabit Ethernet Framing on page 781
- Configuring 10-Gigabit Ethernet Notification of Link Down Alarm on page 783
- Configuring Point-to-Point Protocol over Ethernet on page 785
- Configuring Ethernet Ring Protection Switching on page 799
- Example Ethernet Configurations on page 813

Chapter 33

Configuring Ethernet Interfaces

You can configure the following properties specific to aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces:

- Ethernet Interfaces Overview on page 585
- Configuring Ethernet Physical Interface Properties on page 586
- Configuring J Series Services Router Switching Interfaces on page 589
- MX Series Router Interface Identifiers on page 591
- Enabling Ethernet MAC Address Filtering on page 591
- Configuring Ethernet Loopback Capability on page 593
- Configuring Flow Control on page 594
- Ignoring Layer 3 Incomplete Errors on page 594
- Configuring the Link Characteristics on Ethernet Interfaces on page 595
- Configuring Gratuitous ARP on page 596
- Adjusting the ARP Aging Timer on page 596
- Configuring the Interface Speed on Ethernet Interfaces on page 597
- Configuring the Ingress Rate Limit on page 597
- Configuring Weighted Random Early Detection on page 598

Ethernet Interfaces Overview

Ethernet was developed in the early 1970s at the Xerox Palo Alto Research Center (PARC) as a data-link control layer protocol for interconnecting computers. It was first widely used at 10 megabits per second (Mbps) over coaxial cables and later over unshielded twisted pairs using 10Base-T. More recently, 100Base-TX (Fast Ethernet, 100 Mbps), Gigabit Ethernet (1 gigabit per second [Gbps]), and 10-Gigabit Ethernet (10 Gbps) have become available.

Juniper Networks routers support the following types of Ethernet interfaces:

- Fast Ethernet
- Tri-Rate Ethernet copper
- Gigabit Ethernet
- Gigabit Ethernet intelligent queuing (IQ)

- Gigabit Ethernet IQ2 and IQ2-E
- 10-Gigabit Ethernet IQ2 and IQ2-E
- 10-Gigabit Ethernet
- 10-Gigabit Ethernet dense wavelength-division multiplexing (DWDM)
- Management Ethernet interface, which is an out-of-band management interface within the router
- Internal Ethernet interface, which connects the Routing Engine to the packet forwarding components
- Aggregated Ethernet interface, a logical linkage of Fast Ethernet, Gigabit Ethernet, or 10-Gigabit Ethernet physical connections

Configuring Ethernet Physical Interface Properties

To configure Fast Ethernet-specific physical interface properties, include the `fastether-options` statement at the `[edit interfaces fe-fpc/pic/port]` hierarchy level:

```
[edit interfaces fe-fpc/pic/port]
link-mode (full-duplex | half-duplex);
speed (10m | 100m);
vlan-tagging;
fastether-options {
    802.3ad aex (primary | backup);
    (flow-control | no-flow-control);
    ignore-l3-incompletes;
    ingress-rate-limit rate;
    (loopback | no-loopback);
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
}
```



NOTE: The `speed` statement applies to the management Ethernet interface (`fxp0` or `em0`), the Fast Ethernet 12-port and 48-port Physical Interface Card (PIC) interfaces, the J Series Gigabit Ethernet uPIM interfaces and the MX Series Tri-Rate Ethernet copper interfaces. The Fast Ethernet, `fxp0`, and `em0` interfaces can be configured for 10 Mbps or 100 Mbps (`10m` | `100m`). The J Series Gigabit Ethernet uPIM interfaces and the MX Series Tri-Rate Ethernet copper interfaces can be configured for 10 Mbps, 100 Mbps, or 1 Gbps (`10m` | `100m` | `1g`). The 4-port and 8-port Fast Ethernet PICs support a speed of 100 Mbps only.



NOTE: JUNOS Software supports Ethernet host addresses with no subnets. This enables you to configure an Ethernet interface as a host address (that is, with a network mask of `/32`), without requiring a subnet. Such interfaces can serve as OSPF point-to-point interfaces, and MPLS is also supported.

To configure physical interface properties specific to Gigabit Ethernet and 10-Gigabit Ethernet, include the `gigether-options` statement at the `[edit interfaces ge-fpc/pic/port]` or `[edit interfaces xe-fpc/pic/port]` hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
gigether-options {
  802.3ad aex (primary | backup);
  auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
  local-interface-offline>
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  (loopback | no-loopback);
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
}
```

Additionally, for 10-Gigabit Ethernet DWDM-specific physical interface properties, include the `optics-options` statement at the `[edit interfaces ge-fpc/pic/port]` hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
optics-options {
  wavelength nm;
}
```

To configure Gigabit Ethernet IQ-specific physical interface properties, include the `gigether-options` statement at the `[edit interfaces ge-fpc/pic/port]` hierarchy level. These statements are supported on 10-Gigabit Ethernet IQ2 and IQ2-E PIC. Some of these statements are also supported on Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router). For more information, see “Example: Configuring Gigabit Ethernet Interfaces” on page 813.

```
[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
gigether-options {
  802.3ad aex (primary | backup);
  auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
  local-interface-offline>
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  (loopback | no-loopback);
  (source-filtering | no-source-filtering);
  ethernet-switch-profile {
    (mac-learn-enable | no-mac-learn-enable);
    tag-protocol-id [tpids ];
    ethernet-policer-profile {
      input-priority-map {
        ieee802.1p premium [values ];
      }
      output-priority-map {
        classifier {
          premium {
```

```

        forwarding-class class-name {
            loss-priority (high | low);
        }
    }
}
policer cos-policer-name {
    aggregate {
        bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
        burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
    }
    premium {
        bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
        burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
    }
}
}
}
native-vlan-id number;
}

```

To configure 10-Gigabit Ethernet IQ2 and IQ2-E-specific physical interface properties, include the `lan-phy` or `wan-phy` statement at the [edit interfaces *xe-fpc/pic/port* framing] hierarchy level. For more information, see “Configuring 10-Gigabit Ethernet Framing” on page 781.

```
[edit interfaces]
xe-0/0/0 {
    framing {
        (lan-phy | wan-phy);
    }
}
```

To configure OAM 802.3ah support for Ethernet interfaces, include the `oam` statement at the `[edit protocols]` hierarchy level.

```
oam {
  ethernet {
    link-fault-management {
      interfaces {
        interface-name {
          pdu-interval interval;
          link-discovery (active | passive);
          pdu-threshold count;
        }
      }
    }
  }
}
```

To configure Gigabit Ethernet IQ-specific logical interface properties, include the `input-vlan-map`, `output-vlan-map`, and `layer2-policer` statements:

```
input-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
}
```

```

    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
vlan-tags (Stacked VLAN Tags) inner tpid.vlan-id outer tpid.vlan-id;

```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

To configure aggregated Ethernet-specific physical interface properties, include the `aggregated-ether-options` statement at the [edit interfaces *aex*] hierarchy level:

```

[edit interfaces aex]
aggregated-ether-options {
    ethernet-switch-profile {
        tag-protocol-id tpid;
    }
    (flow-control | no-flow-control);
    lacp mode {
        periodic interval;
    }
    link-protection;
    link-speed speed;
    (loopback | no-loopback);
    minimum-links number;
    source-address-filter {
        mac-address;
    }
    (source-filtering | no-source-filtering);
}

```

Configuring J Series Services Router Switching Interfaces

The J Series routers with multiport Gigabit Ethernet uPIMs supports Ethernet access switching. This functionality provides the ability to switch traffic at Layer 2 in addition to routing traffic at Layer 3.

J Series routers with multiport Gigabit Ethernet uPIMs can be deployed in branch offices as an access or desktop switch with integrated routing capability. The multiport Gigabit Ethernet uPIM provides Ethernet switching, while the Routing Engine provides routing functionality.

Routed traffic is forwarded from any port of the multiport Gigabit Ethernet uPIM to the WAN interface. Switched traffic is forwarded from one port of the multiport Gigabit Ethernet uPIM to another port on the same the multiport Gigabit Ethernet uPIM. Switched traffic is not forwarded from a port on one multiport Gigabit Ethernet uPIM to a port on a different multiport Gigabit Ethernet uPIM. For more information about configuring the multiport Gigabit Ethernet uPIM switching mode, see the *JUNOS System Basics Configuration Guide*.

In access switching mode, only one physical interface is configured for the entire multiport Gigabit Ethernet uPIM. The single physical interface serves as a Virtual Router Interface (VRI). Configuration of the physical port characteristics is done under the single physical interface.

To configure multiport Gigabit Ethernet uPIM Ethernet port properties, include the `switch-port` statement at the `[edit interfaces ge-pim/0/0]` hierarchy level:

```
[edit interfaces ge-pim/0/0]
switch-options {
  switch-port port-number {
    (auto-negotiation | no-auto-negotiation);
    speed 1g;
    link-mode (full-duplex | half-duplex);
  }
}
```

Access switching mode is supported on the 6-port, 8-port, and 16-port Gigabit Ethernet uPIMs.

The multiport Gigabit Ethernet uPIMs are supported on the J2320, J2350, J4350, and J6350 Services Routers.

The 6-port and 8-port multiport Gigabit Ethernet uPIM occupies a single slot and can be installed in any slot. Because the 16-port Gigabit Ethernet uPIM is two slots high, you cannot install a 16-port uPIM in the top slots (slots 1 and 4). Ports are numbered 0 through 5 on the 6-port Gigabit Ethernet uPIM, 0 through 7 on the 8-port Gigabit Ethernet uPIM, and 0 through 15 on the 16-port Gigabit Ethernet uPIM.

Example: Configuring J Series Services Router Switching Interfaces

Configure a single physical interface for the uPIM and set the port parameters for port 0 and port 1:

```
[edit interfaces]
ge-2/0/0 {
  switch-options {
    switch-port 0 {
      no-auto-negotiation;
      speed 1g;
      link-mode full-duplex;
```

```

    }
    switch-port 1 {
        no-auto-negotiation;
        speed 10m;
        link-mode half-duplex;
    }
}

```

MX Series Router Interface Identifiers

The MX Series routers use the convention *ge-fpc/pic/port* to identify interfaces.

fpc identifies the number of the FPC or DPC card on which the physical interface is located. Specifically, it is the number of the slot in which the card is installed. If two Switch Control Boards (SCBs) are installed in an MX960 router, the FPC range is from 0 through 11. If three SCBs are installed, the range is from 0 through 5 and from 7 through 11. On the MX480 router, the range is 0 through 6. On the MX240 router, the range is 0 through 2 with one SCB installed. With two SCBs installed, the DPC can be installed in any of two consecutive interface card slots, the range is from 1 through 2.

For DPCs, the PIC and port numbers are identified on the DPC front panel. Use the PIC and port numbers that correspond to the port you are configuring.

Ports are numbered from 0 through 9 for Gigabit Ethernet and Tri-Rate Ethernet copper interfaces. Port numbers are always 0 for 10-Gigabit Ethernet interfaces.



NOTE: In certain displays, the MX Series routers identify the Packet Forwarding Engine (PFE) rather than the PIC number. PFE 0 corresponds to PIC 0, PFE 1 corresponds to PIC 2, PFE 2 corresponds to PIC 1, and PFE 3 corresponds to PIC 3.

Enabling Ethernet MAC Address Filtering

By default, source address filtering is disabled. On aggregated Ethernet interfaces, Fast Ethernet, Gigabit Ethernet, Gigabit Ethernet IQ, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can enable source address filtering, which blocks all incoming packets to an interface.



NOTE: Source address filtering is not supported on J Series Services Routers.

To enable the filtering, include the **source-filtering** statement:

```
source-filtering;
```

To explicitly disable filtering, include the **no-source-filtering** statement:

```
no-source-filtering;
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]



NOTE: When you integrate a standalone T640 router into a routing matrix, the PIC media access control (MAC) addresses for the integrated T640 router are derived from a pool of MAC addresses maintained by the TX Matrix router. For each MAC address you specify in the configuration of a formerly standalone T640 router, you must specify the same MAC address in the configuration of the TX Matrix router.

Similarly, when you integrate a standalone T1600 router into a routing matrix, the PIC MAC addresses for the integrated T1600 router are derived from a pool of MAC addresses maintained by the TX Matrix Plus router. For each MAC address you specify in the configuration of a formerly standalone T1600 router, you must specify the same MAC address in the configuration of the TX Matrix Plus router.

Filtering Specific MAC Addresses

When source address filtering is enabled, you can configure the interface to receive packets from specific MAC addresses. To do this, specify the MAC addresses in the `source-address-filter` statement:

```
source-address-filter {
  mac-address;
  <additional-mac-address>;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]

You can specify the MAC address as `nn:nn:nn:nn:nn:nn` or `nnnn .nnnn.nnnn`, where `n` is a hexadecimal number. You can configure up to 64 source addresses. To specify more than one address, include the `source-address-filter` statement multiple times.



NOTE: The `source-address-filter` statement is not supported on Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router); instead, include the `accept-source-mac` statement. For more information, see “Configuring MAC Address Filtering” on page 761.

If the remote Ethernet card is changed, the interface cannot receive packets from the new card because it has a different MAC address.

Source address filtering does not work when Link Aggregation Control Protocol (LACP) is enabled. For more information about LACP, see “Configuring Aggregated Ethernet LACP” on page 627.



NOTE: On untagged Gigabit Ethernet interfaces, you should not configure the `source-address-filter` statement at the [edit interfaces *ge-fpc/pic/port* *gigether-options*] hierarchy level and the `accept-source-mac` statement at the [edit interfaces *ge-fpc/pic/port* *gigether-options* unit *logical-unit-number*] hierarchy level simultaneously. If these statements are configured for the same interfaces at the same time, an error message is displayed.

On tagged Gigabit Ethernet interfaces, you should not configure the `source-address-filter` statement at the [edit interfaces [edit interfaces *ge-fpc/pic/port* *gigether-options*] hierarchy level and the `accept-source-mac` statement at the [edit interfaces *ge-fpc/pic/port* *gigether-options* unit *logical-unit-number*] hierarchy level with an identical MAC address specified in both filters. If these statements are configured for the same interfaces with an identical MAC address specified, an error message is displayed.

Configuring Ethernet Loopback Capability

By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system. To place an interface in loopback mode, include the `loopback` statement:

```
loopback;
```



NOTE: If you configure a local loopback on a 1-port 10-Gigabit IQ2 and IQ2-E PIC using the `loopback` statement at the [edit interfaces *interface-name* *gigether-options*] hierarchy level, the transmit-path stops working, causing the remote end to detect a link down.

To return to the default—that is, to disable loopback mode—delete the `loopback` statement from the configuration:

```
[edit]
user@host# delete interfaces fe-fpc/pic/port fastether-options loopback
```

To explicitly disable loopback mode, include the **no-loopback** statement:

```
no-loopback;
```

You can include the **loopback** and **no-loopback** statements at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]

Configuring Flow Control

By default, the routing platform imposes flow control to regulate the amount of traffic sent out on a Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interface. Flow control is not supported on the 4-port Fast Ethernet PIC. This is useful if the remote side of the connection is a Fast Ethernet or Gigabit Ethernet switch.

You can disable flow control if you want the routing platform to permit unrestricted traffic. To disable flow control, include the **no-flow-control** statement:

```
no-flow-control;
```

To explicitly reinstate flow control, include the **flow-control** statement:

```
flow-control;
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* aggregated-ether-options]
- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]

Ignoring Layer 3 Incomplete Errors

By default, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces count Layer 3 incomplete errors. You can configure the interface to ignore Layer 3 incomplete errors.

To ignore Layer 3 incomplete errors, include the **ignore-l3-incompletes** statement:

```
ignore-l3-incompletes;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* fastether-options]
- [edit interfaces *interface-name* gigether-options]

Configuring the Link Characteristics on Ethernet Interfaces

Full-duplex communication means that both ends of the communication can send and receive signals at the same time. *Half-duplex* is also bidirectional communication, but signals can flow in only one direction at a time.

By default, the router's management Ethernet interface, `fxp0` or `em0`, autonegotiates whether to operate in full-duplex or half-duplex mode. J Series Gigabit Ethernet interfaces and Fast Ethernet interfaces, except the J Series ePIM Fast Ethernet interfaces, can operate in either full-duplex or half-duplex mode, and all other interfaces can operate only in full-duplex mode. For Gigabit Ethernet and 10-Gigabit Ethernet, the link partner must also be set to full duplex.



NOTE: For M Series, MX Series, and most T Series routers, the management Ethernet interface is `fxp0`. For TX Matrix Plus routers and T1600 routers configured in a routing matrix, the management Ethernet interface is `em0`.



NOTE: Automated scripts that you have developed for standalone T1600 routers (T1600 routers that are not in a routing matrix) might contain references to the `fxp0` management Ethernet interface. Before reusing the scripts on T1600 routers in a routing matrix, edit the command lines that reference the `fxp0` management Ethernet interface so that the commands reference the `em0` management Ethernet interface instead.



NOTE: When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.



NOTE: On a J Series ePIM Fast Ethernet interface, if you specify half-duplex (or if full-duplex mode is not autonegotiated), the following message is written to the system log: "Half-duplex mode not supported on this PIC, forcing full-duplex mode."



NOTE: When you manually configure Fast Ethernet interfaces on the M Series and T Series routers, link mode and speed must both be configured. If both these values are not configured, the router uses autonegotiation for the link and ignores the user-configured settings.

To explicitly configure an Ethernet interface to operate in either full-duplex or half-duplex mode, include the `link-mode` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]
link-mode (full-duplex | half-duplex);
```

Configuring Gratuitous ARP

Gratuitous Address Resolution Protocol (ARP) requests provide duplicate IP address detection. A gratuitous ARP request is a broadcast request for a router's own IP address. If a router sends an ARP request for its own IP address and no ARP replies are received, the router's assigned IP address is not being used by other nodes. If a router sends an ARP request for its own IP address and an ARP reply is received, the router's assigned IP address is already being used by another node.

By default, the router responds to gratuitous ARP requests. On Ethernet interfaces, you can disable responses to gratuitous ARP requests. To disable responses to gratuitous ARP requests, include the `no-gratuitous-arp-request` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
no-gratuitous-arp-request;
```

To return to the default—that is, to respond to gratuitous ARP requests—delete the `no-gratuitous-arp-request` statement from the configuration:

```
[edit]  
user@host# delete interfaces interface-name no-gratuitous-arp-request
```

Gratuitous ARP replies are reply packets sent to the broadcast MAC address with the target IP address set to be the same as the sender's IP address. When the router receives a gratuitous ARP reply, the router can insert an entry for that reply in the ARP cache.

By default, updating the ARP cache on gratuitous ARP replies is disabled on the router. On Ethernet interfaces, you can enable handling of gratuitous ARP replies on a specific interface by including the `gratuitous-arp-reply` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
gratuitous-arp-reply;
```

To restore the default behavior, include the `no-gratuitous-arp-reply` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
no-gratuitous-arp-reply;
```

Adjusting the ARP Aging Timer

By default, the ARP aging timer is set at 20 minutes. In most network environments, this default value does not cause a problem. However, in environments with many directly attached hosts, such as metro Ethernet, the number of ARP entries to update can be high. In such environments, you might want to increase the amount of time between ARP updates by configuring the ARP aging timer.

To configure the ARP aging timer, include the `aging-timer` statement at the `[edit system arp]` hierarchy level:

```
[edit system arp]
aging-timer minutes;
```

The aging timer range is from 20 through 240 minutes. The timer value you configure takes effect as ARP entries expire. In other words, each subsequent refreshed ARP entry receives the new timer value. The new timer value does not apply to ARP entries that exist at the time you commit the configuration.

For more information about statements you can configure at the [edit system] hierarchy level, see the *JUNOS System Basics Configuration Guide*.

Configuring the Interface Speed on Ethernet Interfaces

For M Series and T Series Fast Ethernet 12-port and 48-port PIC interfaces, the management Ethernet interface (fxp0 or em0), the J Series Gigabit Ethernet uPIM interfaces, and the MX Series Tri-Rate Ethernet copper interfaces, you can explicitly set the interface speed. The Fast Ethernet, fxp0, and em0 interfaces can be configured for 10 Mbps or 100 Mbps (10m | 100m). The J Series Gigabit Ethernet uPIM interfaces and the MX Series Tri-Rate Ethernet copper interfaces can be configured for 10 Mbps, 100 Mbps, or 1 Gbps (10m | 100m | 1g). MX Series routers, with MX-DPC and Tri-Rate Copper SFPs, support 20x1 Copper to provide backwards compatibility with 100/10BASE-T and 1000BASE-T operation through an Serial Gigabit Media Independent Interface (SGMII) interface.



NOTE: On MX Series routers with tri-rate copper SFP interfaces, if the port speed is negotiated to the configured value and the negotiated speed and interface speed do not match, the link will not be brought up.



NOTE: When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.

To explicitly configure the speed, include the **speed** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
speed (10m | 100m | 1g);
```

Configuring the Ingress Rate Limit

On Fast Ethernet 8-port, 12-port, and 48-port PIC interfaces only, you can apply port-based rate limiting to the ingress traffic that arrives at the PIC.

To configure an ingress rate limit on a Fast Ethernet 8-port, 12-port, or 48-port PIC interface, include the **ingress-rate-limit** statement at the [edit interfaces *interface-name* fastether-options] hierarchy level:

```
[edit interfaces interface-name fastether-options]
ingress-rate-limit rate;
```

rate can range in value from 1 through 100 Mbps.

Configuring Weighted Random Early Detection

On M7i, M10i, M40e, M320, M120, and T Series routers, the Ethernet IQ2 and IQ2-E PIC families extend CoS functionality by supporting network congestion avoidance with weighted random early detection (WRED).

Related Topics For information on configuring WRED, see the *JUNOS Class of Service Configuration Guide*.

Chapter 34

Configuring 802.1Q VLANs

For examples of 802.1Q VLAN configuration, see the following sections:

- 802.1Q VLANs Overview on page 599
- Configuring Dynamic 802.1Q VLANs on page 600
- 802.1Q VLAN IDs and Ethernet Interface Types on page 600
- Enabling VLAN Tagging on page 601
- Binding VLAN IDs to Logical Interfaces on page 604
- Configuring VLAN Encapsulation on page 609
- Configuring Extended VLAN Encapsulation on page 610
- Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs on page 612
- Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 614
- Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 615
- Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface on page 617
- Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface on page 618
- Configuring a Logical Interface for Access Mode on page 619
- Configuring a Logical Interface for Trunk Mode on page 620
- Configuring the VLAN ID List for a Trunk Interface on page 620
- Configuring a Trunk Interface on a Bridge Network on page 621

802.1Q VLANs Overview

For Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and aggregated Ethernet interfaces supporting VPLS, the JUNOS Software supports a subset of the IEEE 802.1Q standard for channelizing an Ethernet interface into multiple logical interfaces, allowing many hosts to be connected to the same Gigabit Ethernet switch, but preventing them from being in the same routing or bridging domain.

Configuring Dynamic 802.1Q VLANs

You can configure the router to dynamically create VLANs when a client accesses an interface and requests a VLAN ID that does not yet exist. When a client accesses a VLAN interface, the router instantiates a VLAN dynamic profile that you have associated with the interface. Using the settings in the dynamic profile, the router extracts information about the client from the incoming packet (for example, the interface and unit values), saves this information in the routing table, and creates a VLAN or stacked VLAN ID for the client from a range of VLAN IDs that you configure for the interface.

Dynamically configuring VLANs or stacked VLANs requires the following general steps:

1. Configure a dynamic profile for dynamic VLAN or dynamic stacked VLAN creation.
2. Associate the VLAN or stacked VLAN dynamic profile with the interface.
3. Specify the Ethernet packet type that the VLAN dynamic profile accepts.
4. Define VLAN ranges for use by the dynamic profile when creating VLAN IDs.

For procedures on how to configure dynamic VLANs and dynamic stacked VLANs for client access, see the *JUNOS Subscriber Access Configuration Guide*.

802.1Q VLAN IDs and Ethernet Interface Types

You can partition the router into up to 4095 different VLANs—depending on the router model and the physical interface types—by associating logical interfaces with specific VLAN IDs.

VLAN ID 0 is reserved for tagging the priority of frames. VLAN IDs 1 through 511 are reserved for normal VLANs. VLAN IDs 512 and above are reserved for VLAN circuit cross-connect (CCCs).

For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can configure flexible Ethernet services encapsulation on the physical interface. With flexible Ethernet services encapsulation, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

Table 48 on page 600 lists VLAN ID range by interface type.

Table 48: VLAN ID Range by Interface Type

Interface Type	VLAN ID Range
Aggregated Ethernet for Fast Ethernet	1 through 1023
Aggregate Ethernet for Gigabit Ethernet	1 through 4094
4-port, 8-port, and 12-port Fast Ethernet	1 through 1023

Table 48: VLAN ID Range by Interface Type (*continued*)

Interface Type	VLAN ID Range
48-port Fast Ethernet	1 through 4094
Tri-Rate Ethernet copper	1 through 4094
Gigabit Ethernet	1 through 4094
Gigabit Ethernet IQ	1 through 4094
10-Gigabit Ethernet	1 through 4094
Management and internal Ethernet interfaces	1 through 1023



NOTE: For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the built-in Gigabit Ethernet port on the M7i router), VLAN IDs on a single interface can differ from each other.

Because IS-IS has an 8-bit limit for broadcast multiaccess media, you cannot set up more than 255 adjacencies over Gigabit Ethernet using VLAN tagging. For more information, see the *JUNOS Routing Protocols Configuration Guide*.

Enabling VLAN Tagging

You can configure the router to receive and forward single-tag frames, dual-tag frames, or a mixture of single-tag and dual-tag frames. For more information, see the following sections:

- Configuring Single-Tag Framing on page 602
- Configuring Dual Tagging on page 602
- Configuring Mixed Tagging on page 602
- Configuring Mixed Tagging Support for Untagged Packets on page 603
- Example: Configuring Mixed Tagging on page 603
- Example: Configuring Mixed Tagging to Support Untagged Packets on page 604



NOTE: If you configure VLAN tagging on Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces on M320, M120, and T Series routers, the JUNOS Software creates an internal logical interface that reserves 50 Kbps of bandwidth from Gigabit Ethernet IQ interfaces and 2 Mbps of bandwidth from Gigabit Ethernet IQ2 and IQ2-E interfaces. As a result, the effective available bandwidth for these interface types is now 999.5 Mbps and 998 Mbps, respectively.

Configuring Single-Tag Framing

To configure the router to receive and forward single-tag frames with 802.1Q VLAN tags, include the `vlan-tagging` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
vlan-tagging;
```

Configuring Dual Tagging

To configure the routing platform to receive and forward dual-tag frames with 802.1Q VLAN tags, include the `stacked-vlan-tagging` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
stacked-vlan-tagging;
```

Configuring Mixed Tagging

Mixed tagging is supported for Gigabit Ethernet interfaces on Gigabit Ethernet IQ2 and IQ2-E, and IQ or IQE PICs on M Series and T Series routers, for all MX Series router Gigabit and 10-Gigabit Ethernet interfaces, and for aggregated Ethernet interfaces with member links in IQ2 and IQ2-E PICs or in MX Series DPCs. Mixed tagging lets you configure two logical interfaces on the same Ethernet port, one with single-tag framing and one with dual-tag framing.



NOTE: Mixed tagging is not supported on Fast Ethernet interfaces or on J Series Services Routers.

To configure mixed tagging, include the `flexible-vlan-tagging` statement at the `[edit interfaces ge-fpc/pic/port]` hierarchy level. You must also include the `vlan-tags` statement with `inner` and `outer` options or the `vlan-id` statement at the `[edit interfaces ge-fpc/pic/port unit logical-unit-number]` hierarchy level:

```
[edit interfaces ge-fpc/pic/port]  
flexible-vlan-tagging;  
unit logical-unit-number {  
  vlan-id number;  
  family family {  
    address address;  
  }  
}  
unit logical-unit-number {  
  vlan-tags (Stacked VLAN Tags) inner tpid.vlan-id outer tpid.vlan-id;  
  family family {  
    address address;  
  }  
}
```



NOTE: When you configure the physical interface MTU for mixed tagging, you must increase the MTU to 4 bytes more than the MTU value you would configure for a standard VLAN-tagged interface.

For example, if the MTU value is configured to be 1018 on a VLAN-tagged interface, then the MTU value on a flexible VLAN tagged interface must be 1022—4 bytes more. The additional 4 bytes accommodates the future addition of a stacked VLAN tag configuration on the same physical interface.

If the same physical interface MTU value is configured on both the VLAN and flexible VLAN-tag routers, the L2 circuit configuration does not come up and a MTU mismatch is logged. However, normal traffic flow is unaffected.

Configuring Mixed Tagging Support for Untagged Packets

For 1-, 4-, and 8-port Gigabit Ethernet IQ2 and IQ2-E PICs, for 1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs, for all MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces configured for 802.1Q flexible VLAN tagging, and for aggregated Ethernet interfaces on IQ2 and IQ2-E PICs or MX Series DPCs, you can configure mixed tagging support for untagged packets on a port. Untagged packets are accepted on the same mixed VLAN-tagged port. To accept untagged packets, include the `native-vlan-id` statement and the `flexible-vlan-tagging` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces ge-fpc/pic/port]
flexible-vlan-tagging;
native-vlan-id number;
```

The logical interface on which untagged packets are to be received must be configured with the same native VLAN ID as that configured on the physical interface. To configure the logical interface, include the `vlan-id` statement (matching the `native-vlan-id` statement on the physical interface) at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level.

Example: Configuring Mixed Tagging

The following example configures mixed tagging. Dual-tag and single-tag logical interfaces are under the same physical interface:

```
[edit interfaces ge-3/0/1]
flexible-vlan-tagging;
unit 0 {
  vlan-id 232;
  family inet {
    address 10.66.1.2/30;
  }
}
unit 1 {
  vlan-tags outer 0x8100.222 inner 0x8100.221;
  family inet {
    address 10.66.1.2/30;
  }
}
```

```
}
}
```

For information about binding VLAN IDs to logical interfaces, see “Binding VLAN IDs to Logical Interfaces” on page 604. For information about configuring dual VLAN tags using the `vlan-tag` statement, see “Stacking a VLAN Tag” on page 648.

Example: Configuring Mixed Tagging to Support Untagged Packets

The following example configures untagged packets to be mapped to logical unit number 0:

```
[edit interfaces ge-0/2/0]
flexible-vlan-tagging;
native-vlan-id 232;
unit 0 {
  vlan-id 232;
  family inet {
    address 10.66.1.2/30;
  }
}
unit 1 {
  vlan-tags outer 0x8100.222 inner 0x8100.221;
  family inet {
    address 10.66.1.2/30;
  }
}
```

Binding VLAN IDs to Logical Interfaces

The following sections describe how to configure logical interfaces to receive and forward VLAN-tagged frames:

- Binding VLAN IDs to Logical Interfaces Overview on page 604
- Binding a VLAN ID to a Logical Interface on page 605
- Binding a Range of VLAN IDs to a Logical Interface on page 606
- Binding a List of VLAN IDs to a Logical Interface on page 607

Binding VLAN IDs to Logical Interfaces Overview

To configure a logical interface to receive and forward VLAN-tagged frames, you must bind a VLAN ID, a range of VLAN IDs, or a list of VLAN IDs to the logical interface. Table 49 on page 605 lists the configuration statements you use to bind VLAN IDs to logical interfaces, organized by scope of the VLAN IDs used to match incoming packets:

Table 49: Configuration Statements Used to Bind VLAN IDs to Logical Interfaces

Scope of VLAN ID Matching	Type of VLAN Framing Supported on the Logical Interface	
	Single-Tag Framing	Dual-Tag Framing
VLAN ID	<code>vlan-id <i>vlan-id</i>;</code>	<code>vlan-tags outer <i>tpid</i>.<<i>vlan-id</i>> inner <i>tpid</i><i>vlan-id</i>;</code>
VLAN ID Range	<code>vlan-id-range <i>vlan-id-vlan-id</i>;</code>	<code>vlan-tags outer <<i>tpid</i>.><i>vlan-id</i> inner-range <i>tpid.vlan-id-vlan-id</i>;</code>
VLAN ID List	<code>vlan-id-list [<i>vlan-id vlan-id-vlan-id</i>];</code>	<code>vlan-tags outer <<i>tpid</i>.><i>vlan-id</i> inner-list [<i>vlan-id vlan-id-vlan-id</i>];</code>

You can include the statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]



NOTE: The inner-list option of the `vlan-tags` statement does not support Tag Protocol ID (TPID) values.

Binding a VLAN ID to a Logical Interface

A logical interface that you have associated (bound) to a particular VLAN ID will receive and forward incoming frames that contain a matching VLAN ID.

Binding a VLAN ID to a Single-Tag Logical Interface

To bind a VLAN ID to a single-tag logical interface, include the `vlan-id` statement:

```
vlan-id vlan-id;
```

You can include the statement at the following hierarchy levels:

- [edit interfaces *ethernet-interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]

To configure an Ethernet interface to support single-tag logical interfaces, include the `vlan-tagging` statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the `flexible-vlan-tagging` statement instead.

Binding a VLAN ID to a Dual-Tag Logical Interface

To bind a VLAN ID to a dual-tag logical interface, include the `vlan-tags` statement:

```
vlan-tags inner <tpid.>vlan-id outer <tpid.>vlan-id;
```

You can include the statement at the following hierarchy levels:

- [edit interfaces *ethernet-interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]

To configure an Ethernet interface to support dual-tag logical interfaces, include the `stacked-vlan-tagging` statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the `flexible-vlan-tagging` statement instead.

Binding a Range of VLAN IDs to a Logical Interface

A VLAN range can be used by service providers to interconnect multiple VLANs belonging to a particular customer over multiple sites. Using a VLAN ID range conserves switch resources and simplifies configuration.

Binding a Range of VLAN IDs to a Single-Tag Logical Interface

To bind a range of VLAN IDs to a single-tag logical interface, include the `vlan-id-range` statement:

```
vlan-id-range vlan-id-vlan-id;
```

You can include the statement at the following hierarchy levels:

- [edit interfaces *ethernet-interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]

To configure an Ethernet interface to support single-tag logical interfaces, include the `vlan-tagging` statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the `flexible-vlan-tagging` statement instead.

Binding a Range of VLAN IDs to a Dual-Tag Logical Interface

To bind a range of VLAN IDs to a dual-tag logical interface, include the `vlan-tags` statement. Use the `inner-list` option to specify the VLAN IDs as an inclusive range by separating the starting VLAN ID and ending VLAN ID with a hyphen.

```
vlan-tags inner-list [ vlan-id vlan-id-vlan-id ] outer <tpid.>vlan-id;
```

You can include the statement at the following hierarchy levels:

- [edit interfaces *ethernet-interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]

To configure an Ethernet interface to support dual-tag logical interfaces, include the `stacked-vlan-tagging` statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the `flexible-vlan-tagging` statement instead.

Example: Binding Ranges VLAN IDs to Logical Interfaces

The following example configures two different ranges of VLAN IDs on two different logical ports:

```
[edit interfaces]
ge-3/0/0 {
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 500-600;
  }
}
ge-3/0/1 {
  flexible-vlan-tagging;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id-range 200-300;
  }
  unit 1 {
    encapsulation vlan-bridge;
    vlan-tags outer 1000 inner-range 100-200;
  }
}
```

Binding a List of VLAN IDs to a Logical Interface

In JUNOS Release 9.5 and later, on MX Series routers you can bind a list of VLAN IDs to a single logical interface, eliminating the need to configure a separate logical interface for every VLAN or VLAN range. A logical interface that accepts packets tagged with any VLAN ID specified in a VLAN ID list is called a *VLAN-bundled* logical interface.

You can use VLAN-bundled logical interfaces to configure circuit cross-connects between Layer 2 VPN routing instances or Layer 2 circuits. Using VLAN-bundled logical interfaces simplifies configuration and reduces use of system resources such as logical interfaces, next hops, and circuits.

As an alternative to configuring multiple logical interfaces (one for each VLAN ID and one for each range of VLAN IDs), you can configure a single VLAN-bundled logical interface based on a list of VLAN IDs.

Binding a List of VLAN IDs to a Single-Tag Logical Interface

To bind a list of VLAN IDs to a single-tag logical interface, include the `vlan-id-list` statement. Specify the VLAN IDs in the list individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.

```
vlan-id-list [ vlan-id vlan-id-vlan-id ];
```

You can include the statement at the following hierarchy levels:

- [edit interfaces *ethernet-interface-name* unit *logical-unit-number*]

- [edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]

To configure an Ethernet interface to support single-tag logical interfaces, include the `vlan-tagging` statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the `flexible-vlan-tagging` statement instead.

Binding a List of VLAN IDs to a Dual-Tag Logical Interface

To bind a list of VLAN IDs to a dual-tag logical interface, include the `vlan-tags` statement. Use the `inner-list` option to specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both:

```
vlan-tags inner-list [vlan-id vlan-id-vlan-id ] outer <tpid.>vlan-id;
```



NOTE: The inner-list option of the `vlan-tags` statement does not support Tag Protocol ID (TPID) values.

You can include the statement at the following hierarchy levels:

- [edit interfaces *ethernet-interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *ethernet-interface-name* unit *logical-unit-number*]

To configure an Ethernet interface to support dual-tag logical interfaces, include the `stacked-vlan-tagging` statement at the [edit interfaces *ethernet-interface-name*] hierarchy level. To support mixed tagging, include the `flexible-vlan-tagging` statement instead.

Example: Binding Lists of VLAN IDs to Logical Interfaces

The following example configures two different lists of VLAN IDs on two different logical ports:

```
[edit interfaces]
ge-1/1/0 {
  vlan-tagging; # Only for single-tagging
  encapsulation flexible-ethernet-services;
  unit 10 {
    encapsulation vlan-ccc;
    vlan-id-list [20 30-40 45];
  }
}
ge-1/1/1 {
  flexible-vlan-tagging; # Only for mixed tagging
  encapsulation flexible-ethernet-services;
  unit 10 {
    encapsulation vlan-ccc;
    vlan-id-list [1 10 20 30-40];
  }
  unit 20 {
```



```

        encapsulation vlan-ccc;
        vlan-tags outer 200 inner-list [50–60 80 90–100];
    }
}

```

In the example configuration above, **ge-1/1/0** supports single-tag logical interfaces, and **ge-1/1/1** supports mixed tagging. The single-tag logical interfaces **ge-1/1/0.10** and **ge-1/1/1.20** each bundle lists of VLAN IDs. The dual-tag logical interface **ge-1/1/1.20** bundles lists of inner VLAN IDs.

Configuring VLAN Encapsulation

Gigabit Ethernet IQ, Gigabit Ethernet PICs with small form-factor pluggable optics (SFPs), and MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces with VLAN tagging enabled can use flexible Ethernet services, VLAN CCC, or VLAN virtual private LAN service (VPLS) encapsulation.

Aggregated Ethernet interfaces configured for VPLS can use Ethernet VPLS or VLAN VPLS.

To configure the encapsulation on a Gigabit Ethernet IQ or Gigabit Ethernet physical interface, include the **encapsulation** statement at the [edit interfaces *interface-name*] hierarchy level, specifying **flexible-ethernet-services**, **vlan-ccc**, or **vlan-vpls**:

```

[edit interfaces interface-name]
  encapsulation (flexible-ethernet-services | vlan-ccc | vlan-vpls);

```

To configure the encapsulation on an aggregated Ethernet interface, include the **encapsulation** statement at the [edit interfaces *interface-name*] hierarchy level, specifying **flexible-ethernet-services**, **ethernet-vpls**, or **vlan-vpls**:

```

[edit interfaces interface-name]
  encapsulation (flexible-ethernet-services | ethernet-vpls | vlan-vpls);

```

Ethernet interfaces in VLAN mode can have multiple logical interfaces. In CCC and VPLS modes, VLAN IDs from 1 through 511 are reserved for normal VLANs, and VLAN IDs 512 through 4094 are reserved for CCC or VPLS VLANs. For 4-port Fast Ethernet interfaces, you can use VLAN IDs 512 through 1024 for CCC or VPLS VLANs.

For encapsulation type **flexible-ethernet-services**, all VLAN IDs are valid.

In general, you configure an interface's encapsulation at the [edit interfaces *interface-name*] hierarchy level. However, for some encapsulation types, including flexible Ethernet services, Ethernet VLAN CCC and VLAN VPLS, you can also configure the encapsulation type that is used inside the VLAN circuit itself. To do this, include the **encapsulation** statement:

```

  encapsulation (vlan-ccc | vlan-tcc | vlan-vpls);

```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]

- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

You cannot configure a logical interface with VLAN CCC or VLAN VPLS encapsulation unless you also configure the physical device with the same encapsulation or with flexible Ethernet services encapsulation. In general, the logical interface must have a VLAN ID of 512 or higher; if the VLAN ID is 511 or lower, it will be subject to the normal destination filter lookups in addition to source address filtering. However if you configure flexible Ethernet services encapsulation, this VLAN ID restriction is removed.

Example: Configuring VLAN Encapsulation on a Gigabit Ethernet Interface

Configure VLAN CCC encapsulation on a Gigabit Ethernet interface:

```
interfaces ge-2/1/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 600;
  }
}
```

Example: Configuring VLAN Encapsulation on an Aggregated Ethernet Interface

Configure VLAN CCC encapsulation on an aggregated Gigabit Ethernet interface:

```
interfaces ae0 {
  vlan-tagging;
  encapsulation vlan-vpls;
  unit 0 {
    vlan-id 100;
  }
}
```

Configuring Extended VLAN Encapsulation

Gigabit Ethernet, 4-port Fast Ethernet, MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, 10-Gigabit Ethernet, and aggregated Ethernet interfaces with VLAN tagging enabled can use extended VLAN CCC or VLAN VPLS, which allow 802.1Q tagging. To configure the encapsulation on a physical interface, include the `encapsulation` statement at the [edit interfaces *interface-name*] hierarchy level, specifying `extended-vlan-ccc` or `extended-vlan-vpls`:

```
[edit interfaces interface-name]
encapsulation (extended-vlan-ccc | extended-vlan-vpls);
```

For extended VLAN CCC and extended VLAN VPLS encapsulation, all VLAN IDs 1 and higher are valid. VLAN ID 0 is reserved for tagging the priority of frames.



NOTE: For extended VLAN CCC, the VLAN IDs on ingress and egress interfaces must be the same. For back-to-back connections, all VLAN IDs must be the same.

Example: Configuring Extended VLAN Encapsulation on a Gigabit Ethernet Interface

Configure extended VLAN CCC encapsulation on Gigabit Ethernet ingress and egress interfaces:

```

interfaces ge-0/0/0 {
  vlan-tagging;
  encapsulation extended-vlan-ccc;
  unit 0 {
    vlan-id 2;
    family ccc;
  }
}
interfaces ge-1/0/0 {
  vlan-tagging;
  encapsulation extended-vlan-ccc;
  unit 0 {
    vlan-id 2;
    family ccc;
  }
}

```

Example: Configuring Extended VLAN Encapsulation on an Aggregated Ethernet Interface

Configure extended VLAN VPLS encapsulation on an aggregated Ethernet interface:

```

interfaces ae0 {
  vlan-tagging;
  encapsulation extended-vlan-vpls;
  unit 0 {
    vlan-id 100;
  }
}

```

Guidelines for Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs

For MX Series routers, you can bind a list of VLAN IDs to a logical interface, configure a Layer 2 VPN routing instance or Layer 2 circuit on the logical interface, and then use the logical interface to configure a circuit cross-connect (CCC) to another Layer 2 VPN routing instance or Layer 2 circuit.

A CCC allows you to configure transparent connections between two circuits so that packets from the source circuit are delivered to the destination circuit with, at most, the Layer 2 address being changed. You configure a CCC by connecting circuit interfaces of the same type. For more information, see “Configuring Circuit and Translational Cross-Connects” on page 223.



NOTE: The JUNOS Software supports binding of Ethernet logical interfaces to lists of VLAN IDs on MX Series routers only. For all other routers, you can bind an Ethernet logical interface to only a single VLAN ID or to a single range of VLAN IDs.

The following configuration guidelines apply to bundling lists of VLAN IDs to Ethernet logical interfaces used to configure CCCs:

- Guidelines for Configuring Physical Link-Layer Encapsulation to Support CCCs on page 612
- Guidelines for Configuring Logical Link-Layer Encapsulation to Support CCCs on page 612

Guidelines for Configuring Physical Link-Layer Encapsulation to Support CCCs

To enable a physical interface to support VLAN-bundled logical interfaces that you will use to configure a CCC, you must configure one of the following physical link-layer encapsulation types:

- `extended-vlan-ccc`—For Ethernet interfaces with standard TPID tagging.
- `flexible-ethernet-services`—For supported Gigabit Ethernet interfaces for which you want to configure multiple per-unit Ethernet encapsulations.

To specify the physical link-layer encapsulation type, include the **encapsulation (Physical Interface)** statement. For information about configuring the interface encapsulation on a physical interface, see “Configuring Interface Encapsulation on Physical Interfaces” on page 106.

Guidelines for Configuring Logical Link-Layer Encapsulation to Support CCCs

For VLAN-bundled logical interfaces that you use to configure a CCC, specific logical link-layer encapsulation types are used inside the circuits themselves.

Table 50 on page 613 describes the logical link-layer encapsulation types used within circuits connected using VLAN-bundled logical interfaces of the same type.

Table 50: Encapsulation Inside Circuits CCC-Connected by VLAN-Bundled Logical Interfaces

Encapsulation Inside the Circuit	Layer 2 Circuit Joined by Configuring an Interface-to-Interface CCC Connection	
	Layer 2 VPN Routing Instance	Layer 2 Circuit
Syntax	encapsulation-type (ethernet ethernet-vlan);	encapsulation vlan-ccc;
Hierarchy Level	[edit routing-instances <i>routing-instance-name</i> protocols l2vpn], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> protocols l2vpn]	[edit interfaces <i>ethernet-interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>ethernet-interface-name</i> unit <i>logical-unit-number</i>]
Usage Guidelines	See the <i>JUNOS VPNs Configuration Guide</i> .	See “Configuring Interface Encapsulation on Logical Interfaces” on page 160, “Configuring Circuit and Translational Cross-Connects” on page 223, and “Defining the Encapsulation for Switching Cross-Connects” on page 225.
For a Single-Tag Logical Interface	The MX Series router automatically uses ethernet as the Layer 2 protocol used to encapsulate incoming traffic. Although the connection spans multiple VLANs, the VLANs are bundled and therefore can be encapsulated as a single VLAN. NOTE: With ethernet encapsulation, the circuit signal processing does not check that the VLAN ID list is the same at both ends of the CCC connection.	Configure the MX Series router to use vlan-ccc as the logical link-layer encapsulation type.
For a Dual-Tag Logical Interface	Configure the MX Series router to use ethernet-vlan as the Layer 2 protocol to encapsulate incoming traffic. With ethernet-vlan encapsulation, circuit signal processing checks that the VLAN ID list is the same at both ends of the CCC connection. If a VLAN ID list mismatch is detected, you can view the error condition in the show interfaces command output.	The MX Series router automatically uses vlan-ccc as the logical link-layer encapsulation type, regardless of the value configured.
Related Topics	<ul style="list-style-type: none"> ■ Binding VLAN IDs to Logical Interfaces on page 604 ■ Defining the Encapsulation for Switching Cross-Connects on page 225 	

Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface

This topic describes how to configure a Layer 2 VPN routing instance on a logical interface bound to a list of VLAN IDs.

The topic consists of the following tasks:

- Configuring a VLAN-Bundled Logical Interface on page 614
- Specifying the Interface Over Which VPN Traffic Travels to the CE Router on page 614
- Specifying the Interface to Handle Traffic for a CCC on page 615

Configuring a VLAN-Bundled Logical Interface

To configure a VLAN-bundled logical interface, specify the list of VLAN IDs by including the `vlan-id-list` statement or the `vlan-tags` statement on a provider edge (PE) router:

```
interfaces {
  ethernet-interface-name {
    vlan-tagging; # Support single- or dual-tag logical interfaces
    flexible-vlan-tagging; # Support mixed tagging
    encapsulation (extended-vlan-ccc | flexible-ethernet-services);
    unit logical-unit-number {
      vlan-id-list [vlan-id vlan-id-vlan-id]; # For single-tag
      vlan-tags outer <tpid.>vlan-id inner-list [vlan-id vlan-id-vlan-id]; # For dual-tag
    }
    ...
  }
}
```

You can include the statements at the following hierarchy levels:

- [edit]
- [edit logical-systems *logical-system-name*]

Specifying the Interface Over Which VPN Traffic Travels to the CE Router

To configure a Layer 2 VPN routing instance on a PE router, include the `instance-type` statement and specify the value `l2vpn`. To specify an interface connected to the router, include the `interface` statement and specify the VLAN-bundled logical interface:

```
instance-type l2vpn;
interface logical-interface-name;
```

You can include the statements at the following hierarchy levels:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

Specifying the Interface to Handle Traffic for a CCC

To configure the VLAN-bundled logical interface as the interface to handle traffic for a circuit connected to the Layer 2 VPN routing instance, include the following statements:

```
protocols {
  l2vpn {
    (control-word | no-control-word);
    encapsulation-type (ethernet | ethernet-vlan);
    site site-name {
      site-identifier identifier;
      interface logical-interface-name { # VLAN-bundled logical interface
        . . . interface-options . . .
      }
    }
  }
}
```

You can include the statements at the same hierarchy level at which you include the `instance-type l2vpn` and `interface logical-interface-name` statements:

- [edit routing-instances *routing-instance-name*]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*]

To enable a Layer 2 VPN routing instance on a PE router, include the `l2vpn` statement. For more information, see the *JUNOS VPNs Configuration Guide*.

The `encapsulation-type` statement specifies the Layer 2 protocol used for traffic from the customer edge (CE) router. If the Layer 2 VPN routing instance is being connected to a single-tag Layer 2 circuit, specify `ethernet` as the encapsulation type. If the Layer 2 VPN routing instance is being connected to a dual-tag Layer 2 circuit, specify `ethernet-vlan` as the encapsulation type.

To specify the interface to handle traffic for a circuit connected to the Layer 2 VPN routing instance, include the `interface` statement and specify the VLAN-bundled logical interface.

Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface

This topic describes how to configure a Layer 2 circuit on a logical interface bound to a list of VLAN IDs.

The topic consists of the following tasks:

- Configuring a VLAN-Bundled Logical Interface on page 616
- Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit on page 616

Configuring a VLAN-Bundled Logical Interface

To configure a VLAN-bundled logical interface, specify the list of VLAN IDs by including the `vlan-id-list` statement or the `vlan-tags` statement:

```
interfaces {
  ethernet-interface-name {
    vlan-tagging; # Support single- or dual-tag logical interfaces
    flexible-vlan-tagging; # Support mixed tagging
    encapsulation (extended-vlan-ccc | flexible-ethernet-services);
    unit logical-unit-number {
      encapsulation vlan-ccc; # Required for single-tag
      vlan-id-list [vlan-id vlan-id-vlan-id]; # For single-tag
      vlan-tags outer <tpid.>vlan-id inner-list [vlan-id vlan-id-vlan-id]; # For dual-tag
    }
    ...
  }
}
```

You can include the statements at the following hierarchy levels:

- [edit]
- [edit logical-systems *logical-system-name*]

For a single-tag logical interface, include the `encapsulation` statement and specify `vlan-ccc` so that CCC circuit encapsulation is used inside the Layer 2 circuit.



NOTE: In the case of a dual-tag logical interface, the JUNOS Software automatically uses the `vlan-ccc` encapsulation type.

Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit

To configure the VLAN-bundled logical interface as the interface to handle traffic for a circuit connected to the Layer 2 circuit, include the following statements:

```
l2circuit {
  neighbor address {
    interface logical-interface-name {
      virtual-circuit-id number;
      no-control-word;
    }
  }
}
```

You can include the statements at the following hierarchy levels:

- [edit protocols]
- [edit logical-systems *logical-system-name* protocols]

To enable a Layer 2 circuit, include the `l2circuit` statement.

To configure the router as a neighbor for a Layer 2 circuit, specify the neighbor address using the `neighbor` statement.

To specify the interface to handle traffic for a circuit connected to the Layer 2 circuit, include the `interface` statement and specify the VLAN-bundled logical interface.

Example: Configuring a Layer 2 VPN Routing Instance on a VLAN-Bundled Logical Interface

The following configuration shows that the single-tag logical interface `ge-1/0/5.0` bundles a list of VLAN IDs, and the logical interface `ge-1/1/1.0` supports IPv4 traffic using IP address 10.30.1.130 and can participate in an MPLS path.

```
[edit interfaces]
ge-1/0/5 {
  vlan-tagging;
  encapsulation extended-vlan-ccc;
  unit 0 { # VLAN-bundled logical interface
    vlan-id-list [513 516 520-525];
  }
}
ge-1/1/1 {
  unit 0 {
    family inet {
      address 10.30.1.1/30;
    }
    family mpls;
  }
}
```

The following configuration shows the type of traffic supported on the Layer 2 VPN routing instance:

```
[edit protocols]
rsvp {
  interface all;
  interface lo0.0;
}
mpls {
  label-switched-path lsp {
    to 10.255.69.128;
  }
  interface all;
}
bgp {
  group g1 {
    type internal;
    local-address 10.255.69.96;
    family l2vpn {
      signaling;
    }
    neighbor 10.255.69.128;
  }
}
ospf {
```

```

traffic-engineering;
area 0.0.0.0 {
    interface lo0.0;
    interface ge-1/1/1.0;
}
}

```

The following configuration shows that the VLAN-bundled logical interface is the interface over which VPN traffic travels to the CE router and handles traffic for a CCC to which the VPN connects.

```

[edit routing-instances]
red {
    instance-type l2vpn;
    interface ge-1/0/5.0; # VLAN-bundled logical interface
    route-distinguisher 10.255.69.96:100;
    vrf-target target:1:1;
    protocols {
        l2vpn {
            encapsulation-type ethernet; # For single-tag VLAN logical interface
            site CE_ultima {
                site-identifier 1;
                interface ge-1/0/5.0;
            }
        }
    }
}
}

```



NOTE: Because the VLAN-bundled logical interface supports single-tag frames, Ethernet is the Layer 2 protocol used to encapsulate incoming traffic. Although the connection spans multiple VLANs, the VLANs are bundled and therefore can be encapsulated as a single VLAN.

However, with Ethernet encapsulation, the circuit signal processing does not check that the VLAN ID list is the same at both ends of the CCC connection.

Example: Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface

The following configuration shows that the single-tag logical interface `ge-1/0/5.0` bundles a list of VLAN IDs, and the logical interface `ge-1/1/1.0` supports IPv4 traffic using IP address 10.30.1.1/30 and can participate in an MPLS path.

```

[edit interfaces]
ge-1/0/5 {
    vlan-tagging;
    encapsulation extended-vlan-ccc;
    unit 0 { # VLAN-bundled logical interface
        vlan-id-list [513 516 520-525];
    }
}
ge-1/1/1 {
    unit 0 {

```

```

        family inet {
            address 10.30.1.1/30;
        }
        family mpls;
    }
}

```

The following configuration shows the type of traffic supported on the Layer 2 VPN routing instance, and shows that the VLAN-bundled logical interface handles traffic for a CCC to which the Layer 2 circuit connects:

```

[edit protocols]
rsvp {
    interface all;
    interface lo0.0;
}
mpls {
    label-switched-path lsp {
        to 10.255.69.128;
    }
    interface all;
}
ospf {
    traffic-engineering;
    area 0.0.0.0 {
        interface lo0.0;
        interface ge-1/1/1.0;
    }
}
ldp {
    interface ge-1/1/1.0;
    interface ge-1/0/5.0; # VLAN-bundled logical interface
    interface lo0.0;
}
l2circuit {
    neighbor 10.255.69.128 {
        interface ge-1/0/5.0 { # VLAN-bundled logical interface
            virtual-circuit-id 3;
            no-control-word;
        }
    }
}
}

```

Configuring a Logical Interface for Access Mode

Enterprise network administrators can configure a single logical interface to accept untagged packets and forward the packets within a specified bridge domain. A logical interface configured to accept untagged packets is called an *access interface* or *access port*. Access interface configuration is supported on MX Series routers only.

```
interface-mode access;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family bridge]

- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge]

When an untagged or tagged packet is received on an access interface, the packet is accepted, the VLAN ID is added to the packet, and the packet is forwarded within the bridge domain that is configured with the matching VLAN ID.

Example: Configuring a Logical Interface for Access Mode

The following example configures a logical interface as an access port with a VLAN ID of 20:

```
[edit interfaces ge-1/2/0]
unit 1 {
  family bridge {
    interface-mode access;
    vlan-id 20;
  }
}
```

Configuring a Logical Interface for Trunk Mode

As an alternative to configuring a logical interface for each VLAN, enterprise network administrators can configure a single logical interface to accept untagged packets or packets tagged with any VLAN ID specified in a list of VLAN IDs. Using a VLAN ID list conserves switch resources and simplifies configuration. A logical interface configured to accept packets tagged with any VLAN ID specified in a list is called a *trunk interface* or *trunk port*. Trunk interface configuration is supported on MX Series routers only. Trunk interfaces support integrated routing and bridging (IRB).

To configure a logical interface to accept any packet tagged with a VLAN ID that matches the list of VLAN IDs, include the **interface-mode** statement and specify the **trunk** option:

```
interface-mode trunk;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family bridge]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge]

Configuring the VLAN ID List for a Trunk Interface

To configure the list of VLAN IDs to be accepted by the trunk port, include the **vlan-id-list** statement and specify the list of VLAN IDs. You can specify individual VLAN IDs with a space separating the ID numbers, specify a range of VLAN IDs with a dash separating the ID numbers, or specify a combination of individual VLAN IDs and a range of VLAN IDs.

```
vlan-id-list (Interface in Bridge Domain) [number number-number];
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family bridge interface-mode trunk]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family bridge interface-mode trunk]

When a packet is received that is tagged with a VLAN ID specified in the trunk interface list of VLAN IDs, the packet is accepted and forwarded within the bridge domain that is configured with the matching VLAN ID.

When a packet is received that is tagged with a VLAN ID not specified in the trunk interface list of VLAN IDs, the native VLAN ID is pushed in front of the existing VLAN tag or tags and the packet is forwarded within the bridge domain that is configured with the matching VLAN ID.

When an untagged packet is received on a trunk interface, the native VLAN ID is added to the packet and the packet is forwarded within the bridge domain that is configured with the matching VLAN ID.

A bridge domain configured with a matching VLAN ID must be configured before the trunk interface is configured. To learn more about configuring bridge domains, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring a Trunk Interface on a Bridge Network

On MX Series routers, you can configure a trunk interface on a bridge network.

The following output sample shows trunk port configuration on a bridge network:

```
user@host# run show interfaces
ge-0/0/0 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id 1;
  }
}
ge-2/0/0 {
  unit 0 {
    family bridge {
      interface-mode trunk;
      vlan-id-list 1-200;
    }
  }
}
ge-2/0/1 {
  flexible-vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 0 {
    encapsulation vlan-bridge;
    vlan-id 1;
  }
}
```

```
}  
}
```

Chapter 35

Configuring Aggregated Ethernet Interfaces

This chapter contains the following sections:

- Aggregated Ethernet Interfaces Overview on page 623
- Configuring Aggregated Ethernet Interfaces on page 624
- Configuring Ethernet Link Aggregation on page 625
- Configuring Aggregated Ethernet Link Protection on page 626
- Setting the Number of Aggregated Ethernet Interfaces on the Chassis on page 627
- Configuring Aggregated Ethernet LACP on page 627
- Configuring Tagged Aggregated Ethernet Interfaces on page 632
- Configuring Untagged Aggregated Ethernet Interfaces on page 633
- Configuring Aggregated Ethernet Link Speed on page 634
- Configuring Aggregated Ethernet Minimum Links on page 635
- Configuring Hierarchical Scheduler on Aggregated Ethernet Interfaces on page 635
- Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers on page 636

Aggregated Ethernet Interfaces Overview

Link aggregation of Ethernet interfaces is defined in the IEEE 802.3ad standard. The JUNOS implementation of 802.3ad balances traffic across the member links within an aggregated Ethernet bundle based on the Layer 3 information carried in the packet. This implementation uses the same load-balancing algorithm used for per-packet load balancing. For information about per-packet load balancing, see the *JUNOS Routing Protocols Configuration Guide*.



NOTE: For information about configuring circuit cross-connects over aggregated Ethernet, see “Examples: Configuring Switching Cross-Connects” on page 233.

Configuring Aggregated Ethernet Interfaces

You configure an aggregated Ethernet virtual link by specifying the link number as a physical device and then associating a set of ports that have the same speed and are in full-duplex mode. The physical interfaces can be Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ, Gigabit Ethernet IQ2 and IQ2-E, or 10-Gigabit Ethernet IQ2 and IQ2-E interfaces. Generally, you cannot use a combination of these interfaces within the same aggregated link; however, you can combine Gigabit Ethernet and Gigabit Ethernet IQ interfaces in a single aggregated Ethernet bundle.

The following routers support a maximum of 16 physical interfaces per single aggregated Ethernet bundle:

- M120
- M320
- MX960
- MX480
- MX240
- All T Series routers

All other routers support a maximum of 8 physical interfaces per aggregated Ethernet bundle.

Aggregated Ethernet interfaces can use interfaces from different FPCs, DPCs, or PICs.

Simple filters are not supported for interfaces in aggregated Ethernet bundles:

- in M Series routers, simple filters are supported in Gigabit Ethernet Enhanced Intelligent Queuing interfaces only, except when the interface is part of an aggregated Ethernet bundle.
- in MX Series routers, simple filters are supported in Enhanced Queuing Dense Port Concentrator (EQ DPC) interfaces only, except when the interface is part of an aggregated Ethernet bundle.

On the aggregated bundle, no IQ-specific capabilities such as MAC accounting, VLAN rewrites, and VLAN queuing are available. For more information about IQ-specific capabilities, see “Configuring Gigabit Ethernet Accounting and Policing” on page 755.

Aggregated Ethernet interfaces can be either tagged or untagged, with LACP enabled or disabled. Aggregated Ethernet interfaces on Ethernet IQ2 and IQ2-E PICs in MX Series routers support the configuration of **flexible-vlan-tagging**, **native-vlan-id**, and rewrite operations on dual-tagged frames, which consist of the following configuration statements:

- `inner-tag-protocol-id`
- `inner-vlan-id`
- `pop-pop`

- pop-swap
- push-push
- swap-push
- swap-swap

In all cases, you must set the number of aggregated Ethernet interfaces on the chassis. You can also set the link speed and the minimum links in a bundle.

Configuring Ethernet Link Aggregation

On Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces on M Series and T Series routers, you can associate a physical interface with an aggregated Ethernet interface. To enable the aggregated link, include the `802.3ad` statement at the `[edit interfaces interface-name fastether-options]` or `[edit interfaces interface-name gigether-options]` hierarchy level:

```
[edit interfaces interface-name (fastether-options | gigether-options)]
802.3ad aex;
```

You specify the interface instance number `x` to complete the link association; `x` can be from 0 through 127, for a total of 128 aggregated interfaces. You must also include a statement defining `aex` at the `[edit interfaces]` hierarchy level. You can optionally specify other physical properties that apply specifically to the aggregated Ethernet interfaces; for details, see “Configuring Ethernet Interfaces” on page 585, and for a sample configuration, see “Example: Configuring Aggregated Ethernet Interfaces” on page 814.



NOTE: In general, aggregated Ethernet bundles support the features available on all supported interfaces that can become a member link within the bundle. As an exception, Gigabit Ethernet IQ features and some newer Gigabit Ethernet features are not supported in aggregated Ethernet bundles.

Gigabit Ethernet IQ and SFP interfaces can be member links, but IQ- and SFP-specific features are not supported on the aggregated Ethernet bundle even if all the member links individually support those features.

To delete an aggregated Ethernet interface from the configuration, issue the `delete interfaces aex` command at the `[edit]` hierarchy level in configuration mode:

```
[edit]
user@host# delete interfaces aex
```

If you delete an aggregated Ethernet interface from the configuration, the JUNOS Software removes the configuration statements related to `aex` and sets this interface to down state. However, the aggregated Ethernet interface is not deleted until you delete the `chassis aggregated-devices ethernet device-count` configuration statement.

Configuring Aggregated Ethernet Link Protection

To support JUNOS QoS features, aggregated Ethernet interfaces support link protection. On aggregated Ethernet interfaces, you designate a primary and backup link. Egress traffic passes only through the designated primary link. This includes transit traffic and locally generated traffic on the router. When the primary link fails, traffic is routed through the backup link. Because some traffic loss is unavoidable, egress traffic is not automatically routed back to the primary link when the primary link is reestablished. Instead, you manually control when traffic should be diverted back to the primary link from the designated backup link by issuing the **request interface revert aex** operational command.

To enable link protection on aggregated Ethernet interfaces, include the **link-protection** statement at the **[edit interfaces aex aggregated-ether-options]** hierarchy level:

```
[edit interfaces]
aex {
  aggregated-ether-options {
    link-protection;
  }
}
```

You also must specify a primary and a secondary, or backup, link. To configure a primary and a backup link, include the primary and backup statement at the **[edit interfaces ge-fpc/pic/port gige-ether-options 802.3ad aex]** hierarchy level or the **[edit interfaces fe-fpc/pic/port fastether-options 802.3ad aex]** hierarchy level:

```
[edit interfaces]
interface-name {
  (gige-ether-options | fastether-options) {
    802.3ad aex (primary | backup);
  }
}
```

To revert back to sending traffic to the primary designated link when traffic is passing through the designated backup link, issue the **request interface revert aex** operational mode command:

```
user@host# request interface revert aex
```

To disable link protection, issue the **delete interfaces aex aggregated-ether-options link-protection** command in configuration mode and commit the configuration:

```
[edit interfaces aex]
user@host# delete interfaces aex aggregated-ether-options link-protection
```

For a sample configuration, see “Example: Configuring Aggregated Ethernet Link Protection” on page 815.

Setting the Number of Aggregated Ethernet Interfaces on the Chassis

By default, no aggregated Ethernet interfaces are created. You must define the number of aggregated Ethernet interfaces by including the `device-count` statement at the `[edit chassis aggregated-devices ethernet]` hierarchy level:

```
[edit chassis]
aggregated-devices {
  ethernet {
    device-count number;
  }
}
[edit interfaces]
aex{
  aggregated-ether-options {
    link-protection;
  }
  unit logical-unit number {
    family inet {
      address address;
    }
  }
}
```

The maximum number of aggregated devices you can configure is 128. The aggregated interfaces are numbered from `ae0` through `ae127`. For information about configuring aggregated devices, see the *JUNOS System Basics Configuration Guide*.

You must also specify the constituent physical links by including the `802.3ad` statement at the `[edit interfaces interface-name fastether-options]` or `[edit interfaces interface-name gigether-options]` hierarchy level; for more information, see “Configuring Ethernet Link Aggregation” on page 625. For a sample configuration, see “Example: Configuring Aggregated Ethernet Interfaces” on page 814.

Configuring Aggregated Ethernet LACP

For aggregated Ethernet interfaces, you can configure the Link Aggregation Control Protocol (LACP). LACP is one method of bundling several physical interfaces to form one logical interface. You can configure both VLAN-tagged and untagged aggregated Ethernet with or without LACP enabled.

LACP exchanges are made between actors and partners. An actor is the local interface in an LACP exchange. A partner is the remote interface in an LACP exchange.

LACP is defined in IEEE 802.3ad, *Aggregation of Multiple Link Segments*.

LACP was designed to achieve the following:

- Automatic addition and deletion of individual links to the aggregate bundle without user intervention
- Link monitoring to check whether both ends of the bundle are connected to the correct group

The JUNOS implementation of LACP provides link monitoring but not automatic addition and deletion of links.

The LACP mode can be active or passive. If the actor and partner are both in passive mode, they do not exchange LACP packets, which results in the aggregated Ethernet links not coming up. If either the actor or partner is active, they do exchange LACP packets. By default, LACP is in passive mode on aggregated Ethernet interfaces. To initiate transmission of LACP packets and response to LACP packets, you must enable LACP active mode.

To enable LACP active mode, include the `lACP` statement at the `[edit interfaces interface-name aggregated-ether-options]` hierarchy level, and specify the `active` option:

```
[edit interfaces interface-name aggregated-ether-options]
lACP {
  active;
}
```

To restore the default behavior, include the `lACP` statement at the `[edit interfaces interface-name aggregated-ether-options]` hierarchy level, and specify the `passive` option:

```
[edit interfaces interface-name aggregated-ether-options]
lACP {
  passive;
}
```

For more information, see the following sections:

- Configuring the LACP Interval on page 628
- Configuring LACP Link Protection on page 629
- Tracing LACP Operations on page 631
- Example: Configuring Aggregated Ethernet LACP on page 631

Configuring the LACP Interval

By default, the actor and partner send LACP packets every second. You can configure the interval at which the interfaces send LACP packets by including the `periodic` statement at the `[edit interfaces interface-name aggregated-ether-options lACP]` hierarchy level:

```
[edit interfaces interface-name aggregated-ether-options lACP]
periodic interval;
```

The interval can be fast (every second) or slow (every 30 seconds). You can configure different periodic rates on active and passive interfaces. When you configure the active and passive interfaces at different rates, the transmitter honors the receiver's rate.



NOTE: Source address filtering does not work when LACP is enabled. For more information about source address filtering, see “Enabling Ethernet MAC Address Filtering” on page 591.

Percentage policers are not supported on aggregated Ethernet interfaces with the CCC protocol family configured. For more information about percentage policers, see the *JUNOS Policy Framework Configuration Guide*.

Generally, LACP is supported on all untagged aggregated Ethernet interfaces. For more information, see “Configuring Untagged Aggregated Ethernet Interfaces” on page 633.

For M Series Multiservice Edge Routers with enhanced Flexible PIC Concentrators (FPCs) and T Series routers, LACP over VLAN-tagged aggregated Ethernet interfaces is supported. For 8-port, 12-port, and 48-port Fast Ethernet PICs, LACP over VLAN-tagged interfaces is not supported.

Configuring LACP Link Protection



NOTE: When using LACP link protection, you can configure only two member links to an aggregated Ethernet interface: one active and one standby.

To force active and standby links within an aggregated Ethernet, you can configure LACP link protection and system priority at the aggregated Ethernet interface level using the `link-protection` and `system-priority` statements. Configuring values at this level results in only the configured interfaces using the defined configuration. LACP interface configuration also enables you to override global (chassis) LACP settings.

LACP link protection also uses port priority. You can configure port priority at the Ethernet interface `[gigether-options]` hierarchy level using the `port-priority` statement. If you choose not to configure port priority, LACP link protection uses the default value for port priority (127).



NOTE: LACP link protection supports per-unit scheduling configuration on aggregated Ethernet interfaces.

Enabling LACP Link Protection

To enable LACP link protection for an aggregated Ethernet interfaces, use the `link-protection` statement at the `[edit interfaces aeX aggregated-ether-options lacp]` hierarchy level:

```
[edit interfaces aeX aggregated-ether-options lacp]
link-protection;
  disable (Link Protection);
  revertive;
```

```

    non-revertive;
}

```

By default, LACP link protection reverts to a higher-priority (lower-numbered) link when that higher-priority link becomes operational or a link is added to the aggregator that is determined to be higher in priority. However, you can suppress link calculation by adding the **non-revertive** statement to the LACP link protection configuration. In nonrevertive mode, once a link is active and collecting and distributing packets, the subsequent addition of a higher-priority (better) link does not result in a switch and the current link remains active.

If LACP link protection is configured to be nonrevertive at the global ([edit chassis] hierarchy) level, you can add the **revertive** statement to the LACP link protection configuration to override the nonrevertive setting for the interface. In revertive mode, the addition of a higher-priority link to the aggregator results in LACP performing a priority recalculation and switching from the current active link to the new active link.



CAUTION: If both ends of an aggregator have LACP link protection enabled, make sure to configure both ends of the aggregator to use the same mode. Mismatching LACP link protection modes can result in lost traffic.

Configuring LACP System Priority

To configure LACP system priority for aggregated Ethernet interfaces on the interface, use the **system-priority** statement at the [edit interfaces aeX aggregated-ether-options lacp] hierarchy level:

```

[edit interfaces aeX aggregated-ether-options lacp]
system-priority;

```

The system priority is a 2-octet binary value that is part of the LACP system ID. The LACP system ID consists of the system priority as the two most-significant octets and the interface MAC address as the six least-significant octets. The system with the numerically lower value for system priority has the higher priority. By default, system priority is 127, with a range of 0 to 65,535.

Configuring LACP Port Priority

To configure LACP port priority for aggregated Ethernet interfaces, use the **port-priority** statement at the [edit interfaces *interface-name* gigether-options 802.3ad aeX lacp] or [edit interfaces *interface-name* fastether-options 802.3ad aeX lacp] hierarchy levels:

```

[edit interfaces interface-name gigether-options 802.3ad aeX lacp]
port-priority priority;

```

The port priority is a 2-octet field that is part of the LACP port ID. The LACP port ID consists of the port priority as the two most-significant octets and the port number as the two least-significant octets. The system with the numerically lower value for port priority has the higher priority. By default, port priority is 127, with a range of 0 to 65,535.

Port aggregation selection is made by each system based on the highest port priority and are assigned by the system with the highest priority. Ports are selected and assigned starting with the highest priority port of the highest priority system and working down in priority from there.

Tracing LACP Operations

To trace the operations of the LACP process, include the `traceoptions` statement at the `[edit protocols lacp]` hierarchy level:

```
[edit protocols lacp]
traceoptions {
  file <filename> <files number> <size size> <world-readable | no-world-readable>;
  flag <flag>;
  no-remote-trace;
}
```

You can specify the following flags in the `protocols lacp traceoptions` statement:

- `all`—All LACP tracing operations
- `configuration`—Configuration code
- `packet`—Packets sent and received
- `process`—LACP process events
- `protocol`—LACP protocol state machine
- `routing-socket`—Routing socket events
- `startup`—Process startup events

For general information about tracing, see the tracing and logging information in the *JUNOS System Basics Configuration Guide*.

Example: Configuring Aggregated Ethernet LACP

Configure aggregated Ethernet LACP over a VLAN-tagged interface:

LACP with VLAN-Tagged Aggregated Ethernet

```
[edit interfaces]
fe-5/0/1 {
  fastether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    lacp {
      active;
    }
  }
  vlan-tagging;
  unit 0 {
```

```

        vlan-id 100;
        family inet {
            address 10.1.1.2/24 {
                vrrp-group 0 {
                    virtual-address 10.1.1.4;
                    priority 200;
                }
            }
        }
    }
}

```

Configure aggregated Ethernet LACP over an untagged interface:

LACP with Untagged Aggregated Ethernet

```

[edit interfaces]
fe-5/0/1 {
    fastether-options {
        802.3ad ae0;
    }
}
ae0 {
    aggregated-ether-options {
        lacp {
            active;
        }
    }
    unit 0 {
        family inet {
            address 10.1.1.2/24 {
                vrrp-group 0 {
                    virtual-address 10.1.1.4;
                    priority 200;
                }
            }
        }
    }
}
}

```

Configuring Tagged Aggregated Ethernet Interfaces

To specify aggregated Ethernet interfaces, include the `vlan-tagging` statement at the `[edit interfaces aeX]` hierarchy level:

```

[edit interfaces aex]
vlan-tagging;

```

You must also include the `vlan-id` statement:

```

vlan-id number;

```

You can include this statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number]`

- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

For more information about the `vlan-tagging` and `vlan-id` statements, see “Configuring 802.1Q VLANs” on page 599.

Configuring Untagged Aggregated Ethernet Interfaces

When you configure an untagged Aggregated Ethernet interface, the existing rules for untagged interfaces apply. These rules are as follows:

- You can configure only one logical interface (unit 0) on the port. The logical unit 0 is used to send and receive LACP or marker protocol data units (PDUs) to and from the individual links.
- You cannot include the `vlan-id` statement in the configuration of the logical interface.

Table 51 on page 633 lists untagged aggregated Ethernet and LACP support by PIC and router.

Table 51: Untagged Aggregated Ethernet and LACP Support by PIC and Platform

PIC Type	M Series	LACP	T Series	LACP
4-port Fast Ethernet PIC Type 1	Yes	Yes	Yes	Yes
1-port Gigabit Ethernet PIC Type 1	Yes	Yes	Yes	Yes
2-port Gigabit Ethernet PIC Type 2	Yes	Yes	Yes	Yes
4-port Gigabit Ethernet PIC Type 2	Yes	Yes	Yes	Yes
1-port 10-Gigabit Ethernet M160	Yes	Yes	NA	NA
10-port Gigabit Ethernet PIC Type 3	Yes (M120, M320)	Yes	Yes	Yes
1-port 10-Gigabit Ethernet PIC Type 3	N/A	NA	Yes	Yes
8-port Gigabit Ethernet PIC Type 3	Yes	Yes	Yes	Yes

The 8-port Fast Ethernet PIC does not support untagged aggregated Ethernet or LACP.

Syslog messages are logged if you try to configure an untagged aggregated Ethernet interface using an unsupported PIC type.

For more information about configuring LACP, see “Configuring Aggregated Ethernet LACP” on page 627.

Example: Configuring Untagged Aggregated Ethernet Interfaces

Configure an untagged aggregated Ethernet interface by omitting the `vlan-tagging` and `vlan-id` statements from the configuration:

```
[edit interfaces]
fe-5/0/1 {
  fastether-options {
    802.3ad ae0;
  }
}
ae0 {
  unit 0 {
    family inet {
      address 13.1.1.2/24 {
        vrrp-group 0 {
          virtual-address 13.1.1.4;
          priority 200;
        }
      }
    }
  }
}
```

Configuring Aggregated Ethernet Link Speed

On aggregated Ethernet interfaces, you can set the required link speed for all interfaces included in the bundle. All interfaces that make up a bundle must be the same speed. If you include in the aggregated Ethernet interface an individual link that has a speed different from the speed you specify in the `link-speed` parameter, an error message will be logged. To set the required link speed, include the `link-speed` statement at the `[edit interfaces interface-name aggregated-ether-options]` hierarchy level:

```
[edit interfaces interface-name aggregated-ether-options]
link-speed speed ;
```

speed can be in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000).

Aggregated Ethernet interfaces on the M120 router can have one of the following speed values:

- 100m—Links are 100 Mbps.
- 10g—Links are 10 Gbps.
- 1g—Links are 1 Gbps.
- OC192—Links are OC192 or STM64c.

Configuring Aggregated Ethernet Minimum Links

On aggregated Ethernet interfaces, you can configure the minimum number of links that must be up for the bundle as a whole to be labeled up. By default, only one link must be up for the bundle to be labeled up.

To configure the minimum number of links, include the `minimum-links` statement at the `[edit interfaces interface-name aggregated-ether-options]` hierarchy level:

```
[edit interfaces interface-name aggregated-ether-options]
  minimum-links number;
```

On M120, M320, MX Series, T Series, and TX Matrix routers with Ethernet interfaces, the valid range for `minimum-links number` is 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

On all other routers, the range of valid values for `minimum-links number` is 1 through 8. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

If the number of links configured in an aggregated Ethernet interface is less than the minimum link value configured under the `aggregated-ether-options` statement, the configuration commit fails and an error message is displayed.

Configuring Hierarchical Scheduler on Aggregated Ethernet Interfaces

On aggregated Ethernet interfaces, you can configure the hierarchical scheduler in non link-protect mode. The M120, MX Series, and T Series routers with aggregated Ethernet IQ2 PICs in non link-protect mode now support the following scheduler functions:

- Per unit scheduler
- Hierarchical scheduler
- Shaping at the physical interface

To configure the hierarchical scheduler on aggregated Ethernet interfaces in the non link-protect mode, include the `hierarchical-scheduler` statement at the `[edit interfaces aeX]` hierarchy level:

```
[edit interfaces aeX hierarchical-scheduler]
```

Prior to JUNOS Release 9.6, the hierarchical scheduler mode required the `aggregated-ether-options` statement `link-protection` option, otherwise a configuration error occurs.

To specify the member link bandwidth derivation based on the equal division model (`scale`) or the replication model (`replicate`) on aggregated Ethernet interfaces, include the `member-link-scheduler (scale | replicate)` option at the `[edit class-of-service interfaces aeX]` hierarchy level. The default setting is `scale`.

```
[edit class-of-service interfaces aeX member-link-scheduler (scale | replicate)]
```



NOTE: In link-protect mode, only one link is active at a time and the other link acts as the backup link, whereas in a non link-protect mode, all the links of the aggregate bundle are active at the same time. There is no backup link. If a link goes down or a new link is added to the bundle, traffic redistribution occurs.

For more information on the hierarchical scheduler (CoS), see the *JUNOS Class of Service Configuration Guide*.

Configuring Symmetrical Load Balancing on an 802.3ad Link Aggregation Group on MX Series Routers

This section describes configuration of symmetrical load balancing on an 802.3ad link aggregation group (LAG) on MX Series routers.

This section includes the following topics:

- Overview of Symmetrical Load Balancing on an 802.3ad LAG on MX Series Routers on page 636
- Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers on page 637
- Example Configurations on page 640

Overview of Symmetrical Load Balancing on an 802.3ad LAG on MX Series Routers

MX Series routers with Aggregated Ethernet PICs support symmetrical load balancing on an 802.3ad LAG. This feature is significant when two MX Series routers are connected transparently through deep packet inspection (DPI) devices over an LAG bundle. DPI devices keep track of flows and require information of a given flow in both forward and reverse directions. Without symmetrical load balancing on an 802.3ad LAG, the DPIs could misunderstand the flow, leading to traffic disruptions. By using this feature, a given flow of traffic (duplex) is ensured for the same devices in both directions.

Symmetrical load balancing on an 802.3ad LAG utilizes a mechanism of interchanging the source and destination addresses for a hash computation of fields, such as source address and destination address. The result of a hash computed on these fields is used to choose the link of the LAG. The hash-computation for the forward and reverse flow must be identical. This is achieved by swapping source fields with destination fields for the reverse flow. The swapped operation is referred to as *complement hash*.

computation or *symmetric-hash complement* and the regular (or unswapped) operation as *symmetric-hash computation* or *symmetric-hash*. The swappable fields are MAC address, IP address, and port.

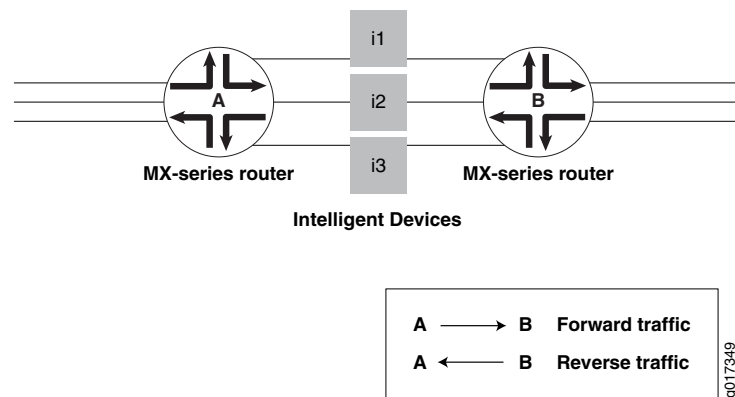
Configuring Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers

You can specify whether symmetric hash or complement hash is done for load-balancing traffic. To configure symmetric hash, use the `symmetric-hash` statement at the `[edit forwarding-options hash-key family inet]` hierarchy level. To configure symmetric hash complement, use the `symmetric-hash <complement>` statement and option at the `[edit forwarding-options hash-key family inet]` hierarchy level.

These operations can also be performed at the PIC level by specifying a *hash key*. To configure a hash key at the PIC level, use the `symmetric-hash` or `symmetric-hash <complement>` statement at the `[edit chassis hash-key family inet]` and `[edit chassis hash-key family multiservice]` hierarchy levels.

Consider the example in Figure 59 on page 637.

Figure 59: Symmetric Load Balancing on an 802.3ad LAG on MX Series Routers



Router A is configured with symmetric hash, and Router B is configured with symmetric hash complement. Thus, for a given flow fx , post hash computation is from Router A to Router B through i2. The reverse traffic for the same flow fx is from Router B to Router A through the same i2 device as its hashing (done after swapping source and destination fields) and returns the same link index; since it is performed on the interchanged source and destination addresses.

However, the link chosen may or may not correspond to what was attached to the DPI. In other words, the hashing result should point to the same links that are connected, so that the traffic flows through the same DPI devices in both directions. To make sure this happens, you need to also configure the counterpart ports (ports that are connected to same DPI-iN) with the identical link index. This is done when configuring a child-link into the LAG bundle. This ensures that the link chosen for a given hash result is always the same on either router.

Note that any two links connected to each other should have the same link index and these link indices must be unique in a given bundle.

**NOTE:**

The following restrictions apply when configuring symmetric load balancing on an 802.3ad LAG on MX Series routers:

- The Packet Forwarding Engine (PFE) can be configured to hash the traffic in either symmetric or complement mode. A single PFE complex cannot work simultaneously in both operational modes and such a configuration can yield undesirable results.
- The per-PFE setting overrides the chassis-wide setting only for the family configured. For the other families, the PFE complex still inherits the chassis-wide setting (when configured) or the default setting.
- Any change in the hash key configuration requires a reboot of the FPC for the changes to take effect.
- This feature supports VPLS, INET, and bridged traffic only.
- This feature cannot work in tandem with the **per-flow-hash-seed** <load-balancing> option. It requires that all the PFE complexes configured in complementary fashion share the same seed. A change in the seed between two counterpart PFE complexes may yield undesired results.

For additional information, see the *VPNs Configuration Guide* and the *System Basics Configuration Guide*.

Example Configuration Statements

To configure 802.3ad LAG parameters at the bundle level:

```
[edit interfaces]
g(x)e-fpc/pic/port {
  together-options {
    802.3ad {
      bundle;
      link-index number;
    }
  }
}
```

where the link-index *number* ranges from 0 through 15.

You can check the link index configured above using the **show interfaces** command:

```
[edit forwarding-options hash-key]
family inet {
  layer-3;
  layer-4;
  symmetric-hash {
    [complement;]
  }
}
family multiservice {
  source-mac;
  destination-mac;
```

```

payload {
  ip {
    layer-3 {
      source-ip-only | destination-ip-only;
    }
    layer-4;
  }
}
symmetric-hash {
  [complement;]
}
}

```

For load-balancing Layer 2 traffic based on Layer 3 fields, you can configure 802.3ad LAG parameters at a per PIC level. These configuration options are available under the chassis hierarchy as follows:

```

[edit chassis]
fpc X {
  pic Y {
    .
    .
    .
    hash-key {
      family inet {
        layer-3;
        layer-4;
        symmetric-hash {
          [complement;]
        }
      }
    }
    family multiservice {
      source-mac;
      destination-mac;
      payload {
        ip {
          layer-3 {
            source-ip-only | destination-ip-only;
          }
          layer-4;
        }
      }
      symmetric-hash {
        [complement;]
      }
    }
  }
  .
  .
  .
}

```

Example Configurations

Example Configurations of Chassis Wide Settings

Router A: user@host> **show configuration forwarding-options hash-key**
 family multiservice {
 payload {
 ip {
 layer-3;
 }
 }
 symmetric hash;
 }

Router B: user@host> **show configuration forwarding-options hash-key**
 family multiservice {
 payload {
 ip {
 layer-3;
 }
 }
 symmetric-hash {
 complement;
 }
 }

Example Configurations of Per PFE Settings

Router A: user@host> **show configuration chassis fpc 2 pic 2 hash-key**
 family multiservice {
 payload {
 ip {
 layer-3;
 }
 }
 symmetric hash;
 }

Router B: user@host> **show configuration chassis fpc 2 pic 3 hash-key**
 family multiservice {
 payload {
 ip {
 layer-3;
 }
 }
 symmetric-hash {
 complement;
 }
 }

Chapter 36

Stacking and Rewriting Gigabit Ethernet VLAN Tags

This section discusses the following topics:

- Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview on page 641
- Stacking and Rewriting Gigabit Ethernet VLAN Tags on page 642
- Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames on page 644
- Configuring Stacked VLAN Tagging on page 645
- Configuring Dual VLAN Tags on page 645
- Configuring Inner and Outer TPIDs and VLAN IDs on page 645
- Stacking a VLAN Tag on page 648
- Removing a VLAN Tag on page 649
- Removing the Outer and Inner VLAN Tags on page 649
- Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag on page 650
- Stacking Two VLAN Tags on page 651
- Rewriting the VLAN Tag on Tagged Frames on page 651
- Rewriting a VLAN Tag on Untagged Frames on page 652
- Rewriting a VLAN Tag and Adding a New Tag on page 655
- Rewriting the Inner and Outer VLAN Tags on page 655
- Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags on page 656

Stacking and Rewriting Gigabit Ethernet VLAN Tags Overview

On Gigabit Ethernet, Gigabit Ethernet IQ, Gigabit Ethernet IQ2 and IQ2-E, 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, and MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces with the VLAN encapsulation type configured to support Layer 2 tunneling protocols such as CCC or VPLS (as described in “Configuring 802.1Q VLANs” on page 599), you can stack and rewrite VLAN tags. Stacking and rewriting VLAN tags allows you to use an additional (outer) VLAN tag to differentiate between customer edge (CE) routers that share one VLAN ID. A frame can be received on an interface, or it can be internal to the system (as a result of the `input-vlan-map` statement).



NOTE: On IQ2 and IQ2-E interfaces and MX Series interfaces, when a VLAN tag is pushed, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the VLAN or VLANs being pushed. If the original packet is untagged, the IEEE bits of the VLAN or VLANs being pushed are set to 0.

Stacking and Rewriting Gigabit Ethernet VLAN Tags

You can configure rewrite operations to stack (**push**), remove (**pop**), or rewrite (**swap**) tags on single-tagged frames and dual-tagged frames. If a port is not tagged, rewrite operations are not supported on any logical interface on that port.

You can configure the following VLAN rewrite operations:

- **pop**—Remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.
- **pop-pop**—For Ethernet IQ2 and IQ2-E interfaces, remove both the outer and inner VLAN tags of the frame.
- **pop-swap**—For Ethernet IQ2 and IQ2-E interfaces, remove the outer VLAN tag of the frame, and replace the inner VLAN tag of the frame with a user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.
- **push**—Add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag.
- **push-push**—For Ethernet IQ2 and IQ2-E interfaces, push two VLAN tags in front of the frame.
- **swap-push**—For Ethernet IQ2 and IQ2-E interfaces, replace the outer VLAN tag of the frame with a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.
- **swap-swap**—For Ethernet IQ2 and IQ2-E interfaces, replace both the inner and the outer VLAN tags of the incoming frame with a user-specified VLAN tag value.

You configure VLAN rewrite operations for logical interfaces in the input VLAN map for incoming frames and in the output VLAN map for outgoing frames. To configure the input VLAN map, include the **input-vlan-map** statement:

```
input-vlan-map {
  ...interface-specific configuration...
}
```

To configure the output VLAN map, include the **output-vlan-map** statement:

```
output-vlan-map {
  ...interface-specific configuration...
}
```

You can include both statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]

- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The type of VLAN rewrite operation permitted depends upon whether the frame is single-tagged or dual-tagged. Table 52 on page 643 shows supported rewrite operations and whether they can be applied to single-tagged frames or dual-tagged frames. The table also indicates the number of tags being added or removed during the operation.

Table 52: Rewrite Operations on Not Tagged, Single-Tagged, and Dual-Tagged Frames

Rewrite Operation	Not Tagged	Single-Tagged	Dual-Tagged	Number of Tags
pop	No	Yes	Yes	– 1
push	Sometimes	Yes	Yes	+ 1
swap	No	Yes	Yes	0
push-push	Sometimes	Yes	Yes	+ 2
swap-push	No	Yes	Yes	+ 1
swap-swap	No	No	Yes	0
pop-pop	No	No	Yes	– 2
pop-swap	No	No	Yes	– 1

The rewrite operations **push** and **push-push** can be valid in certain circumstances on frames that are not tagged. For example, a single-tagged logical interface (interface 1) and a dual-tagged logical interface (interface 2) have the following configurations:

```

Interface 1 [edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
    pop;
}
output-vlan-map {
    push;
}

Interface 2 [edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
    pop-pop;
}
output-vlan-map {
    push-push;
}

```

When a frame is received on the interface as a result of the **input-vlan-map** operation, the frame is not tagged. As it goes out of the second interface, the **output-vlan-map** operation **push-push** is applied to it. The resulting frame will be dual-tagged at the logical interface output.

Depending on the VLAN rewrite operation, you configure the rewrite operation for the interface in the input VLAN map, the output VLAN map, or in both the input VLAN map and the output VLAN map. Table 53 on page 644 shows what rewrite operation combinations you can configure. “None” means that no rewrite operation is specified for the VLAN map.

Table 53: Applying Rewrite Operations to VLAN Maps

Input VLAN Map	Output VLAN Map								
	none	push	pop	swap	push-push	swap-push	swap-swap	pop-pop	swap-pop
none	Yes	No	No	Yes	No	No	Yes	No	No
push	No	No	Yes	No	No	No	No	No	No
pop	No	Yes	No	No	No	No	No	No	No
swap	Yes	No	No	Yes	No	No	No	No	No
push-push	No	No	No	No	No	No	No	Yes	No
swap-push	No	No	No	No	No	No	No	No	Yes
swap-swap	Yes	No	No	No	No	No	Yes	No	No
pop-pop	No	No	No	No	Yes	No	No	No	No
pop-swap	No	No	No	No	No	Yes	No	No	No

As well as knowing if the VLAN rewrite operation is valid, and whether it is applied to the input VLAN map or the output VLAN map, you must also know whether the rewrite operation requires you to include statements to configure the inner and outer TPIDs and inner and outer VLAN IDs in the input VLAN map or output VLAN map. For information about configuring inner and outer TPIDs and inner and outer VLAN IDs, see “Configuring Inner and Outer TPIDs and VLAN IDs” on page 645.

Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames

For Gigabit Ethernet IQ interfaces, aggregated Ethernet with Gigabit Ethernet IQ interfaces, Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), and MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, you can configure frames with particular TPIDs to be processed as tagged frames. To do this, you specify up to eight IEEE 802.1Q TPID values per port; a frame with any of the specified TPIDs is processed as a tagged frame; however, with IQ2 and IQ2-E interfaces, only the first four IEEE 802.1Q TPID values per port are supported. To configure the TPID values, include the `tag-protocol-id` statement:

```
tag-protocol-id [ tpids ];
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* gigether-options ethernet-switch-profile]
- [edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile]

All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces *interface-name* gigether-options ethernet-switch-profile tag-protocol-id [*tpids*]] or [edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile tag-protocol-id [*tpids*]] hierarchy level.

Configuring Stacked VLAN Tagging

To configure stacked VLAN tagging for all logical interfaces on a physical interface, include the **stacked-vlan-tagging** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
  stacked-vlan-tagging;
```

If you include the **stacked-vlan-tagging** statement in the configuration, you must configure dual VLAN tags for all logical interfaces on the physical interface. For more information, see “Stacking a VLAN Tag” on page 648.

Configuring Dual VLAN Tags

To configure dual VLAN tags on a logical interface, include the **vlan-tags** statement:

```
vlan-tags (Stacked VLAN Tags) inner <tpid.>vlan-id outer <tpid.>vlan-id;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The outer tag VLAN ID range is from 1 through 511 for normal interfaces, and from 512 through 4094 for VLAN CCC or VLAN VPLS interfaces. The inner tag is not restricted.

You must also include the **stacked-vlan-tagging** statement in the configuration. See “Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags” on page 656.

Configuring Inner and Outer TPIDs and VLAN IDs

For some rewrite operations, you must configure the inner or outer TPID values and inner or outer VLAN ID values. These values can be applied to either the input VLAN map or the output VLAN map.

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, include the **inner-tag-protocol-id** statement to configure the inner

TPID. For the inner VLAN ID, include the `inner-vlan-id` statement. For the outer TPID, include the `tag-protocol-id` statement. For the outer VLAN ID, include the `vlan-id` statement:

```
input-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
output-vlan-map {
  (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
  inner-tag-protocol-id tpid;
  inner-vlan-id number;
  tag-protocol-id tpid;
  vlan-id number;
}
```

For aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, include the `tag-protocol-id` statement for the outer TPID. For the outer VLAN ID, include the `vlan-id` statement:

```
input-vlan-map {
  (pop | push | swap);
  tag-protocol-id tpid;
  vlan-id number;
}
output-vlan-map {
  (pop | push | swap);
  tag-protocol-id tpid;
  vlan-id number;
}
```

For the input VLAN map, include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]

For the output VLAN map, include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

The VLAN IDs you define in the input VLAN maps are stacked on top of the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see “Configuring 802.1Q VLANs” on page 599.

All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces *interface-name* *gigether-options* ethernet-switch-profile tag-protocol-id [*tpids*]] hierarchy level or [edit interfaces *interface-name*

aggregated-ether-options ethernet-switch-profile tag-protocol-id [*tpids*] hierarchy level. For more information, see “Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames” on page 644.

Table 54 on page 647 and Table 55 on page 647 specify when these statements are required. Table 54 on page 647 indicates valid statement combinations for rewrite operations for the input VLAN map. “No” means the statement must not be included in the input VLAN map for the rewrite operation. “Optional” means the statement may be optionally specified for the rewrite operation in the input VLAN map. “Any” means that you must include the `vlan-id` statement, `tag-protocol-id` statement, `inner-vlan-id` statement, or `inner-tag-protocol-id` statement.

Table 54: Rewrite Operations and Statement Usage for Input VLAN Maps

Input VLAN Map Statements				
Rewrite Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	Optional	Optional	No	No
pop	No	No	No	No
swap	Any	Any	No	No
push-push	Optional	Optional	Optional	optional
swap-push	Optional	Optional	Any	Any
swap-swap	Optional	Optional	Any	Any
pop-swap	No	No	Any	Any
pop-pop	No	No	No	No

Table 55 on page 647 indicates valid statement combinations for rewrite operations for the output VLAN map. “No” means the statement must not be included in the output VLAN map for the rewrite operation. “Optional” means the statement may be optionally specified for the rewrite operation in the output VLAN map.

Table 55: Rewrite Operations and Statement Usage for Output VLAN Maps

Output VLAN Map Statements				
Rewrite Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	No	Optional	No	No
pop	No	No	No	No
swap	No	Optional	No	No
push-push	No	Optional	No	Optional
swap-push	No	Optional	No	Optional

Table 55: Rewrite Operations and Statement Usage for Output VLAN Maps (*continued*)

Output VLAN Map Statements				
swap-swap	No	Optional	No	Optional
pop-swap	No	No	No	Optional
pop-pop	No	No	No	No

The following examples use Table 54 on page 647 and Table 55 on page 647 and show how the **pop-swap** operation can be configured in an input VLAN map and an output VLAN map:

Input VLAN Map with inner-vlan-id Statement, Output VLAN Map with Optional inner-tag-protocol-id Statement

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-swap;
  inner-vlan-id number;
}
output-vlan-map {
  pop-swap;
  inner-tag-protocol-id tpid;
}
```

Input VLAN Map with inner-tag-protocol-id Statement, Output VLAN Map with Optional inner-tag-protocol-id Statement

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-swap;
  inner-tag-protocol-id tpid;
}
output-vlan-map {
  pop-swap;
  inner-tag-protocol-id tpid;
}
```

Input VLAN Map with inner-tag-protocol-id and inner-vlan-id Statements

```
[edit interfaces interface-name unit logical-unit-number]
input-vlan-map {
  pop-swap;
  inner-vlan-id number;
  inner-tag-protocol-id tpid;
}
```

Stacking a VLAN Tag

To stack a VLAN tag on all tagged frames entering or exiting the interface, include the **push**, **vlan-id**, and **tag-protocol-id** statements in the input VLAN map or the output VLAN map:

```
input-vlan-map {
  push;
  vlan-id number;
  tag-protocol-id tpid;
}
```



```
output-vlan-map {
    push;
    tag-protocol-id tpid;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

If you include the **push** statement in an interface’s input VLAN map, see Table 53 on page 644 for information about permissible rewrite operations,

The VLAN IDs you define in the input VLAN maps are stacked on top of the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see “Configuring 802.1Q VLANs” on page 599.

All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces *interface-name* **gether-options** ethernet-switch-profile tag-protocol-id [*tpids*]] hierarchy level. For more information, see “Configuring Inner and Outer TPIDs and VLAN IDs” on page 645.

Removing a VLAN Tag

To remove a VLAN tag from all tagged frames entering or exiting the interface, include the **pop** statement in the input VLAN map or output VLAN map:

```
input-vlan-map {
    pop;
}
output-vlan-map {
    pop;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Removing the Outer and Inner VLAN Tags

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs

on MX Series routers, to remove both the outer and inner VLAN tags of the frame, include the **pop-pop** statement in the input VLAN map or output VLAN map:

```
input-vlan-map {
  pop-pop;
}
output-vlan-map {
  pop-pop;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

See Table 54 on page 647 and Table 55 on page 647 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to remove the outer VLAN tag of the frame and replace the inner VLAN tag of the frame with a user-specified VLAN tag value, include the **pop-swap** statement in the input VLAN map or output VLAN map:

```
input-vlan-map {
  pop-swap;
}
output-vlan-map {
  pop-swap;
}
```

The inner tag becomes the outer tag in the final frame.

You can include this statements a the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

See Table 54 on page 647 and Table 55 on page 647 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Stacking Two VLAN Tags

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to push two VLAN tags in front of tagged frames entering or exiting the interface, include the **push-push** statement in the input VLAN map or the output VLAN map:

```
input-vlan-map {
    push-push;
}
output-vlan-map {
    push-push;
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

See Table 54 on page 647 and Table 55 on page 647 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Rewriting the VLAN Tag on Tagged Frames

To rewrite the VLAN tag on all tagged frames entering the interface to a specified VLAN ID and TPID, include the **swap**, **tag-protocol-id**, and **vlan-id** statements in the input VLAN map:

```
input-vlan-map {
    swap;
    vlan-id number;
    tag-protocol-id tpid;
}
```

To rewrite the VLAN tag on all tagged frames exiting the interface to a specified VLAN ID and TPID, include the **swap** and **tag-protocol-id** statements in the output VLAN map:

```
output-vlan-map {
    swap;
```

```

    vlan-id number;
    tag-protocol-id tpid;
}

```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]

You cannot include both the **swap** statement and the **vlan-id** statement in the output VLAN map configuration. If you include the **swap** statement in the configuration, the VLAN ID in outgoing frames is rewritten to the VLAN ID bound to the logical interface. For more information about binding a VLAN ID to the logical interface, see “Configuring 802.1Q VLANs” on page 599.

The swap operation works on the outer tag only, whether or not you include the **stacked-vlan-tagging** statement in the configuration. For more information, see “Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags” on page 656.

Rewriting a VLAN Tag on Untagged Frames

On M320, M120, and MX Series routers with Gigabit Ethernet IQ, IQ2, and IQ2E PICs, 10-Gigabit Ethernet IQ, IQ2, and IQ2E PICs, and on MX Series 40-port Gigabit Ethernet R, 40-port Gigabit Ethernet R EQ, 4-port 10-Gigabit Ethernet R, and 4-port 10-Gigabit Ethernet R EQ DPCs, you can rewrite VLAN tags on untagged incoming and outgoing frames under **ethernet-ccc** and **ethernet-vpls** encapsulations. On MX Series routers with IQ2 and IQ2-E PICs, you can perform all rewrite VLAN tag operations. These features provide added flexibility.

Consider a network where two provider edges (PE) are connected by a Layer 2 circuit. PE1 is receiving traffic on an untagged port while the corresponding port on PE2 is tagged. In the normal case, packets coming from PE1 will be dropped at PE2 because it is expecting tagged packets. However, if PE1 can push a VLAN tag on the incoming packet before sending it across to PE2, you can ensure that packets are not dropped. To make it work in both directions, PE1 must strip the VLAN tag from outgoing packets. Therefore, a push on the ingress side is always paired with a pop on the egress side.

The rewrite operations represented by the following statement options are supported under **ethernet-ccc** and **ethernet-vpls** encapsulations:

- **push**—A VLAN tag is added to the incoming untagged frame.
- **pop**—VLAN tag is removed from the outgoing frame.
- **push-push**—An outer and inner VLAN tag are added to the incoming untagged frame.
- **pop-pop**—Both the outer and inner VLAN tags of the outgoing frame are removed.

IQ2 and 10-Gigabit Ethernet PICs support all rewrite operations described above. Details on the possible combinations of usage are explained later in this section.



NOTE: The **push-push** and **pop-pop** operations are not supported on the Gigabit Ethernet IQ PIC.

For the **input-vlan-map** statement, only the **push** and **push-push** options are supported because it does not make sense to remove a VLAN tag from an incoming untagged frame. Similarly, only the **pop** and **pop-pop** options are supported for the **output-vlan-map** statement. Also, with the **push** and **push-push** options, the tag parameters have to be explicitly specified. Apart from this, the other rules for configuring the **input-vlan-map** and **output-vlan-map** statements are the same as for tagged frames. Table 56 on page 653 through Table 58 on page 653 explain the rules in more detail.

For the **input-vlan-map** statement, only the **push** and **push-push** options are supported because it does not make sense to remove a VLAN tag from an incoming untagged frame. Similarly, only the **pop** and **pop-pop** options are supported for the **output-vlan-map** statement. Also, with the **push** and **push-push** options, the **vlan-id** parameters (**vlan-id** for **push** and **vlan-id** or **inner-vlan-id** for **push-push**) have to be explicitly specified. TPID however, is optional and the default value of **0x8100** is set if not configured. Apart from this, the other rules for configuring the **input-vlan-map** and **output-vlan-map** statements are the same as for tagged frames.

Table 56: Input VLAN map statements allowed for ethernet-ccc and ethernet-vpls encapsulations

Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
push	Yes	Optional	No	Optional
push-push	Yes	Optional	Yes	Optional

Table 57: Output VLAN map statements allowed for ethernet-ccc and ethernet-vpls encapsulations

Operation	vlan-id	tag-protocol-id	inner-vlan-id	inner-tag-protocol-id
pop	No	No	No	No
pop-pop	No	No	No	No

Table 58: Rules for applying rewrite operations to VLAN maps

Output VLAN Map			
Input VLAN Map	None	pop	pop-pop
None	Yes	No	No
push	No	Yes	No
push-push	No	No	Yes

**Example: push and pop
with Ethernet CCC
Encapsulation**

```

ge-3/1/0 {
  encapsulation ethernet-ccc;
  unit 0 {
    encapsulation ethernet-ccc;
    input-vlan-map {
      push;
      tag-protocol-id 0x8100;
      vlan-id 600;
    }
    output-vlan-map pop;
    family ccc;
  }
}

```

**Example: push-push and
pop-pop with Ethernet
CCC Encapsulation**

```

ge-3/1/0 {
  encapsulation ethernet-ccc;
  unit 0 {
    encapsulation ethernet-ccc;
    input-vlan-map {
      push-push;
      tag-protocol-id 0x8100;
      inner-tag-protocol-id 0x8100;
      vlan-id 600;
      inner-vlan-id 575;
    }
    output-vlan-map pop-pop;
    family ccc;
  }
}

```

**Example: push and pop
with Ethernet VPLS
Encapsulation**

```

ge-3/1/0 {
  encapsulation ethernet-vpls;
  unit 0 {
    encapsulation ethernet-vpls;
    input-vlan-map {
      push;
      tag-protocol-id 0x8100;
      vlan-id 700;
    }
    output-vlan-map pop;
    family vpls;
  }
}

```

**Example: push-push and
pop-pop with Ethernet
VPLS Encapsulation**

```

ge-3/1/0 {
  encapsulation ethernet-vpls;
  unit 0 {
    encapsulation ethernet-vpls;
    input-vlan-map {
      push-push;
      tag-protocol-id 0x8100;
      inner-tag-protocol-id 0x8100;
      vlan-id 600;
      inner-vlan-id 575;
    }
  }
}

```

```

    }
    output-vlan-map pop-pop;
    family vpls;
  }
}

```

You can use the `show interface interface-name` command to display the status of a modified VLAN map for the specified interface.

Rewriting a VLAN Tag and Adding a New Tag

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to replace the outer VLAN tag of the incoming frame with a user-specified VLAN tag value, include the **swap-push** statement in the input VLAN map or output VLAN map:

```

input-vlan-map {
    swap-push;
}
output-vlan-map {
    swap-push;
}

```

A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

See Table 54 on page 647 and Table 55 on page 647 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Rewriting the Inner and Outer VLAN Tags

On Ethernet IQ, IQ2 and IQ2-E interfaces, on MX Series router Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, and on aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, to replace both the inner and the outer VLAN tags of the incoming frame with a user-specified VLAN tag value, include the **swap-swap** statement in the input VLAN map or output VLAN map:

```

input-vlan-map {
    swap-swap;
}
output-vlan-map {

```

```

        swap-swap;
    }

```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* output-vlan-map]
- [edit interfaces *interface-name* unit *logical-unit-number* input-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* output-vlan-map]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* input-vlan-map]

See Table 54 on page 647 and Table 55 on page 647 for information about configuring inner and outer VLAN ID values and inner and outer TPID values required for VLAN maps.

Examples: Stacking and Rewriting Gigabit Ethernet IQ VLAN Tags

Configure a VLAN CCC tunnel in which Ethernet frames enter the tunnel at interface *ge-4/0/0* and exit the tunnel at interface *ge-4/2/0*.

The following examples show how to perform the following tasks:

- Push a TPID and VLAN ID pair on ingress.
- Stack inner and outer VLAN tags.
- Swap a VLAN ID on ingress.
- Swap a VLAN ID on egress.
- Swap a VLAN ID on both ingress and egress.
- Swap the outer VLAN tag and push a new VLAN tag on ingress; pop the outer VLAN tag and swap the inner VLAN tag on egress.
- Swap a VLAN ID and TPID pair for both VLAN tags on ingress and on egress.
- Pop the outer VLAN tag and swap the inner VLAN tag on ingress; swap the outer VLAN tag and push a new VLAN tag on egress.
- Pop two VLAN ID and TPID pairs on ingress; push two VLAN ID and TPID pairs on egress.

Push a TPID and VLAN ID Pair on Ingress

```

[edit interfaces]
ge-4/0/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  gigether-options {
    ethernet-switch-profile {
      tag-protocol-id 0x9909;
    }
  }
  unit 0 {
    encapsulation vlan-ccc;
    vlan-id 512;
  }
}

```



```

        input-vlan-map {
            push;
            tag-protocol-id 0x9909;
            vlan-id 520;
        }
        output-vlan-map pop;
    }
}
ge-4/2/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    unit 0 {
        encapsulation vlan-ccc;
        vlan-id 515;
        input-vlan-map {
            swap-push;
            vlan-id 520;
            inner-vlan-id 512;
        }
        output-vlan-map {
            pop-swap;
        }
    }
    [edit protocols]
    mpls {
        interface ge-4/0/0.0;
        interface ge-4/2/0.0;
    }
    connections {
        interface-switch vlan-tag-push {
            interface ge-4/0/0.0;
            interface ge-4/2/0.0;
        }
    }
}

```

Stack Inner and Outer VLAN Tags

```

[edit interfaces]
ge-0/2/0 {
    stacked-vlan-tagging;
    mac 00.01.02.03.04.05;
    gigether-options {
        loopback;
    }
    unit 0 {
        vlan-tags outer 0x8100.200 inner 0x8100.200;
    }
}

```

Swap a VLAN ID on Ingress

```

[edit interfaces]
ge-4/0/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    gigether-options {
        ethernet-switch-profile {
            tag-protocol-id 0x9100;
        }
    }
}

```

```

    }
  }
  ...
  unit 1 {
    encapsulation vlan-ccc;
    vlan-id 1000;
    input-vlan-map {
      swap;
      tag-protocol-id 0x9100;
      vlan-id 2000;
    }
  }
}
ge-4/2/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  ...
  unit 1 {
    encapsulation vlan-ccc;
    vlan-id 2000;
    input-vlan-map {
      swap;
      tag-protocol-id 0x9100;
      vlan-id 1000;
    }
  }
}
[edit protocols]
mpls {
  ...
  interface ge-4/0/0.1;
  interface ge-4/2/0.1;
}
connections {
  ...
  interface-switch vlan-tag-swap {
    interface ge-4/2/0.1;
    interface ge-4/0/0.1;
  }
}
}

```

**Swap a VLAN ID on
Egress**

```

[edit interfaces]
ge-4/0/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  ...
  unit 1 {
    encapsulation vlan-ccc;
    vlan-id 1000;
  }
}
ge-4/2/0 {
  vlan-tagging;
  encapsulation vlan-ccc;
  together-options {

```

```

        ethernet-switch-profile {
            tag-protocol-id 0x8800;
        }
    }
    ...
    unit 1 {
        encapsulation vlan-ccc;
        vlan-id 2000;
        output-vlan-map {
            swap;
            tag-protocol-id 0x8800;
        }
    }
}
[edit protocols]
mpls {
    ...
    interface ge-4/0/0.1;
    interface ge-4/2/0.1;
}
connections {
    ...
    interface-switch vlan-tag-swap {
        interface ge-4/2/0.1;
        interface ge-4/0/0.1;
    }
}

```

**Swap a VLAN ID on Both
Ingress and Egress**

```

[edit interfaces]
ge-4/0/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    together-options {
        ethernet-switch-profile {
            tag-protocol-id [ 0x8800 0x9100 ];
        }
    }
}
...
unit 1 {
    encapsulation vlan-ccc;
    vlan-id 1000;
    input-vlan-map {
        swap;
        tag-protocol-id 0x9100;
        vlan-id 2000;
    }
}
}
ge-4/2/0 {
    vlan-tagging;
    encapsulation vlan-ccc;
    together-options {
        ethernet-switch-profile {
            tag-protocol-id [ 0x8800 0x9100 ];
        }
    }
}

```

```

    }
    unit 1 {
        encapsulation vlan-ccc;
        vlan-id 2000;
        output-vlan-map {
            swap;
            tag-protocol-id 0x8800;
        }
    }
}
[edit protocols]
mpls {
    ...
    interface ge-4/0/0.1;
    interface ge-4/2/0.1;
}
connections {
    ...
    interface-switch vlan-tag-swap {
        interface ge-4/2/0.1;
        interface ge-4/0/0.1;
    }
}

```

Swap the Outer VLAN Tag and Push a New VLAN Tag on Ingress; Pop the Outer VLAN Tag and Swap the Inner VLAN Tag on Egress

```

[edit interfaces]
ge-1/1/0 {
    unit 1 {
        vlan-id 200;
        input-vlan-map {
            swap-push;
            tag-protocol-id 0x9100;
            vlan-id 400;
            inner-tag-protocol-id 0x9100;
            inner-vlan-id 500;
        }
        output-vlan-map {
            pop-swap;
            inner-tag-protocol-id 0x9100;
        }
    }
}

```

Swap a TPID and VLAN ID Pair for Both VLAN Tags on Ingress and on Egress

```

[edit interfaces]
ge-1/1/0 {
    unit 0 {
        vlan-tags {
            inner 0x9100.425;
            outer 0x9200.525;
        }
        input-vlan-map {
            swap-swap;
            tag-protocol-id 0x9100;
            vlan-id 400;
            inner-tag-protocol-id 0x9100;
            inner-vlan-id 500;
        }
    }
}

```

```

    }
    output-vlan-map {
        swap-swap;
        tag-protocol-id 0x9200;
        inner-tag-protocol-id 0x9100;
    }
}

```

Pop the Outer VLAN Tag and Swap the Inner VLAN Tag on Ingress; Swap the Outer VLAN Tag and Push a New VLAN Tag on Egress

```

[edit interfaces]
ge-1/1/0 {
    unit 0 {
        vlan-tags {
            inner 0x9100.425;
            outer 0x9200.525;
        }
        input-vlan-map {
            pop-swap;
            tag-protocol-id 0x9100;
            vlan-id 400;
        }
        output-vlan-map {
            swap-push;
            tag-protocol-id 0x9200;
            inner-tag-protocol-id 0x9100;
        }
    }
}

```

Pop a TPID and VLAN ID Pair on Ingress; Push a VLAN ID and TPID Pair on Egress

```

[edit interfaces]
ge-1/1/0 {
    unit 0 {
        vlan-tags {
            inner 0x9100.425;
            outer 0x9200.525;
        }
        input-vlan-map {
            pop-pop;
        }
        output-vlan-map {
            push-push;
            tag-protocol-id 0x9200;
            inner-tag-protocol-id 0x9100;
        }
    }
}

```

POP an Outer VLAN Tag to Connect an Untagged VPLS Interface to Tagged VPLS Interfaces

```

[edit interfaces]
ge-1/1/0 {
    vlan-tagging;
    encapsulation extended-vlan-vpls;
    unit 0 {
        vlan-id 0;
        input-vlan-map {
            push;
        }
    }
}

```

```
        vlan-id 0;  
    }  
    output-vlan-map pop;  
    family vpls;  
}  
}
```

Chapter 37

Configuring Layer 2 Bridging Interfaces

This section contains the following topics:

- Layer 2 Bridging Interfaces Overview on page 663
- Configuring Layer 2 Bridging Interfaces on page 663

Layer 2 Bridging Interfaces Overview

Bridging operates at Layer 2 of the OSI reference model while routing operates at Layer 3. A set of logical ports configured for bridging can be said to constitute a bridging domain.

A bridging domain can be created by configuring a routing instance and specifying the instance-type as **bridge**.

Integrated routing and bridging (IRB) is the ability to:

- Route a packet if the destination MAC address is the MAC address of the router and the packet **ethertype** is IPv4, IPv6, or MPLS.
- Switch all multicast and broadcast packets within a bridging domain at layer 2.
- Route a copy of the packet if the destination MAC address is a multicast address and the **ethertype** is IPv4 or IPv6.
- Switch all other unicast packets at Layer 2.
- Handle supported Layer 2 control packets such as STP and LACP.
- Handle supported Layer 3 control packets such as OSPF and RIP.

Configuring Layer 2 Bridging Interfaces

You can configure an IRB logical interface at the [edit interfaces *ge-fpc /pic/port* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces ge-fpc/pic/port]  
unit logical-unit-number {  
}
```

You can configure Layer 3 information on the IRB logical interface by including the **irb** statement at the [edit interfaces] hierarchy level:

```
[edit interfaces]
irb {
  unit logical-unit-number {
    family inet {
      address address {
      }
    }
  }
}
```

For examples of Layer 2 bridging configuration, see the *JUNOS Routing Protocols Configuration Guide*.

Example: Configuring Layer 2 Bridging Interfaces

The following example configures an IRB logical interface and Layer 3 information on the interface.

```
[edit interfaces]
ge-1/0/0 {
  unit 0 {
  }
}
irb {
  unit 0 {
    family inet {
      address 192.168.12.1/28;
    }
  }
}
```


Chapter 38

Configuring TCC and Layer 2.5 Switching

This section contains the following topics:

- TCC and Layer 2.5 Switching Overview on page 665
- Configuring VLAN TCC Encapsulation on page 665
- Configuring Ethernet TCC on page 666

TCC and Layer 2.5 Switching Overview

Translational cross-connect (TCC) is a switching concept that allows you to forward traffic between a variety of Layer 2 protocols or circuits. It is similar to its predecessor, CCC. However, while CCC requires the same Layer 2 encapsulations on both sides of a router (such as Point-to-Point Protocol [PPP] or Frame Relay-to-Frame Relay), TCC lets you connect different types of Layer 2 protocols interchangeably. With TCC, combinations such as PPP-to-ATM and Ethernet-to-Frame Relay cross-connections are possible.

Configuring VLAN TCC Encapsulation

VLAN TCC encapsulation allows circuits to have different media on either side of the forwarding path. VLAN TCC encapsulation supports TPID 0x8100 only. You must include configuration statements at the logical and physical interface hierarchy levels.

To configure VLAN TCC encapsulation, include the **encapsulation** statement and specify the **vlan-tcc** option:

```
[edit interfaces interface-name unit logical-unit-number]  
encapsulation vlan-tcc;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Additionally, configure the logical interface by including the **proxy** and **remote** statements:

```
proxy {  
    inet-address;
```

```

}
remote {
  (inet-address | mac-address);
}

```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family tcc]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family tcc]

The proxy address is the IP address of the non-Ethernet TCC neighbor for which the TCC router is acting as a proxy.

The remote address is the IP or MAC address of the remote router. The **remote** statement provides ARP capability from the TCC switching router to the Ethernet neighbor. The MAC address is the physical Layer 2 address of the Ethernet neighbor.

When VLAN TCC encapsulation is configured on the logical interface, you also must specify flexible Ethernet services on the physical interface. To specify flexible Ethernet services, include the **encapsulation** statement at the [edit interfaces *interface-name*] hierarchy level and specify the **flexible-ethernet-services** option:

```

[edit interfaces interface-name]
encapsulation flexible-ethernet-services;

```

Extended VLAN TCC encapsulation supports TPIDs 0x8100 and 0x9901. Extended VLAN TCC is specified at the physical interface level. When configured, all units on that interface must use VLAN TCC encapsulation, and no explicit configuration is needed on logical interfaces.

One-port Gigabit Ethernet, 2-port Gigabit Ethernet, and 4-port Fast Ethernet PICs with VLAN tagging enabled can use VLAN TCC encapsulation. To configure the encapsulation on a physical interface, include the **encapsulation** statement at the [edit interfaces *interface-name*] hierarchy level and specify the **extended-vlan-tcc** option:

```

[edit interfaces interface-name]
encapsulation extended-vlan-tcc;

```

For VLAN TCC encapsulation, all VLAN IDs from 1 through 1024 are valid. VLAN ID 0 is reserved for tagging the priority of frames.

Extended VLAN TCC is not supported on 4-port Gigabit Ethernet PICs.

Configuring Ethernet TCC

For Layer 2.5 virtual private networks (VPNs) using an Ethernet interface as the TCC router, you can configure an Ethernet TCC.

To configure an Ethernet TCC, include the **encapsulation** statement and specify the **ethernet-tcc** option at the [edit interfaces *interface-name*] hierarchy level:

```

[edit interfaces interface-name]

```

```
encapsulation ethernet-tcc;
```

For Ethernet TCC encapsulation, you must also configure the logical interface by including the **proxy** and **remote** statements:

```
proxy {
    inet-address;
}
remote {
    (inet-address | mac-address);
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family tcc]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family tcc]

The proxy address is the IP address of the non-Ethernet TCC neighbor for which the TCC router is acting as a proxy.

The remote address is the IP or MAC address of the remote router. The **remote** statement provides ARP capability from the TCC switching router to the Ethernet neighbor. The MAC address is the physical Layer 2 address of the Ethernet neighbor.

Ethernet TCC is supported on interfaces that carry IPv4 traffic only. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC and extended VLAN CCC are not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC and extended VLAN TCC are not supported.


Example: Configuring an Ethernet TCC or Extended VLAN TCC

Configure a full-duplex Layer 2.5 translational cross-connect between Router A and Router C, using a Juniper Networks router, Router B, as the TCC interface. Ethernet TCC encapsulation provides an Ethernet wide area circuit for interconnecting IP traffic. (See the topology in Figure 60 on page 667.)

The Router A-to-Router B circuit is PPP, and the Router B-to-Router C circuit accepts packets carrying standard TPID values.

If traffic flows from Router A to Router C, the JUNOS Software strips all PPP encapsulation data from incoming packets and adds Ethernet encapsulation data before forwarding the packets. If traffic flows from Router C to Router A, the JUNOS Software strips all Ethernet encapsulation data from incoming packets and adds PPP encapsulation data before forwarding the packets.

Figure 60: Topology of Layer 2.5 Translational Cross-Connect



On Router B

```

interfaces ge-0/0/0 {
  encapsulation ethernet-tcc;
  unit 0 {
    family tcc {
      proxy {
        inet-address 10.10.10.3;
      }
      remote {
        inet-address 10.10.10.2;
      }
    }
  }
}

```

Configure a full-duplex Layer 2.5 translational cross-connect between Router A and Router C, using a Juniper Networks router, Router B, as the TCC interface. Extended VLAN TCC encapsulation provides an Ethernet wide area circuit for interconnecting IP traffic. (See the topology in Figure 60 on page 667.)

Configuring an Extended VLAN TCC The Router A-to-Router B circuit is PPP, and the Router B-to-Router C circuit is Ethernet with VLAN tagging enabled.

On Router B

```

interfaces ge-0/0/0 {
  vlan-tagging;
  encapsulation extended-vlan-tcc;
  unit 0 {
    vlan-id 1;
    family tcc {
      proxy {
        inet-address 10.10.10.3/24;
      }
      remote {
        inet-address 10.10.10.2/24;
      }
    }
  }
}

```

Chapter 39

Configuring Static ARP Table Entries

This section contains the following topics:

- Static ARP Table Entries Overview on page 669
- Configuring Static ARP Table Entries on page 669

Static ARP Table Entries Overview

For Fast Ethernet, Gigabit Ethernet, Tri-Rate Ethernet copper, and 10-Gigabit Ethernet interfaces, you can configure static ARP table entries, defining mappings between IP and MAC addresses.

Configuring Static ARP Table Entries

To configure static ARP table entries, include the **arp** statement:

```
arp ip-address (mac | multicast-mac) mac-address <publish>;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address*]

The IP address that you specify must be part of the subnet defined in the enclosing **address** statement.

To associate a multicast MAC address with a unicast IP address, include the **multicast-mac** statement.

Specify the MAC address as six hexadecimal bytes in one of the following formats: *nnnn.nnnn.nnnn* or *nn:nn:nn:nn:nn:nn*; for example, 0011.2233.4455 or 00:11:22:33:44:55.

For unicast MAC addresses only, if you include the **publish** option, the router replies to proxy ARP requests.



NOTE: By default, an ARP policer is installed that is shared among all the Ethernet interfaces on which you have configured the `family inet` statement. By including the `arp` statement at the `[edit interfaces interface-name unit logical-unit-number family inet policer]` hierarchy level, you can apply a specific ARP-packet policer to an interface. For more information, see “Applying Policers” on page 194.

When you need to conserve IP addresses, you can configure an Ethernet interface to be unnumbered by including the `unnumbered-address` statement at the `[edit interfaces interface-name unit logical-unit-number family inet]` hierarchy level. For more information, see “Configuring an Unnumbered Interface” on page 185.



NOTE: The JUNOS Software supports the IPv6 static neighbor discovery cache entries, similar to the static ARP entries in IPv4.

Example: Configuring Static ARP Table Entries

Configure two static ARP table entries on the router’s management interface:

```
[edit interfaces]
fxp0 {
  unit 0 {
    family inet {
      address 10.10.0.11/24 {
        arp 10.10.0.99 mac 0001.0002.0003;
        arp 10.10.0.101 mac 00:11:22:33:44:55 publish;
      }
    }
  }
}
```

Chapter 40

Configuring Unrestricted Proxy ARP

This section contains the following topics:

- Unrestricted Proxy ARP Overview on page 671
- Configuring Unrestricted Proxy ARP on page 672

Unrestricted Proxy ARP Overview

By default, the JUNOS Software responds to an ARP request only if the destination address of the ARP request is local to the incoming interface.

For Ethernet interfaces, you can configure unrestricted proxy ARP, which enables the router to respond to any ARP request, on condition that the router has an active route to the destination address of the ARP request. The route is not limited to the incoming interface of the request, nor is it required to be a direct route.

You might want to configure unrestricted proxy ARP for routers that are acting as provider edge (PE) devices in Ethernet Layer 2 LAN switching domains.



WARNING: If you configure unrestricted proxy ARP, the proxy router replies to ARP requests for the target IP address on the same interface as the incoming ARP request. This behavior is appropriate for cable modem termination system (CMTS) environments, but might cause Layer 2 reachability problems if you enable unrestricted proxy ARP in other environments.

When an IP client broadcasts the ARP request across the Ethernet wire, the end node with the correct IP address responds to the ARP request and provides the correct MAC address. If the unrestricted proxy ARP feature is enabled, the router response is redundant and might fool the IP client into determining that the destination MAC address within its own subnet is the same as the address of the router.

While the destination address can be remote, the source address of the ARP request must be on the same subnet as the interface upon which the ARP request is received. For security reasons, this rule applies to both unrestricted and restricted proxy ARP.

In most situations, you should not configure the router to perform unrestricted proxy ARP. Do so only for special situations, such as when cable modems are used. Figure 61 on page 672 and Figure 62 on page 672 show examples of situations in which you might want to configure unrestricted proxy ARP.

In Figure 61 on page 672, the edge device is not running any IP protocols. In this case, you configure the core router to perform unrestricted proxy ARP. The edge device is the client of the proxy.

In Figure 62 on page 672, the Broadcast Remote Access Server (B-RAS) routers are not running any IP protocols. In this case, you configure unrestricted proxy ARP on the B-RAS interfaces. This allows the core device to behave as though it is directly connected to the end users.

Figure 61: Edge Device Case for Unrestricted Proxy ARP

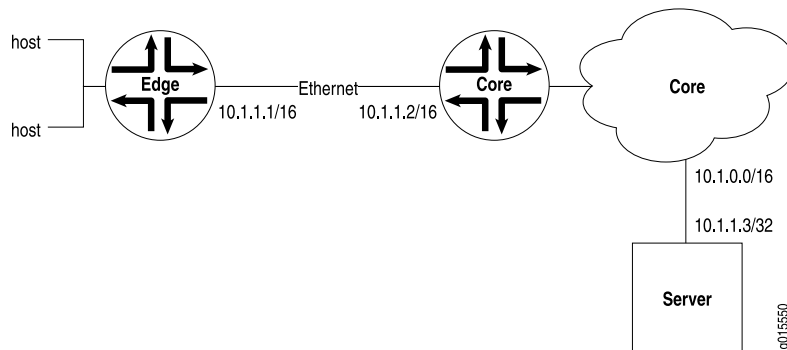



Figure 62: Core Device Case for Unrestricted Proxy ARP



Configuring Unrestricted Proxy ARP

To configure unrestricted proxy ARP, include the `proxy-arp` statement:

```
proxy-arp;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]

- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

To return to the default—that is, to disable unrestricted proxy ARP—delete the `proxy-arp` statement from the configuration:

```
[edit]
user@host# delete interfaces interface-name unit logical-unit-number proxy-arp
```

You can track the number of unrestricted proxy ARP requests processed by the router by issuing the `show system statistics arp operational mode` command.

Chapter 41

Configuring MAC Address Validation on Static Ethernet Interfaces

This section contains the following topics:

- MAC Address Validation on Static Ethernet Interfaces Overview on page 675
- Configuring MAC Address Validation on Static Ethernet Interfaces on page 675

MAC Address Validation on Static Ethernet Interfaces Overview

MAC address validation enables the router to validate that received packets contain a trusted IP source and an Ethernet MAC source address.

MAC address validation is supported on AE, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces (with or without VLAN tagging) on MX Series routers only.

There are two types of MAC address validation that you can configure:

- Loose—Forwards packets when both the IP source address and the MAC source address match one of the trusted address tuples.

Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not support the MAC address of the tuple

Continues to forward packets when the source address of the incoming packet does not match any of the trusted IP addresses.

- Strict—Forwards packets when both the IP source address and the MAC source address match one of the trusted address tuples.

Drops packets when the MAC address does not match the tuple's MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.

Configuring MAC Address Validation on Static Ethernet Interfaces

To configure MAC address validation on static Ethernet interfaces, include the `mac-validate (loose | strict)` statement in the `[edit interfaces interface-name unit logical-unit-number family family]` hierarchy:

```
[edit interfaces interface-name unit logical-unit-number family family]  
mac-validate (loose | strict);
```

Example of Strict MAC Validation on a Static Ethernet Interface

This example shows strict MAC address validation on a static Ethernet interface without VLAN tagging.

```
[edit interfaces]  
ge-2/1/9 {  
  unit 0 {  
    proxy-arp;  
    family inet {  
      mac-validate strict;  
      address 88.22.100.1/24 {  
        arp 88.22.100.3 mac 00:00:58:16:64:03;  
      }  
    }  
  }  
}
```

Chapter 42

Enabling Passive Monitoring on Ethernet Interfaces

This section contains the following topics:

- Passive Monitoring on Ethernet Interfaces Overview on page 677
- Enabling Passive Monitoring on Ethernet Interfaces on page 677

Passive Monitoring on Ethernet Interfaces Overview

The Monitoring Services I and Monitoring Services II PICs are designed to enable IP services. If you have a Monitoring Services PIC and a 4-port Fast Ethernet, 4-port Gigabit Ethernet PIC with SFPs, 10-port Gigabit Ethernet PIC with SFPs, or 1-port 10-Gigabit Ethernet PIC installed in an M40e, M160, MX Series, or T Series router, you can monitor IP version 4 (IPv4) traffic from another router.

Enabling Passive Monitoring on Ethernet Interfaces

On Ethernet interfaces, enable packet flow monitoring by including the `passive-monitor-mode` statement at the `[edit interfaces interface-name]` hierarchy level:

```
[edit interfaces interface-name]  
passive-monitor-mode;
```

When you configure an interface in passive monitoring mode, the Packet Forwarding Engine silently drops packets coming from that interface and destined to the router itself. Passive monitoring mode also stops the Routing Engine from transmitting any packet from that interface. Packets received from the monitored interface can be forwarded to monitoring interfaces. If you include the `passive-monitor-mode` statement in the configuration:

- Gigabit and Fast Ethernet interfaces can support both per-port passive monitoring and per-VLAN passive monitoring. The destination MAC filter on the receive port of the Ethernet interfaces is disabled.
- Ethernet encapsulation options are not allowed.

On monitoring services interfaces, enable packet flow monitoring by including the `family` statement at the `[edit interfaces mo-fpc/pic/port unit logical-unit-number]` hierarchy level, specifying the `inet` option:

```
[edit interfaces mo-fpc/pic/port unit logical-unit-number]  
family inet;
```

For conformity with cflowd record structure, you must include the `receive-options-packets` and `receive-ttl-exceeded` statements at the `[edit interfaces mo-fpc/pic/port unit logical-unit-number family inet]` hierarchy level:

```
[edit interfaces mo-fpc/pic/port unit logical-unit-number family inet]  
receive-options-packets;  
receive-ttl-exceeded;
```

For the monitoring services interface, you can configure multiservice physical interface properties. For more information, see “Configuring Multiservice Physical Interface Properties” on page 138 and the *JUNOS Services Interfaces Configuration Guide*.

Chapter 43

Configuring IEEE 802.1ag OAM Connectivity-Fault Management

This section contains the following topics:

- IEEE 802.1ag OAM Connectivity Fault Management Overview on page 680
- Creating the Maintenance Domain on page 682
- Configuring Maintenance Intermediate Points on page 683
- Creating the Maintenance Association on page 684
- Configuring the Maintenance Association Short Name Format on page 685
- Configuring the Continuity Check on page 685
- Configuring the Continuity Check Hold Interval on page 685
- Configuring the Continuity Check Interval on page 686
- Configuring the Continuity Check Loss Threshold on page 686
- Configuring a Maintenance End Point on page 686
- Enabling Maintenance End Point Automatic Discovery on page 686
- Configuring the Maintenance End Point Direction on page 686
- Configuring the Maintenance End Point Interface on page 687
- Configuring the Maintenance End Point Priority on page 687
- Configuring a Remote Maintenance End Point on page 688
- Configuring a Remote Maintenance End Point Action Profile on page 688
- Configuring a Connectivity-Fault Management Action Profile on page 688
- Configuring a CFM Action Profile Action on page 688
- Configuring a CFM Interface Down Action Profile Action on page 688
- Configuring the Linktrace Path Age Timer on page 689
- Configuring the Linktrace Database Size on page 690
- Configuring Ethernet Local Management Interface on page 690
- Configuring Port Status TLV and Interface Status TLV on page 697
- Configuring M120 and MX Series Routers for CCC Encapsulated Packets on page 708

IEEE 802.1ag OAM Connectivity Fault Management Overview

Ethernet interfaces on M120, M320, MX Series, and T Series routers support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM). The IEEE 802.1ag specification provides for Ethernet connectivity-fault management (CFM). The goal of CFM is to monitor an Ethernet network that may comprise one or more service instances. JUNOS Software supports IEEE 802.1ag connectivity fault management.

Network entities such as operators, providers, and customers may be part of different administrative domains. Each administrative domain is mapped into one maintenance domain. Maintenance domains are configured with different level values to keep them separate. Each domain provides enough information for the entities to perform their own management, perform end-to-end monitoring, and still avoid security breaches.



NOTE: As a requirement for Ethernet OAM 802.1ag to work, distributed periodic packet management (PPM) runs on the Routing Engine and Packet Forwarding Engine (PFE) by default. You can only disable PPM on the PFE. To disable PPM on the PFE, include the `ppm <no-delegate-processing>` statement at the `[edit routing-options ppm]` hierarchy level.

IEEE 802.1ag OAM supports graceful Routing Engine switchover (GRES). IEEE 802.1ag OAM is supported on untagged, single tagged, and stacked VLAN interfaces.

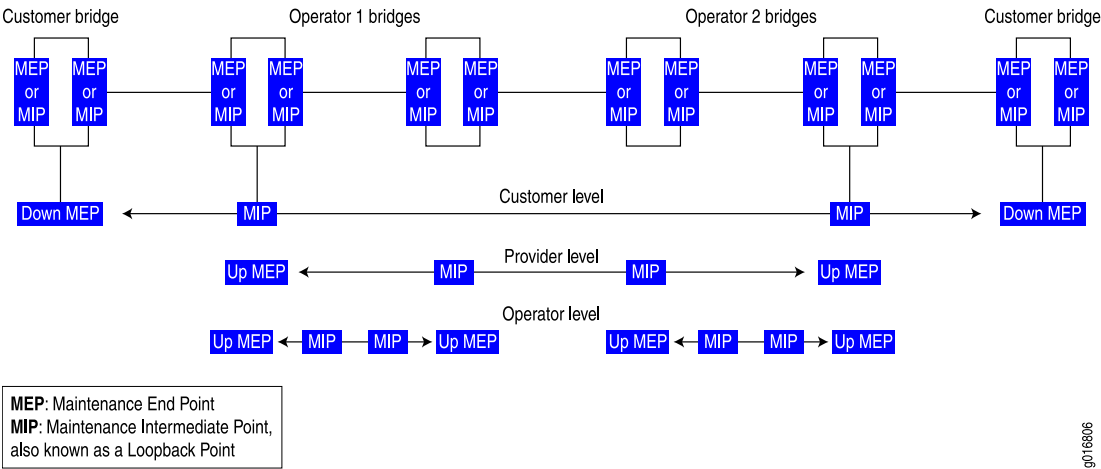
This section includes the following topics:

- Connectivity Fault Management Key Elements on page 680

Connectivity Fault Management Key Elements


Figure 63 on page 681 shows the relationships among the customer, provider, and operator Ethernet bridges, maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs).

Figure 63: Relationship Among MEPs, MIPs, and Maintenance Domain Levels



A maintenance association is a set of MEPs configured with the same maintenance association identifier and maintenance domain level. Figure 64 on page 681 shows the hierarchical relationships between the Ethernet bridge, maintenance domains, maintenance associations, and MEPs.

Figure 64: Relationship Among Bridges, Maintenance Domains, Maintenance Associations, and MEPs



Continuity Check Protocol

The continuity check protocol is used for fault detection by a MEP within a maintenance association. The MEP periodically sends continuity check multicast messages. The receiving MEPs use the continuity check messages to build a MEP database of all MEPs in the maintenance association.

The continuity check protocol packets use the ethertype value 0x8902 and the multicast destination MAC address 01:80:c2:00:00:32.

Linktrace Protocol

The linktrace protocol is used for path discovery between a pair of maintenance points. Linktrace messages are triggered by an administrator using the **traceroute** command to verify the path between a pair of MEPs under the same maintenance association. Linktrace messages can also be used to verify the path between an MEP and an MIP under the same maintenance domain. The operation of IEEE 802.1ag linktrace request and response messages is similar to the operation of Layer 3 **traceroute** commands. For more information about the **traceroute** command, see the *JUNOS System Basics Configuration Guide*.

Creating the Maintenance Domain

To enable CFM on an Ethernet interface, maintenance domains, maintenance associations, and MEPs must be created and configured.

To create a maintenance domain, include the **maintenance-domain** *domain-name* statement at the [edit protocols oam ethernet connectivity-fault-management] hierarchy level.

Give the maintenance domain a name. Names can be in one of several formats:

- Configuring the Maintenance Domain Name Format on page 682
- Configuring the Maintenance Domain Level on page 682
- Configuring MIP for Bridge Domains of a Virtual Switch on page 683

Configuring the Maintenance Domain Name Format

You can specify the maintenance domain name format as one of the following:

- A plain ASCII character string.
- A domain name service (DNS) format, a MAC address plus a two-octet identifier in the range from 0 through 65,535, or none.
- A MAC address plus a two-octet identifier in the range from 0 through 65,535.
- Or none.

If none is specified, the maintenance domain name is not used.

The default name format is an ASCII character string.

To configure the maintenance domain name format, include the **name-format** (*character-string* | none | dns | mac+2octet) statement at the [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name*] hierarchy level.

Configuring the Maintenance Domain Level

The maintenance domain level is a mandatory parameter that indicates the nesting relationship between various maintenance domains. The level is embedded in each

of the CFM frames. CFM messages within a given level are processed by MEPs at that same level. For example, the operator domain can be level 0, the provider domain can be level 3, and the customer domain can be level 7.

To configure the maintenance domain level, include the `level number` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name]` hierarchy level.

Configuring MIP for Bridge Domains of a Virtual Switch

The default maintenance domain configuration allows MIP configuration for bridge domains for a default virtual switch or user-defined virtual switch. You can use the `virtual-switch` statement with the `bridge-domain` substatement to specify which MIPs to enable.

A bridge domain must be specified by name only if it is configured with `vlan-id id` under the `virtual-switch` statement.

If a bridge domain is configured with a range of VLAN IDs, then the VLAN IDs must be explicitly listed after the bridge domain name.

To configure a bridge domain under a user-defined virtual switch, include the `virtual-switch name` statement and `bridge-domain` substatement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name default-x]` hierarchy level.

```
virtual-switch name {
  bridge-domain {
    name-1;
    name-2 {
      vlan-id [id1 id2 ... idn];
    }
  }
}
```

Configuring Maintenance Intermediate Points

MX Series routers support maintenance intermediate points (MIPs) for the Ethernet OAM 802.1ag CFM protocol at a bridge-domain level. This enables you to define a maintenance domain for each default level. The MIPs names are created as `default-level-number` at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain]` hierarchy level. Use the `bridge-domain`, `routing-instance`, `domain-instance`, and `mip-half-function` MIP options to specify the MIP configuration.



NOTE: Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the `mip-half-function` value for all maintenance domains and maintenance associations be the same.

To display MIP configurations, use the `show oam ethernet connectivity-fault-management mip (bridge-domain | instance-name | interface-name)` command.

The following sections describe MIP configuration:

- Configuring the Maintenance Domain Routing Instances Bridge Domain on page 684
- Configuring the Maintenance Domain Routing Instance on page 684
- Configuring the Maintenance Domain Instance on page 684
- Configuring the Maintenance Domain MIP Half Function on page 684

Configuring the Maintenance Domain Routing Instances Bridge Domain

The VLAN corresponds to the bridge domain.

To configure the bridge domain, include the `bridge-domain` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain routing-instances name]` hierarchy level.

Configuring the Maintenance Domain Routing Instance

To configure the routing instance, include the `routing-instance` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain]` hierarchy level.

Configuring the Maintenance Domain Instance

To configure the maintenance domain instance, include the `instance` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain]` hierarchy level.

Configuring the Maintenance Domain MIP Half Function

MIP Half Function (MHF) divides MIP functionality into two unidirectional segments, improves visibility with minimal configuration, and improves network coverage by increasing the number of points that can be monitored. MHF extends monitoring capability by responding to loop-back and link-trace messages to help isolate faults.

Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains or maintenance associations, it is essential that the *MIP half function* value for all maintenance domains and maintenance associations be the same. To configure the MIP half function, include the `mip-half-function` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain]` hierarchy level.

Creating the Maintenance Association

To create a maintenance association, include the `maintenance-association` *ma-name* statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name]` hierarchy level.

Configuring the Maintenance Association Short Name Format

Maintenance association names can be in one of the following formats:

- As a plain ASCII character string
- As the VLAN identifier of the VLAN you primarily associate with the maintenance association
- As a two-octet identifier in the range from 0 through 65,535
- As a name in the format specified by RFC 2685

The default short name format is an ASCII character string.

To configure the maintenance association short name format, include the `short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id)` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name]` hierarchy level.

Configuring the Continuity Check

You can configure the following continuity check protocol parameters:

- `hold-interval minutes`
- `interval (time)`
- `loss-threshold number`

To enable the continuity check protocol, include the `continuity-check` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name]` hierarchy level.

Related Topics See Configuring the Continuity Check Hold Interval on page 685, Configuring the Continuity Check Interval on page 686, and Configuring the Continuity Check Loss Threshold on page 686.

Configuring the Continuity Check Hold Interval

You can specify the continuity check hold interval. The hold interval is the number of minutes to wait before flushing the MEP database if no updates occur. The default value is 10 minutes.

To configure the hold interval, include the `hold-interval` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name continuity-check]` hierarchy level.

Configuring the Continuity Check Interval

You can specify the continuity check message (CCM) interval. The interval is the time between the transmission of CCMs. You can specify 10 minutes (**10m**), 1 minute (**1m**), 10 seconds (**10s**), 1 second (**1s**), 100 milliseconds (**100ms**), or 10 milliseconds (**10ms**). The default value is 1 minute.



NOTE: For the continuity check message interval to be configured for 10 milliseconds, periodic packet management (PPM) runs on the Routing Engine and Packet Forwarding Engine (PFE) by default. You can only disable PPM on the PFE. To disable PPM on the PFE, use the `no-delegate-processing` statement at the `[edit routing-options ppm]` hierarchy level.

To configure the interval, include the `interval` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name continuity-check]` hierarchy level.

Configuring the Continuity Check Loss Threshold

You can specify the number of continuity check messages that can be lost before marking the MEP as down. The default value is three (PDUs).

To configure the loss threshold, include the `loss-threshold` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name continuity-check]` hierarchy level.

Configuring a Maintenance End Point

To configure the maintenance end point, include the `mep mep-id` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name]` hierarchy level.

Enabling Maintenance End Point Automatic Discovery

You can enable the MEP to accept continuity check messages from all remote MEPs of the same maintenance association.

To configure automatic discovery, include the `auto-discovery` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id]` hierarchy level.

Configuring the Maintenance End Point Direction

You can specify the direction in which CFM packets are transmitted for the MEP.

Direction up continuity check messages (CCMs) are transmitted out of every logical interface which is part of the same bridging or VPLS instance except for the interface configured on this MEP.

Direction down CCMs are transmitted only out of the interface configured on this MEP.



NOTE: Ports in the spanning tree protocol (STP) blocking state do not block CFM packets destined to a down MEP. Ports in STP blocking state without the continuity check protocol configured, do block CFM packets.

To configure the MEP direction, include the `direction` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id]` hierarchy level.

Configuring the Maintenance End Point Interface

You must specify the interface to which the MEP is attached. It can be a physical interface, logical interface or trunk interface.

On MX Series routers, you can enable the MEP on a specific VLAN of a trunk interface.

To configure the interface, include the `interface interface-name` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id]` hierarchy level.

MEP Interface Configuration

This example shows the MEP interface configuration statements:

```
[edit protocols oam ethernet connectivity-fault-management maintenance-domain
  domain-name maintenance-association ma-name]
mep mep-id {
  direction (up | down);
  interface (ge | xe)-(fpc/pic/port | fpc/pic/port.domain | fpc/pic/port.domain vlan
    vlan-id);
  auto-discovery;
  priority number;
}
```

Configuring the Maintenance End Point Priority

You can specify the IEEE 802.1 priority bits that are used by continuity check and link trace messages.

To configure the priority, include the `priority` statement at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id]` hierarchy level.

Configuring a Remote Maintenance End Point

You can configure a remote MEP from which CCM messages are expected. If autodiscovery is not enabled, the remote MEP must be configured under the **mep** statement. If the remote MEP is not configured under the **mep** statement, the CCMs from the remote MEP are treated as errors.

To configure the remote MEP, include the **remote-mep** statement at the [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id*] hierarchy level.

Configuring a Remote Maintenance End Point Action Profile

You can specify the name of the action profile to use for the remote MEP.

To configure the action profile, include the **action-profile** *profile-name* statement at the [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id* remote-mep *mep-id*] hierarchy level. The profile must already be defined at the [edit protocols oam ethernet connectivity-fault-management] hierarchy level.

Configuring a Connectivity-Fault Management Action Profile

You can configure an action profile and specify the action to be taken when connectivity to a remote MEP fails.

To configure the action profile name, include the **action-profile** statement at the [edit protocols oam ethernet connectivity-fault-management] hierarchy level.

Configuring a CFM Action Profile Action

You can configure the action to be taken when connectivity to a remote MEP fails.

To configure the action profile action, include the **default-action** statement at the [edit protocols oam ethernet connectivity-fault-management maintenance-domain *domain-name* maintenance-association *ma-name* mep *mep-id* remote-mep *mep-id* action-profile] hierarchy level.



NOTE: The action profile is supported only on the physical interface level, and not on the logical interface level.

Configuring a CFM Interface Down Action Profile Action

You can configure the action to be taken when connectivity to a remote MEP fails and the interface goes down.

To enable the interface down action, include the `interface-down` statement at the `[edit protocols oam ethernet connectivity-fault-management action-profile profile-name default-action]` hierarchy level.

```
[edit]
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        action-profile bring-down {
          default-actions {
            interface-down;
          }
        }
      }
      maintenance-domain md1 {
        level 0;
        maintenance-association ma1 {
          continuity-check {
            interval 100 ms;
          }
          mep 4001 {
            interface ge-4/1/0;
            direction down;
            remote-mep 1 {
              action-profile bring-down;
            }
          }
        }
      }
    }
  }
}
```

The action profile is supported on the physical interface level and the logical interface level.



NOTE: Associating an action-profile with the action of `interface-down` on an up MEP CFM session which is running over a CCC interface (l2circuit/l2vpn) is not advisable and could result in a deadlock situation.

Configuring the Linktrace Path Age Timer

If no response to a `linktrace` request is received, the request and response entries are deleted after the age timer expires. To configure the linktrace age time use the `linktrace` statement `age time` option at the `[edit protocols oam ethernet connectivity-fault-management]` hierarchy level. The age is configured in minutes or seconds.

Related Topics See Linktrace Protocol on page 682 and `linktrace`.

Configuring the Linktrace Database Size

Configure the number of linktrace reply entries to be stored per linktrace request. To configure the linktrace database size use the `linktrace` statement `path-database-size` option at the `[edit protocols oam ethernet connectivity-fault-management]` hierarchy level.

The linktrace database is displayed using the `show oam ethernet connectivity-fault-management path-database` command.

Related Topics See Linktrace Protocol on page 682 and linktrace.

Configuring Ethernet Local Management Interface

This section contains the following topics:


- Ethernet Local Management Interface Overview on page 690
- Configuring the Ethernet Local Management Interface on page 692
- Example E-LMI Configuration on page 693

Ethernet Local Management Interface Overview

MX Series routers with Gigabit Ethernet (ge), 10-Gigabit Ethernet (xe), or Aggregated Ethernet (ae) interfaces support the Ethernet Local Management Interface (E-LMI). The E-LMI specification is available at the Metro Ethernet Forum. E-LMI procedures and protocols are used for enabling automatic configuration of the customer edge (CE) to support Metro Ethernet services. The E-LMI protocol also provides user-to-network interface (UNI) and Ethernet virtual connection (EVC) status information to the CE. The UNI and EVC information enables automatic configuration of CE operation based on the Metro Ethernet configuration.

The E-LMI protocol operates between the CE device and the provider edge (PE) device. It runs only on the PE-CE link and notifies the CE of connectivity status and configuration parameters of Ethernet services available on the CE port. The scope of the E-LMI protocol is shown in Figure 65 on page 690.

Figure 65: Scope of the E-LMI Protocol



The E-LMI implementation on MX Series routers includes only the PE side of the E-LMI protocol.

E-LMI interoperates with an OAM protocol, such as Connectivity Fault Management (CFM), that runs within the provider network to collect OAM status. CFM runs at the provider maintenance level (UNI-N to UNI-N with up MEPs at the UNI). E-LMI relies on the CFM for end-to-end status of EVCs across CFM domains (SVLAN domain or VPLS).

The E-LMI protocol relays the following information:

- Notification to the CE of the addition/deletion of an EVC (active, not active, or partially active)
- Notification to the CE of the availability state of a configured EVC
- Communication of UNI and EVC attributes to the CE:
 - UNI attributes:
 - UNI identifier (a user-configured name for UNI)
 - CE-VLAN ID/EVC map type (all-to-one bundling, service multiplexing with bundling, or no bundling)
 - Bandwidth profile is not supported (including the following features):
 - CM (coupling mode)
 - CF (color flag)
 - CIR (committed Information rate)
 - CBR (committed burst size)
 - EIR (excess information rate)
 - EBS (excess burst size)
 - EVC attributes:
 - EVC reference ID
 - EVC status type (active, not active, or partially active)
 - EVC type (point-to-point or multipoint-to-multipoint)
 - EVC ID (a user-configured name for EVC)
 - Bandwidth profile (not supported)
 - CE-VLAN ID/EVC map

E-LMI on MX Series routers supports the following EVC types:

- QinQ SVLAN (point-to-point or multipoint-to-multipoint)—Requires an end-to-end CFM session between UNI-Ns to monitor the EVS status.
- VPLS (BGP or LDP) (point-to-point or multipoint-to-multipoint)—Either VPLS pseudowire status or end-to-end CFM sessions between UNI-Ns can be used to monitor EVC status.
- L2 circuit/L2VPN (point-to-point)—Either VPLS pseudowire status or end-to-end CFM sessions between UNI-Ns can be used to monitor EVC status.



NOTE: I2-circuit and I2vpn are not supported.

Configuring the Ethernet Local Management Interface

To configure E-LMI, perform the following steps:

- Configuring an OAM Protocol (CFM) on page 692
- Assigning the OAM Protocol to an EVC on page 692
- Enabling E-LMI on an Interface and Mapping CE VLAN IDs to an EVC on page 692

Configuring an OAM Protocol (CFM)

For information on configuring the OAM protocol (CFM), see “Configuring IEEE 802.1ag OAM Connectivity-Fault Management” on page 679.

Assigning the OAM Protocol to an EVC

To configure an EVC, you must specify a name for the EVC using the `evcsevc-id` statement at the `[edit protocols oam ethernet]` hierarchy level. You can set the EVC protocol for monitoring EVC statistics to `cfm` or `vpls` using the `evc-protocol` statement and its options at the `[edit protocols oam ethernet evcs]` hierarchy level.

You can set the number of remote UNIs in the EVC using the `remote-uni-count` *number* statement at the `[edit protocols oam ethernet evcs evcs-protocol]` hierarchy level. The `remote-uni-count` defaults to 1. Configuring a value greater than 1 makes the EVC multipoint-to-multipoint. If you enter a value greater than the actual number of endpoints, the EVC status will display as partially active even if all endpoints are up. If you enter a `remote-uni-count` less than the actual number of endpoints, the status will display as active, even if all endpoints are not up.

Example: Assigning OAM protocol to an EVC

You can configure an EVC using the `evcs` statement at the `[edit protocols oam ethernet]` hierarchy level:

```
[edit protocols oam ethernet]
evcs evc-id {
  evc-protocol (<cfm (<management-domain name> <management-association name> )
|<< vpls (routing-instance name>>>)) {
    remote-uni-count <number>; # Optional, defaults to 1
    multipoint-to-multipoint;
    # Optional, defaults to point-to-point if remote-uni-count is 1
  }
}
```

Enabling E-LMI on an Interface and Mapping CE VLAN IDs to an EVC

To configure E-LMI, include the `lmi` statement at the `[edit protocols oam ethernet]` hierarchy level:

```
[edit protocols oam ethernet]
```

```

lmi (Ethernet OAM) {
    polling-verification-timer value;
    # Polling verification timer (T392), defaults to 15 seconds
    status-counter count; # Status counter (N393), defaults to 4
    interface name {
        evc evc-id {
            default-evc;
            vlan-list [ vlan-ids ];
        }
        evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
        polling-verification-time value; # Optional, defaults to global value
        status-counter count; # Optional, defaults to global value
        uni-id value; # Optional, defaults to interface-name
    }
}

```

You can set the status counter to count consecutive errors using the **status-counter *count*** statement at the [edit protocols oam ethernet lmi] hierarchy level. The status counter is used to determine if E-LMI is operational or not. The default value is 4.

You can set the **polling-verification-timer *value*** statement at the [edit protocols oam ethernet lmi] hierarchy level. The default value is 15 seconds.

You can enable an interface and set its options for use with E-LMI using the **interface *name*** statement at the [edit protocols oam ethernet lmi] hierarchy level. Only **ge**, **xe**, and **ae** interfaces are supported. You can use the interface **uni-id** option to specify a name for the UNI. If **uni-id** is not configured, it defaults to the name variable of **interface *name***.

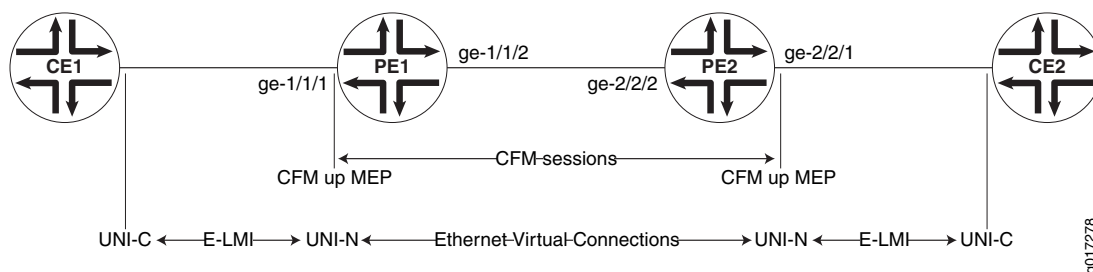
You can specify the CE-VLAN ID/EVC map type using the **evc-map-type *type*** interface option. The options are **all-to-one-bundling**, **bundling**, or **service-multiplexing**. Service multiplexing is with no bundling. The default type is **all-to-one-bundling**.

To specify the EVC that an interface uses, use the **evc *evc-id*** statement at the [edit protocols oam ethernet lmi interface *name*] hierarchy level. You can specify an interface as the default EVC interface using the **default-evc** statement at the [edit protocols oam ethernet lmi interface *name* evc *evc-id*] hierarchy level. All VLANs that are not mapped to any other EVCs are mapped to this EVC. Only one EVC can be configured as the default.

You can map a list of VLANs to an EVC using the **vlan-list *vlan-id-list*** statement at the [edit protocols oam ethernet lmi interface *name* evc *evc-id*] hierarchy level.

Example E-LMI Configuration

Figure 66 on page 694 illustrates the E-LMI configuration for a point-to-point EVC (SVLAN) monitored by CFM. In this example, VLANs 1 through 2048 are mapped to **evc1** (SVLAN 100) and 2049 through 4096 are mapped to **evc2** (SVLAN 200). Two CFM sessions are created to monitor these EVCs.

Figure 66: E-LMI Configuration for a Point-to-Point EVC (SVLAN) Monitored by CFM**Configuring PE1**

```
[edit]
interfaces {
  ge-1/1/1 {
    unit 0 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 1-2048;
      }
    }
    unit 1 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 2049-4096;
      }
    }
  }
  ge-1/1/2 {
    unit 0 {
      vlan-id 100;
      family bridge {
        interface-mode trunk;
        inner-vlan-id-list 1-2048;
      }
    }
    unit 1 {
      vlan-id 200;
      family bridge {
        interface-mode trunk;
        inner-vlan-id-list 2049-4096;
      }
    }
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain md {
          level 0;
          maintenance-association 1 {
            name-format vlan;
          }
        }
      }
    }
  }
}
```

```
mep 1 {  
    direction up;  
    interface ge-1/1/1.0 vlan 1;  
}  
  
maintenance-association 2049 {  
    name-format vlan;  
    mep 1 {  
        direction up;  
        interface ge-1/1/1.1 vlan 2049;  
    }  
}  
  
} }  
  
evcs {  
    evc1 {  
        evc-protocol cfm management-domain md management-association 1;  
        remote-uni-count 1;  
    }  
    evc2 {  
        evc-protocol cfm management-domain md management-association 2049;  
        remote-uni-count 1;  
    }  
}  
  
lmi {  
    interface ge-1/1/1 {  
        evc evc1 {  
            vlan-list 1-2048;  
        }  
        evc evc2 {  
            vlan-list 2049-4096;  
        }  
        evc-map-type bundling;  
        uni-id uni-ce1;  
    }  
}  
  
} }  
  
}
```

Configuring PE2

```
[edit]
interfaces {
  ge-2/2/1 {
    unit 0 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 1-2048;
      }
    }
    unit 1 {
      family bridge {
        interface-mode trunk;
        vlan-id-list 2049-4096;
      }
    }
  }
}
```

```

    }
  }
}
ge-2/2/2 {
  unit 0 {
    vlan-id 100;
    family bridge {
      interface-mode trunk;
      inner-vlan-id-list 1-2048;
    }
  }
  unit 1 {
    vlan-id 200;
    family bridge {
      interface-mode trunk;
      inner-vlan-id-list 2049-4095;
    }
  }
}
}
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain md {
          level 0;
          maintenance-association 1 {
            name-format vlan;
            mep 1 {
              direction up;
              interface ge-2/2/1.0 vlan 1;
            }
          }
          maintenance-association 2049 {
            name-format vlan;
            mep 1 {
              direction up;
              interface ge-2/2/1.1 vlan 2049;
            }
          }
        }
      }
    }
  }
  evcs {
    evc1 {
      evc-protocol cfm management-domain md management-association 1;
      remote-uni-count 1;
    }
    evc2 {
      evc-protocol cfm management-domain md management-association 2049;
      uni-count 2;
    }
  }
}
lmi {
  interface ge-2/2/1 {
    evc evc1 {
      vlan-list 1-2048;
    }
  }
}

```



```

    }
    evc evc2 {
        vlan-list 2049-4095;
    }
    evc-map-type bundling;
    uni-id uni-ce2;
}
}
}
}
}
}
}

```

Configuring Two UNIs Sharing the Same EVC

```

[edit protocols]
oam {
  ethernet {
    connectivity-fault-management { ...}
    evcs {
      evc1 {
        evc-protocol cfm management-domain md management-association 1;
        remote-uni-count 1;
      }
    }
    lmi {
      interface ge-2/2/1 {
        evc evc1 {
          vlan-list 0-4095;
        }
        evc-map-type all-to-one-bundling;
        uni-id uni-ce1;
      }
      interface ge-2/3/1 {
        evc evc1 {
          vlan-list 0-4095;
        }
        evc-map-type all-to-one-bundling;
        uni-id uni-ce2;
      }
    }
  }
}
}
}
}
}
}
}

```

Configuring Port Status TLV and Interface Status TLV

This section contains the following topics:

- Overview of TLVs on page 698
- Various TLVs for CFM PDUs on page 698
- Support for Additional Optional TLVs on page 700
- MAC Status Defects on page 705
- Configuring Remote MEP Action Profile Support on page 707

Overview of TLVs

Type, Length, and Value (TLVs) are described in the IEEE 802.1ag standard for CFM as a method of encoding variable-length and/or optional information in a PDU. TLVs are not aligned to any particular word or octet boundary. TLVs follow each other with no padding between them.

Table 59 on page 698 shows the TLV format and indicates if it is required or optional.

Table 59: Format of TLVs

Parameter	Octet (sequence)	Description
Type	1	Required. If 0, no Length or Value fields follow. If not 0, at least the Length field follows the Type field.
Length	2–3	Required if the Type field is not 0. Not present if the Type field is 0. The 16 bits of the Length field indicate the size, in octets, of the Value field. 0 in the Length field indicates that there is no Value field.
Value	4	Length specified by the Length field. Optional. Not present if the Type field is 0 or if the Length field is 0.

Various TLVs for CFM PDUs

Table 60 on page 698 shows a set of TLVs defined by IEEE 802.1ag for various CFM PDU types. Each TLV can be identified by the unique value assigned to its type field. Some type field values are reserved.

Table 60: Type Field Values for Various TLVs for CFM PDUs

TLV or Organization	Type Field
End TLV	0
Sender ID TLV	1
Port Status TLV	2
Data TLV	3
Interface Status TLV	4
Reply Ingress TLV	5
Reply Egress TLV	6
LTM Egress Identifier TLV	7
LTR Egress Identifier TLV	8
Reserved for IEEE 802.1	9 to 30

Table 60: Type Field Values for Various TLVs for CFM PDUs *(continued)*

TLV or Organization	Type Field
Organization-Specific TLV	31
Defined by ITU-T Y.1731	32 to 63
Reserved for IEEE 802.1	64 to 255

Not every TLV is applicable for all types of CFM PDUs.

- TLVs applicable for continuity check message (CCM):
 - End TLV
 - Sender ID TLV
 - Port Status TLV
 - Interface Status TLV
 - Organization-Specific TLV
- TLVs applicable for loopback message (LBM):
 - End TLV
 - Sender ID TLV
 - Data TLV
 - Organization-Specific TLV
- TLVs applicable for loopback reply (LBR):
 - End TLV
 - Sender ID TLV
 - Data TLV
 - Organization-Specific TLV
- TLVs applicable for linktrace message (LTM):
 - End TLV
 - LTM Egress Identifier TLV
 - Sender ID TLV
 - Organization-Specific TLV
- TLVs applicable for linktrace reply (LTR):
 - End TLV
 - LTR Egress Identifier TLV
 - Reply Ingress TLV

- Reply Egress TLV
- Sender ID TLV
- Organization-Specific TLV

The following TLVs are currently supported in the applicable CFM PDUs:

- End TLV
- Reply Ingress TLV
- Reply Egress TLV
- LTR Egress Identifier TLV
- LTM Egress Identifier TLV
- Data TLV

Support for Additional Optional TLVs

The following additional optional TLVs are supported:

- Port Status TLV
- Interface Status TLV

MX Series routers support configuration of port status TLV and interface status TLV. Configuring the Port Status TLV allows the operator to control the transmission of the Port Status TLV in CFM PDUs.



NOTE: Although Port Status TLV configuration statements are visible in the CLI on M120 and M320 routers, Port Status TLV cannot be configured on these systems. Port Status TLV can be enabled on a MEP interface only if it is a bridge logical interface, which is not possible on these systems.

Port Status TLV

The Port Status TLV indicates the ability of the bridge port on which the transmitting MEP resides to pass ordinary data, regardless of the status of the MAC. The value of this TLV is driven by the MEP variable `enableRmepDefect`, as shown in Table 62 on page 701. The format of this TLV is shown in Table 61 on page 700.

Any change in the Port Status TLVs value triggers one extra transmission of that bridge ports MEP CCMs.

Table 61: Port Status TLV Format

Parameter	Octet (Sequence)
Type = 2	1

Table 61: Port Status TLV Format (*continued*)

Parameter	Octet (Sequence)
Length	2–3
Value (See Table 62 on page 701)	4

Table 62: Port Status TLV Values

Mnemonic	Ordinary Data Passing Freely Through the Port	Value
psBlocked	No: enableRmepDefect = false	1
psUp	Yes: enableRmepDefect = true	2

The MEP variable `enableRmepDefect` is a boolean variable indicating whether frames on the service instance monitored by the maintenance associations if this MEP are enabled to pass through this bridge port by the Spanning Tree Protocol and VLAN topology management. It is set to TRUE if:

- The bridge port is set in a state where the traffic can pass through it.
- The bridge port is running multiple instances of the spanning tree.
- The MEP interface is not associated with a bridging domain.

Configuring Port Status TLV

JUNOS Software provides configuration support for the Port Status TLV, allowing you to control the transmission of this TLV in CCM PDUs. The JUNOS Software provides this configuration at the continuity-check level. By default, the CCM does not include the Port Status TLV. To configure the Port Status TLV, use the `port-status-tlv` statement at the [edit protocols oam ethernet connectivity-fault-management maintenance-domain *identifier* maintenance-association *identifier* continuity-check] hierarchy level.



NOTE: Port Status TLV configuration is not mandated by IEEE 802.1ag. The JUNOS Software provides it in order to give more flexibility to the operator; however it receives and processes CCMs with a Port Status TLV, regardless of this configuration.

An example of the configuration statements follows:

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain identifier {
          level number;
          maintenance-association identifier {
            continuity-check {
              interval number,

```

```

        loss-threshold number;
        hold-interval number;
        port-status-tlv; # Sets Port Status TLV
    }
}
}
}
}
}
}
}

```

You cannot enable Port Status TLV transmission in the following two cases:

- If the MEP interface under the maintenance-association is not of type bridge.
- If the MEP is configured on a physical interface.

Displaying the Received Port Status TLV

The JUNOS Software saves the last received Port Status TLV from a remote MEP. If the received Port Status value does not correspond to one of the standard values listed in Table 62 on page 701, then the **show** command displays it as "unknown." You can display the last saved received Port Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up
Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

Remote MEP identifier: 1001, State: ok
MAC address: 00:19:e2:b0:74:00, Type: Learned
Interface: ge-2/0/0.0
Last flapped: Never
Remote defect indication: false
Port status TLV: none # RX PORT STATUS
Interface status TLV: none

```

Displaying the Transmitted Port Status TLV

The JUNOS Software saves the last transmitted Port Status TLV from a local MEP. If the transmission of the Port Status TLV has not been enabled, then the **show** command displays "none." You can display the last saved transmitted Port Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up # TX PORT STATUS
Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

Remote MEP identifier: 1001, State: ok
MAC address: 00:19:e2:b0:74:00, Type: Learned
Interface: ge-2/0/0.0
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: none

```

Interface Status TLV

The Interface Status TLV indicates the status of the interface on which the MEP transmitting the CCM is configured, or the next-lower interface in the IETF RFC 2863 IF-MIB. The format of this TLV is shown in Table 63 on page 703. The enumerated values are shown in Table 64 on page 703.

Table 63: Interface Status TLV Format

Parameter	Octet (Sequence)
Type = 4	1
Length	2–3
Value (See Table 64 on page 703)	4

Table 64: Interface Status TLV Values

Mnemonic	Interface Status	Value
isUp	up	1
isDown	down	2
isTesting	testing	3
isUnknown	unknown	4
isDormant	dormant	5
isNotPresent	notPresent	6
isLowerLayerDown	lowerLayerDown	7

Configuring Interface Status TLV

The JUNOS Software provides configuration support for the Interface Status TLV, thereby allowing operators to control the transmission of this TLV in CCM PDUs through configuration at the continuity-check level.



NOTE: This configuration is not mandated by IEEE 802.1ag; rather it is provided to give more flexibility to the operator. The JUNOS Software receives and processes CCMs with the Interface Status TLV, regardless of this configuration.

The interface status TLV configuration is shown below:

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        maintenance-domain identifier {
          level number;
          maintenance-association identifier {
            continuity-check {
              interval number;
              loss-threshold number;
              hold-interval number;
              interface-status-tlv; # Sets the interface status TLV
            }
          }
        }
      }
    }
  }
}
```



NOTE: The JUNOS Software supports transmission of only three out of seven possible values for the Interface Status TLV. The supported values are 1, 2, and 7. However, the JUNOS Software is capable of receiving any value for the Interface Status TLV.

Displaying the Received Interface Status TLV

The JUNOS Software saves the last received Interface Status TLV from the remote MEP. If the received Interface Status value does not correspond to one of the standard values listed in Table 63 on page 703, then the **show** command displays "unknown."

You can display this last saved Interface Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001
```



```

Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up
Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

Remote MEP identifier: 1001, State: ok
MAC address: 00:19:e2:b0:74:00, Type: Learned
Interface: ge-2/0/0.0
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: none # displays the Interface Status TLV state

```

Displaying the Transmitted Interface Status TLV

The JUNOS Software saves the last transmitted Interface Status TLV from a local MEP. If the transmission of Interface Status TLV has not been enabled, then the **show** command displays "none."

You can display the last transmitted Interface Status TLV using the **show oam ethernet connectivity-fault-management mep-database maintenance-domain *identifier* maintenance-association *identifier* local-mep *identifier* remote-mep *identifier*** command, as in the following example:

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 local-mep 2001 remote-mep 1001

```

```

Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3 frames
MEP identifier: 2001, Direction: down, MAC address: 00:19:e2:b2:81:4a
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up
Interface name: ge-2/0/0.0, Interface status: Active, Link status: Up

Remote MEP identifier: 1001, State: ok
MAC address: 00:19:e2:b0:74:00, Type: Learned
Interface: ge-2/0/0.0
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: none

```

MAC Status Defects

The JUNOS Software provides MAC status defect information, indicating that one or more of the remote MEPs is reporting a failure in its Port Status TLV or Interface Status TLV. It indicates "yes" if either some remote MEP is reporting that its interface is not isUp (for example, at least one remote MEPs interface is unavailable), or if all remote MEPs are reporting a Port Status TLV that contains some value other than psUp (for example, all remote MEPs Bridge Ports are not forwarding data). There are two **show** commands you can use to view the MAC Status Defects indication.

Use the `mep-database` command to display MAC status defects:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md6 maintenance-association ma6
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 500, Direction: down, MAC address: 00:05:85:73:7b:39
Auto-discovery: enabled, Priority: 0
Interface status TLV: up, Port status TLV: up
Interface name: xe-5/0/0.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM          : no
  Erroneous CCM received                 : no
  Cross-connect CCM received             : no
  RDI sent by some MEP                  : no
  Some remote MEP's MAC in error state   : yes # MAC Status Defects yes/no
Statistics:
  CCMs sent                             : 1658
  CCMs received out of sequence          : 0
  LBMs sent                             : 0
  Valid in-order LBRs received            : 0
  Valid out-of-order LBRs received        : 0
  LBRs received with corrupted data       : 0
  LBRs sent                              : 0
  LTMs sent                             : 0
  LTMs received                         : 0
  LTRs sent                             : 0
  LTRs received                         : 0
  Sequence number of next LTM request     : 0
  1DMs sent                             : 0
  Valid 1DMs received                    : 0
  Invalid 1DMs received                   : 0
  DMMs sent                             : 0
  DMRs sent                             : 0
  Valid DMRs received                    : 0
  Invalid DMRs received                   : 0
Remote MEP count: 1
  Identifier  MAC address  State  Interface
    200      00:05:85:73:39:4a    ok    xe-5/0/0.0
```

Use the `interfaces` command to display MAC status defects:

```
user@host> show oam ethernet connectivity-fault-management interfaces detail
Interface name: xe-5/0/0.0, Interface status: Active, Link status: Up
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
Interface status TLV: up, Port status TLV: up
MEP identifier: 500, Direction: down, MAC address: 00:05:85:73:7b:39
MEP status: running
Defects:
  Remote MEP not receiving CCM          : no
  Erroneous CCM received                 : no
  Cross-connect CCM received             : no
  RDI sent by some MEP                  : no
  Some remote MEP's MAC in error state   : yes # MAC Status Defects
yes/no
Statistics:
  CCMs sent                             : 1328
  CCMs received out of sequence          : 0
```

```

LBMs sent : 0
Valid in-order LBRs received : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent : 0
LTMs sent : 0
LTMs received : 0
LTRs sent : 0
LTRs received : 0
Sequence number of next LTM request : 0
1DMs sent : 0
Valid 1DMs received : 0
Invalid 1DMs received : 0
DMRs sent : 0
DMRs received : 0
Valid DMRs received : 0
Invalid DMRs received : 0
Remote MEP count: 1
Identifier MAC address State Interface
200 00:05:85:73:39:4a ok xe-5/0/0.0

```

Configuring Remote MEP Action Profile Support

Based on values of `interface-status-tlv` and `port-status-tlv` in the received CCM packets, a specific action, such as `interface-down`, can be taken using the `action-profile` options. Multiple action profiles can be configured on the router, but only one action profile can be assigned to a remote MEP.

The action profile can be configured with at least one event to trigger the action; but the action will be triggered if any one of these events occurs. It is not necessary for all of the configured events to occur to trigger action.

An action-profile can be applied only at the remote MEP level.

The following example shows an action profile configuration with explanatory comments added:

```

[edit protocols oam ethernet connectivity-fault-management]
action-profile tlv-action {
  event {
    # If interface status tlv with value specified in the config is received
    interface-status-tlv down|lower-layer-down;
    # If port status tlv with value specified in the config is received
    port-status-tlv blocked;
    # If connectivity is lost to the peer */
    adjacency-loss;
  }
  action {
    # Bring the interface down */
    interface-down;
  }
  default-actions interface-down;
}
# domains
maintenance-domain identifier {
  # maintenance domain level (0-7)

```

```

    level number;
    # association
    maintenance-association identifier {
        mep identifier {
            interface ge-x/y/z.w;
            remote-mep identifier {
                # Apply the action-profile for the remote MEP
                action-profile tlv-action;
            }
        }
    }
}

```

Monitoring a Remote MEP Action Profile

You can use the `show oam ethernet connectivity-fault-management mep-database` command to view the action profile status of a remote MEP, as in the following example:

**show oam ethernet
connectivity-fault-
management
mep-database
remote-mep
(Action Profile Event)**

```

user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md5 maintenance-association ma5 remote-mep 200
Maintenance domain name: md5, Format: string, Level: 5
Maintenance association name: ma5, Format: string
Continuity-check status: enabled, Interval: 1s, Loss-threshold: 3 frames
MEP identifier: 100, Direction: down, MAC address: 00:05:85:73:e8:ad
Auto-discovery: enabled, Priority: 0
Interface status TLV: none, Port status TLV: none # last status TLVs transmitted
by the router
Interface name: ge-1/0/8.0, Interface status: Active, Link status: Up

Remote MEP identifier: 200, State: ok # displays the remote MEP name and state

MAC address: 00:05:85:73:96:1f, Type: Configured
Interface: ge-1/0/8.0
Last flapped: Never
Remote defect indication: false
Port status TLV: none
Interface status TLV: lower-layer-down
Action profile: juniper # displays remote MEP's action profile identifier
Last event: Interface-status-tlv lower-layer-down # last remote MEP event

# to trigger action
Action: Interface-down, Time: 2009-03-27 14:25:10 PDT (00:00:02 ago)
# action occurrence time

```

Configuring M120 and MX Series Routers for CCC Encapsulated Packets

This section includes the following topics:

- Overview of IEEE 802.1ag CFM OAM Support for CCC Encapsulated Packets on page 709
- CFM Features Supported on Layer 2 VPN Circuits on page 709
- Configuring CFM for CCC Encapsulated Packets on page 709

Overview of IEEE 802.1ag CFM OAM Support for CCC Encapsulated Packets

Layer 2 virtual private network (L2VPN) is a type of virtual private network service used to transport customer's private Layer 2 traffic (for example, Ethernet, ATM or Frame Relay) over the service provider's shared IP/MPLS infrastructure. The service provider edge (PE) router must have an interface with circuit cross-connect (CCC) encapsulation to switch the customer edge (CE) traffic to the public network.


The IEEE 802.1ag Ethernet Connectivity Fault Management (CFM) is an OAM standard used to perform fault detection, isolation, and verification on virtual bridge LANs. M120 and MX Series routers provide CFM support for bridge/VPLS/routed interfaces and support 802.1ag Ethernet OAM for CCC encapsulated packets.

CFM Features Supported on Layer 2 VPN Circuits

CFM features supported on L2VPN circuits are as follows:

- Creation of up/down MEPs at any level on the CE-facing logical interfaces.
- Creation of MIPs at any level on the CE-facing logical interfaces.
- Support for continuity check, loopback, and linkrace protocol.
- Support for the Y1731 Ethernet Delay measurement protocol.
- Support for action profiles to bring the CE-facing logical interfaces down when loss of connectivity is detected.

Figure 67: Layer 2 VPN Topology



To monitor the L2VPN circuit, a CFM up MEP (Level 6 in Figure 67 on page 709) can be configured on the CE-facing logical interfaces of provider edge routers PE1 and PE2. To monitor the CE-PE attachment circuit, a CFM down MEP can be configured on the customer logical interfaces of CE1-PE1 and CE2-PE2 (Level 0 in Figure 67 on page 709).

Configuring CFM for CCC Encapsulated Packets

The only change from the existing CLI configuration is the introduction of a new command to create a MIP on the CE-facing interface of the PE router.

```
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
```

```

# Define a maintenance domains for each default level.
#; These names are specified as DEFAULT_level_number
maintenance-domain DEFAULT_x {
    # L2VPN CE interface
    interface (<ge> | <xe>) fpc/pic/port.domain;
}
{
    level number;
    maintenance-association identifier {
        mep mep-id {
            direction { up | down };
            # L2 VPN CE interface on which encapsulation family CCC is configured.
            interface (<ge> | <xe>) fpc/pic/port.domain;
            auto-discovery;
            priority number;
        }
    }
}
}
}
}
}
}

```

Chapter 44

Configuring ITU-T Y.1731 Ethernet Service OAM

This section contains the following topics:

- Ethernet Frame Delay Measurements Overview on page 711
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 717
- Guidelines for Starting an ETH-DM Session on page 718
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 720
- Configuring Routers to Support an ETH-DM Session on page 724
- Starting an ETH-DM Session on page 727
- Managing ETH-DM Statistics and ETH-DM Frame Counts on page 729
- Example: One-Way Ethernet Frame Delay Measurement on page 732

Ethernet Frame Delay Measurements Overview

This topic contains the following information:

- ITU-T Y.1731 Frame Delay Measurement Feature on page 711
- One-Way Ethernet Frame Delay Measurement on page 713
- Two-Way Ethernet Frame Delay Measurement on page 714
- Choosing Between One-Way and Two-Way ETH-DM on page 715
- Restrictions for Ethernet Frame Delay Measurement on page 716

ITU-T Y.1731 Frame Delay Measurement Feature

The IEEE 802.3-2005 standard for Ethernet Operations, Administration, and Maintenance (OAM) defines a set of link fault management mechanisms to detect and report link faults on a single point-to-point Ethernet LAN.

JUNOS Software supports key OAM standards that provide for automated end-to-end management and monitoring of Ethernet service by service providers:

- *IEEE Standard 802.1ag*, also known as “Connectivity Fault Management (CFM)”.
- *ITU-T Recommendation Y.1731*, which uses different terminology than IEEE 802.1ag and defines Ethernet service OAM features for fault monitoring, diagnostics, and performance monitoring.

These capabilities allow operators to offer binding service-level agreements (SLAs) and generate new revenues from rate- and performance-guaranteed service packages that are tailored to the specific needs of their customers.

Ethernet CFM

The IEEE 802.1ag standard for connectivity fault management (CFM) defines mechanisms to provide for end-to-end Ethernet service assurance over any path, whether a single link or multiple links spanning networks composed of multiple LANs.

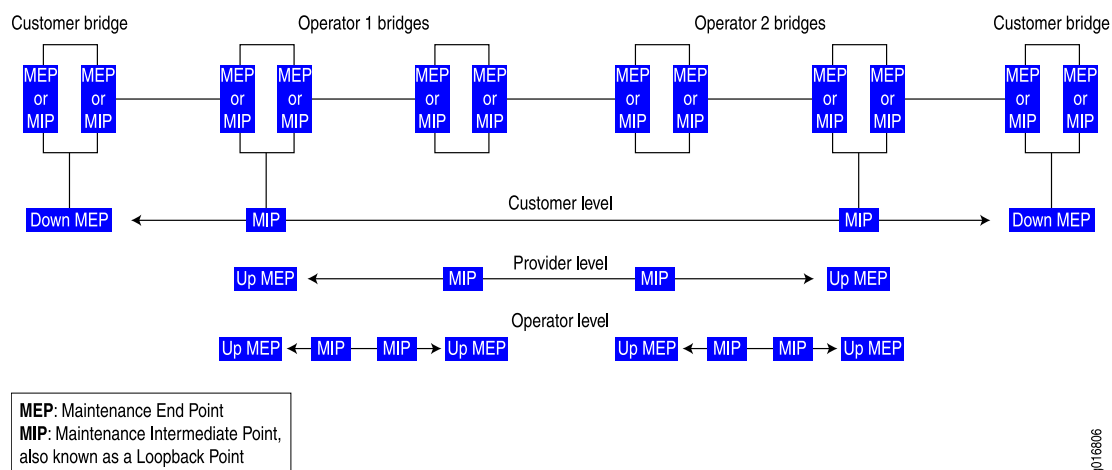
For Ethernet interfaces on M320, MX Series, and T Series routers, JUNOS Software supports the following key elements of the Ethernet CFM standard:

- Fault monitoring using the IEEE 802.1ag Ethernet OAM Continuity Check protocol
- Path discovery and fault verification using the IEEE 802.1ag Ethernet OAM Linktrace protocol
- Fault isolation using the IEEE 802.1ag Ethernet OAM Loopback protocol

In a CFM environment, network entities such as network operators, service providers, and customers may be part of different administrative domains. Each administrative domain is mapped into one maintenance domain. Maintenance domains are configured with different level values to keep them separate. Each domain provides enough information for the entities to perform their own management and end-to-end monitoring, and still avoid security breaches.

Figure 68 on page 712 shows the relationships among the customer, provider, and operator Ethernet bridges, maintenance domains, maintenance association end points (MEPs), and maintenance intermediate points (MIPs).

Figure 68: Relationship of MEPs, MIPs, and Maintenance Domain Levels



Ethernet Frame Delay Measurement

Two key objectives of OAM functionality are to measure quality-of-service attributes such as frame delay and frame delay variation (also known as “frame jitter”). Such measurements can enable you to identify network problems before customers are impacted by network defects.

JUNOS Software supports Ethernet frame delay measurement between MEPs configured on Ethernet physical or logical interfaces on Dense Port Concentrators (DPCs) in MX Series routers. Ethernet frame delay measurement provides fine control to operators for triggering delay measurement on a given service and can be used to monitor SLAs. Ethernet frame delay measurement also collects other useful information, such as worst and best case delays, average delay, and average delay variation. The JUNOS Software implementation of Ethernet frame delay measurement (ETH-DM) is fully compliant with the ITU-T Recommendation Y.1731, *OAM Functions and Mechanisms for Ethernet-based Networks*. The recommendation defines OAM mechanisms for operating and maintaining the network at the Ethernet service layer, which is called the “ETH layer” in ITU-T terminology.

One-Way Ethernet Frame Delay Measurement

In one-way ETH-DM mode, a series of frame delay and frame delay variation values are calculated based on the time elapsed between the time a measurement frame is sent from the initiator MEP at one router and the time when the frame is received at the receiver MEP at the other router.

1DM Transmission When you start a one-way frame delay measurement, the router sends 1DM frames—frames that carry the protocol data unit (PDU) for a one-way delay measurement—from the initiator MEP to the receiver MEP at the rate and for the number of frames you specify. The router marks each 1DM frame as drop-ineligible and inserts a timestamp of the transmission time into the frame.

1DM Reception When an MEP receives a 1DM frame, the router that contains the receiver MEP measures the one-way delay for that frame (the difference between the time the frame was received and the timestamp contained in the frame itself) and the delay variation (the difference between the current and previous delay values).

One-Way ETH-DM Statistics The router that contains the receiver MEP stores each set of one-way delay statistics in the ETH-DM database. The ETH-DM database collects up to 100 sets of statistics for any given CFM session (pair of peer MEPs). You can access these statistics at any time by displaying the ETH-DM database contents.

One-Way ETH-DM Frame Counts Each router counts the number of one-way ETH-DM frames sent and received:

- For an initiator MEP, the router counts the number of 1DM frames sent.
- For a receiver MEP, the router counts the number of valid 1DM frames received and the number of invalid 1DM frames received.

Each router stores ETH-DM frame counts in the CFM database. The CFM database stores CFM session statistics and, for interfaces that support ETH-DM, any ETH-DM frame counts. You can access the frame counts at any time by displaying CFM

database information for Ethernet interfaces assigned to MEPs or for MEPs in CFM sessions.

Synchronization of System Clocks

The accuracy of one-way delay calculations depends on close synchronization of the system clocks at the initiator MEP and receiver MEP.

The accuracy of one-way delay variation is not dependent on system clock synchronization. Because delay variation is simply the difference between consecutive one-way delay values, the out-of-phase period is eliminated from the frame jitter values.



NOTE: For a given one-way Ethernet frame delay measurement, frame delay and frame delay variation values are available only on the router that contains the receiver MEP.

Two-Way Ethernet Frame Delay Measurement

In two-way ETH-DM mode, frame delay and frame delay variation values are based on the time difference between when the initiator MEP transmits a request frame and receives a reply frame from the responder MEP, subtracting the time elapsed at the responder MEP.

DMM Transmission

When you start a two-way frame delay measurement, the router sends delay measurement message (DMM) frames—frames that carry the PDU for a two-way ETH-DM request—from the initiator MEP to the responder MEP at the rate and for the number of frames you specify. The router marks each DMM frame as drop-ineligible and inserts a timestamp of the transmission time into the frame.

DMR Transmission

When an MEP receives a DMM frame, the responder MEP responds with a delay measurement reply (DMR) frame, which carries ETH-DM reply information and a copy of the timestamp contained in the DMM frame.

DMR Reception

When an MEP receives a valid DMR, the router that contains the MEP measures the two-way delay for that frame based on the following sequence of timestamps:

1. TI_{TxDMM}

Time at which the initiator MEP transmits a two-way ETH-DM DMM frame to the responder MEP.

2. $TR_{Rx DMM}$

Time at which the responder MEP receives a DMM frame from the initiator MEP.

3. $TR_{Tx DMR}$

Time at which the responder MEP transmits a two-way ETH-DM (DMR) frame, associated with a specific DMM frame, to the initiator MEP.

4. $TI_{Rx DMR}$

Time at which the initiator MEP receives a DMR frame from the responder MEP.

A two-way frame delay is calculated as follows:

$$[T_{I_{RxDMR}} - T_{I_{TxDMM}}] - [T_{R_{TxDMR}} - T_{R_{RxDMM}}]$$

In other words, frame delay is the difference between the time at which the initiator MEP sends a DMM frame and the time at which the initiator MEP receives the associated DMR frame from the responder MEP, minus the time elapsed at the responder MEP.

The delay variation is the difference between the current and previous delay values.

Two-Way ETH-DM Statistics

The router that contains the initiator MEP stores each set of two-way delay statistics in the ETH-DM database. The ETH-DM database collects up to 100 sets of statistics for any given CFM session (pair of peer MEPs). You can access these statistics at any time by displaying the ETH-DM database contents.

Two-Way ETH-DM Frame Counts

Each router counts the number of two-way ETH-DM frames sent and received:

- For an initiator MEP, the router counts the number DMM frames transmitted, the number of valid DMR frames received, and the number of invalid DMR frames received.
- For a responder MEP, the router counts the number of DMR frames sent.

Each router stores ETH-DM frame counts in the CFM database. The CFM database stores CFM session statistics and, for interfaces that support ETH-DM, any ETH-DM frame counts. You can access the frame counts at any time by displaying CFM database information for Ethernet interfaces assigned to MEPs or for MEPs in CFM sessions.



NOTE: For a given two-way Ethernet frame delay measurement, frame delay and frame delay variation values are available only at the router that contains the initiator MEP.

Choosing Between One-Way and Two-Way ETH-DM

One-way frame delay measurement requires that the system clocks at the initiator MEP and receiver MEP are closely synchronized. Two-way frame delay measurement does not require synchronization of the two systems. If it is not practical for the clocks to be synchronized, two-way frame delay measurements are more accurate.

When two systems are close to each other, their one-way delay values are very high compared to their two-way delay values. This is because one-way delay measurement requires that the timing for the two systems be synchronized at a very granular level, and MX Series routers currently do not support this granular synchronization.

Restrictions for Ethernet Frame Delay Measurement

The following restrictions apply to the Ethernet frame delay measurement feature:

- The Ethernet frame delay measurement feature is supported only for MEPs configured on Ethernet physical or logical interfaces on DPCs in MX Series routers. The ETH-DM feature is not supported on aggregated Ethernet interfaces or label-switched interface (LSI) pseudowires.
- Hardware-assisted timestamping for ETH-DM frames in the reception path is only supported for MEP interfaces on Enhanced DPCs and Enhanced Queuing DPCs in MX Series routers. For information about hardware-assisted timestamping, see “Guidelines for Configuring Routers to Support an ETH-DM Session” on page 717 and “Enabling the Hardware-Assisted Timestamping Option” on page 726.
- Ethernet frame delay measurements can be triggered only when the distributed periodic packet management daemon (**ppmd**) is enabled. For more information about this limitation, see “Guidelines for Configuring Routers to Support an ETH-DM Session” on page 717 and “Ensuring that Distributed **ppmd** Is Not Disabled” on page 725.
- You can monitor only one session at a time to the same remote MEP or MAC address. For more information about starting an ETH-DM session, see “Starting an ETH-DM Session” on page 727.
- ETH-DM statistics are collected at only one of the two peer routers in the ETH-DM session. For a one-way ETH-DM session, you can display frame ETH-DM statistics at the receiver MEP only, using ETH-DM-specific **show** commands. For a two-way ETH-DM session, you can display frame delay statistics at the initiator MEP only, using the same ETH-DM-specific **show** commands. For more information, see “Managing ETH-DM Statistics and ETH-DM Frame Counts” on page 729.
- ETH-DM frame counts are collected at both MEPs and are stored in the respective CFM databases.
- If graceful Routing Engine switchover (GRES) occurs, any collected ETH-DM statistics are lost, and ETH-DM frame counts are reset to zeroes. GRES enables a router with dual Routing Engines to switch from a master Routing Engine to a backup Routing Engine without interruption to packet forwarding. For more information, see the *JUNOS High Availability Configuration Guide*.
- Accuracy of frame delay statistics is compromised when the system is changing (such as from reconfiguration). We recommend performing Ethernet frame delay measurements on a stable system.

Related Topics

- Guidelines for Configuring Routers to Support an ETH-DM Session on page 717
- Guidelines for Starting an ETH-DM Session on page 718
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 720
- Example: One-Way Ethernet Frame Delay Measurement on page 732

Guidelines for Configuring Routers to Support an ETH-DM Session

Keep the following guidelines in mind when configuring routers to support an Ethernet frame delay measurement (ETH-DM) session:

- Configuration Requirements for ETH-DM on page 717
- Configuration Options for ETH-DM on page 717

Configuration Requirements for ETH-DM

You can obtain ETH-DM information for a link that meets the following requirements:

- The measurements can be performed between peer maintenance association endpoints (MEPs) on two routers.
- The two MEPs must be configured on two Ethernet physical interfaces or on two Ethernet logical interfaces. For more information, see “Configuring a Maintenance End Point” on page 686 and “Configuring a Remote Maintenance End Point” on page 688.
- The two MEPs must be configured—on their respective routers—under the same maintenance association (MA) identifier. For more information, see “Creating the Maintenance Association” on page 684.
- On both routers, the MA must be associated with the same maintenance domain (MD) name. For more information, see “Creating the Maintenance Domain” on page 682.
- On both routers, periodic packet management (PPM) must be running on the Routing Engine and Packet Forwarding Engine, which is the default configuration. You can disable PPM on the Packet Forwarding Engine only. However, the Ethernet frame delay measurement feature requires that distributed PPM remain enabled on the Packet Forwarding Engine of both routers. For more information about `ppmd`, see the *JUNOS Routing Protocols Configuration Guide*.
- If the PPM process (`ppmd`) is disabled on the Packet Forwarding Engine, you must re-enable it. Re-enabling distributed `ppmd` entails restarting the `ethernet-connectivity-fault-management` process, which causes all connectivity fault management (CFM) sessions to re-establish. For more information about CFM sessions, see “Configuring Ethernet Local Management Interface” on page 690.



NOTE: The Ethernet frame delay measurement feature is supported only for MEPs configured on Ethernet physical or logical interfaces on DPCs in MX Series routers. The ETH-DM feature is not supported on aggregated Ethernet interfaces or LSI pseudowires.

Configuration Options for ETH-DM

By default, the ETH-DM feature calculates frame delays using software-based timestamping of the ETH-DM PDU frames sent and received by the MEPs in the session. As an option that can increase the accuracy of ETH-DM calculations when

the DPC is loaded with heavy traffic in the receive direction, you can enable hardware-assisted timestamping of session frames in the receive direction.



NOTE: Hardware-assisted timestamping for ETH-DM frames is only supported for MEP interfaces on Enhanced DPCs and Enhanced Queuing DPCs in MX Series routers.

- Related Topics**
- Ethernet Frame Delay Measurements Overview on page 711
 - Configuring Routers to Support an ETH-DM Session on page 724

Guidelines for Starting an ETH-DM Session

Keep the following guidelines in mind when preparing to start an Ethernet frame delay measurement (ETH-DM) session:

- ETH-DM Session Prerequisites on page 718
- ETH-DM Session Parameters on page 718
- Restrictions for an ETH-DM Session on page 719

ETH-DM Session Prerequisites

Before you can start an ETH-DM session, you must configure two MX Series routers to support ETH-DM by defining the two CFM-enabled physical or logical Ethernet interfaces on each router. This entails creating and configuring CFM maintenance domains, maintenance associations, and maintenance association end points on each router. For more information about enabling CFM on an Ethernet interface, see “Creating the Maintenance Domain” on page 682.



NOTE: The Ethernet frame delay measurement feature is supported only for maintenance association end points configured on Ethernet physical or logical interfaces on DPCs in MX Series routers. The ETH-DM feature is not supported on aggregated Ethernet interfaces or LSI pseudowires.

For specific information about configuring routers to support ETH-DM, see “Guidelines for Configuring Routers to Support an ETH-DM Session” on page 717 and “Configuring Routers to Support an ETH-DM Session” on page 724.

ETH-DM Session Parameters

You can initiate a one-way or two-way ETH-DM session by entering the **monitor ethernet delay-measurement** operational command at a router that contains one end of the service for which you want to measure frame delay. The command options specify the ETH-DM session in terms of the CFM elements:

- The type of ETH-DM measurement (one-way or two-way) to be performed.

- The Ethernet service for which the ETH-DM measurement is to be performed:
 - CFM maintenance domain—Name of the existing maintenance domain (MD) for which you want to measure Ethernet frame delays. For more information, see “Creating the Maintenance Domain” on page 682.
 - CFM maintenance association—Name of an existing maintenance association (MA) within the maintenance domain. For more information, see “Creating the Maintenance Association” on page 684.
 - Remote CFM maintenance association end point—The unicast MAC address or the numeric identifier of the remote maintenance association end point (MEP)—the physical or logical interface on the remote router that resides in the specified MD and is named in the specified MA—with which to perform the ETH-DM session. For more information, see “Configuring a Maintenance End Point” on page 686 and “Configuring a Remote Maintenance End Point” on page 688.
- Optional specifications:
 - Count—You can specify the number of ETH-DM requests to send for this frame delay measurement session. The range is from 1 through 65,535 frames. The default value is 10 frames.

NOTE: Although you can trigger frame delay collection for up to 65,535 ETH-DM requests at a time, a router stores only the last 100 frame delay statistics per CFM session (pair of peer MEPs).

 - Frame interval—You can specify the number of seconds to elapse between ETH-DM frame transmittals. The default value is 1 second.

For more detailed information about the parameters you can specify to start an ETH-DM session, see the `monitor ethernet delay-measurement` operational command description in the *JUNOS System Basics and Services Command Reference*.

Restrictions for an ETH-DM Session

The following restrictions apply to an ETH-DM session:

- You cannot run multiple simultaneous ETH-DM sessions with the same remote MEP or MAC address.
- For a given ETH-DM session, you can collect frame delay information for a maximum of 65,535 frames.
- For a given CFM session (pair of peer MEPs), the ETH-DM database stores a maximum of 100 statistics, with the older statistics being “aged out” as newer statistics are collected for that pair of MEPs.
 - For one-way delay measurements collected within the same CFM session, the 100 most recent ETH-DM statistics can be retrieved at any point of time at the router on which the receiver MEP is defined.
 - For two-way delay measurements collected within the same CFM session, the 100 most recent ETH-DM statistics can be retrieved at any point of time at the router on which the initiator MEP is defined.

Depending on the number of frames exchanged in the individual ETH-DM sessions, the ETH-DM database can contain statistics collected through multiple ETH-DM sessions.

- If graceful Routing Engine switchover (GRES) occurs, any collected ETH-DM statistics are lost, and ETH-DM frame counts are reset to zeroes. GRES enables a router with dual Routing Engines to switch from a master Routing Engine to a backup Routing Engine without interruption to packet forwarding. For more information, see the *JUNOS High Availability Configuration Guide*.
- Accuracy of frame delay data is compromised when the system is changing (such as from reconfiguration). We recommend performing Ethernet frame delay measurements on a stable system.

Related Topics

- Ethernet Frame Delay Measurements Overview on page 711
- Starting an ETH-DM Session on page 727
- Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 720
- `monitor ethernet delay-measurement` operational command

Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts

This topic contains the following information:

- ETH-DM Statistics on page 720
- ETH-DM Statistics Retrieval on page 722
- ETH-DM Frame Counts on page 722
- ETH-DM Frame Count Retrieval on page 723

ETH-DM Statistics

Ethernet frame delay statistics are the frame delay and frame delay variation values determined by the exchange of frames containing ETH-DM protocol data units (PDUs).

- For a one-way ETH-DM session, statistics are collected in an ETH-DM database at the router that contains the receiver MEP. For a detailed description of one-way Ethernet frame delay measurement, including the exchange of one-way delay PDU frames, see “Ethernet Frame Delay Measurements Overview” on page 711.
- For a two-way ETH-DM session, statistics are collected in an ETH-DM database at the router that contains the initiator MEP. For a detailed description of two-way Ethernet frame delay measurement, including the exchange of two-way delay PDU frames, see “Ethernet Frame Delay Measurements Overview” on page 711.

A CFM database stores CFM-related statistics and—for Ethernet interfaces that support ETH-DM—the 100 most recently collected ETH-DM statistics for that pair of MEPs. You can view ETH-DM statistics by using the `delay-statistics` or `mep-statistics` form of the `show oam ethernet connectivity-fault-management` command to display the CFM statistics for the MEP that collects the ETH-DM statistics you want to view.

Table 65 on page 721 describes the ETH-DM statistics calculated in an ETH-DM session.

Table 65: ETH-DM Statistics

Field Name	Field Description
One-way delay (µsec) [†]	<p>For a one-way ETH-DM session, the frame delay, in microseconds, collected at the receiver MEP.</p> <p>To display frame delay statistics for a given one-way ETH-DM session, use the delay-statistics or mep-statistics form of the show oam ethernet connectivity-fault-management command at the receiver MEP for that session.</p>
Two-way delay (µsec)	<p>For a two-way ETH-DM session, the frame delay, in microseconds, collected at the initiator MEP.</p> <p>When you start a two-way frame delay measurement, the CLI output displays each DMR frame receipt timestamp and corresponding DMM frame delay and delay variation collected as the session progresses.</p> <p>To display frame delay statistics for a given two-way ETH-DM session, use the delay-statistics or mep-statistics form of the show oam ethernet connectivity-fault-management command at the initiator MEP for that session.</p>
Average delay [†]	<p>When you start a two-way frame delay measurement, the CLI output includes a runtime display of the average two-way frame delay among the statistics collected for the ETH-DM session only.</p> <p>When you display ETH-DM statistics using a show command, the Average delay field displays the average one-way and two- frame delays among all ETH-DM statistics collected at the CFM session level.</p> <p>For example, suppose you start two one-way ETH-DM sessions for 50 counts each, one after the other. If, after both measurement sessions complete, you use a show command to display 100 ETH-DM statistics for that CFM session, the Average delay field displays the average frame delay among all 100 statistics.</p>
Average delay variation [†]	<p>When you start a two-way frame delay measurement, the CLI output includes a runtime display of the average two-way frame delay variation among the statistics collected for the ETH-DM session only.</p> <p>When you display ETH-DM statistics using a show command, the Average delay variation field displays the average one-way and two- frame delay variations among all ETH-DM statistics collected at the CFM session level.</p>
Best-case delay [†]	<p>When you start a two-way frame delay measurement, the CLI output includes a runtime display of the lowest two-way frame delay value among the statistics collected for the ETH-DM session only.</p> <p>When you display ETH-DM statistics using a show command, the Best case delay field displays the lowest one-way and two-way frame delays among all ETH-DM statistics collected at the CFM session level.</p>
Worst-case delay [†]	<p>When you start a two-way frame delay measurement, the CLI output includes a runtime display of the highest two-way frame delay value among the statistics collected for the ETH-DM session only.</p> <p>When you display ETH-DM statistics using a show command, the Worst case delay field displays the highest one-way and two-way frame delays among all statistics collected at the CFM session level.</p>

Table 65: ETH-DM Statistics (*continued*)

Field Name	Field Description
------------	-------------------

[†]When you start a one-way frame delay measurement, the CLI output displays NA (“not available”) for this field. One-way ETH-DM statistics are collected at the remote (receiver) MEP. Statistics for a given one-way ETH-DM session are available only by displaying CFM statistics for the receiver MEP.

ETH-DM Statistics Retrieval

At the receiver MEP for a one-way session, or at the initiator MEP for a two-way session, you can display all ETH-DM statistics collected at a CFM session level by using the following operational commands:

- `show oam ethernet connectivity-fault-management delay-statistics maintenance-domain md-name maintenance-association ma-name <local-mep mep-id> <remote-mep mep-id> <count count>`
- `show oam ethernet connectivity-fault-management mep-statistics maintenance-domain md-name maintenance-association ma-name <local-mep mep-id> <remote-mep mep-id> <count count>`

ETH-DM Frame Counts

The number of ETH-DM PDU frames exchanged in a ETH-DM session are stored in the CFM database on each router.

Table 66 on page 722 describes the ETH-DM frame counts collected in an ETH-DM session.

Table 66: ETH-DM Frame Counts

Field Name	Field Description
1DMs sent	Number of one-way delay measurement (1DM) PDU frames sent to the peer MEP in this session. Stored in the CFM database of the MEP initiating a one-way frame delay measurement.
Valid 1DMs received	Number of valid 1DM frames received. Stored in the CFM database of the MEP receiving a one-way frame delay measurement.
Invalid 1DMs received	Number of invalid 1DM frames received. Stored in the CFM database of the MEP receiving a one-way frame delay measurement.
DMMs sent	Number of delay measurement message (DMM) PDU frames sent to the peer MEP in this session. Stored in the CFM database of the MEP initiating a two-way frame delay measurement.
DMRs sent	Number of delay measurement reply (DMR) frames sent (in response to a received DMM). Stored in the CFM database of the MEP responding to a two-way frame delay measurement.

Table 66: ETH-DM Frame Counts (*continued*)

Field Name	Field Description
Valid DMRs received	Number of valid DMR frames received. Stored in the CFM database of the MEP initiating a two-way frame delay measurement.
Invalid DMRs received	Number of invalid DMR frames received. Stored in the CFM database of the MEP initiating a two-way frame delay measurement.

ETH-DM Frame Count Retrieval

Frame Counts Stored in CFM Databases	<p>Each router counts the number of ETH-DM frames sent or received and stores the counts in a CFM database.</p> <p>You can display ETH-DM frame counts for MEPs assigned to specified Ethernet interfaces or for specified MEPs in CFM sessions by using the following operational commands:</p> <ul style="list-style-type: none"> ■ <code>show oam ethernet connectivity-fault-management interfaces (detail extensive)</code> ■ <code>show oam ethernet connectivity-fault-management mep-database maintenance-domain <i>md-name</i> maintenance-association <i>ma-name</i> <local-mep <i>mep-id</i>> <remote-mep <i>mep-id</i>></code>
One-Way ETH-DM Frame Counts	<p>For a one-way ETH-DM session, delay statistics are collected at the receiver MEP only, but frame counts are collected at both MEPs. As indicated in Table 66 on page 722, one-way ETH-DM frame counts are tallied from the perspective of each router in the session:</p> <ul style="list-style-type: none"> ■ At the initiator MEP, the router counts the number of 1DM frames sent. ■ At the receiver MEP, the router counts the number of valid 1DM frames received and the number of invalid 1DM frames received. <p>You can also view one-way ETH-DM frame counts—for a receiver MEP—by using the <code>show oam ethernet connectivity-fault-management mep-statistics</code> command to display one-way statistics and frame counts together.</p>
Two-Way ETH-DM Frame Counts	<p>For a two-way ETH-DM session, delay statistics are collected at the initiator MEP only, but frame counts are collected at both MEPs. As indicated in Table 66 on page 722, two-way ETH-DM frame counts are tallied from the perspective of each router in the session:</p> <ul style="list-style-type: none"> ■ At the initiator MEP, the router counts the number of DMM frames sent, valid DMR frames received, and invalid DMR frames received. ■ At the responder MEP, the router counts the number of DMR frames sent. <p>You can also view two-way ETH-DM frame counts—for an initiator MEP—by using the <code>show oam ethernet connectivity-fault-management mep-statistics</code> command to display two-way statistics and frame counts together.</p>

- Related Topics**
- Ethernet Frame Delay Measurements Overview on page 711
 - Managing ETH-DM Statistics and ETH-DM Frame Counts on page 729
 - Example: One-Way Ethernet Frame Delay Measurement on page 732
 - `clear oam ethernet connectivity-fault-management statistics` command
 - `show oam ethernet connectivity-fault-management mep-statistics` command
 - `show oam ethernet connectivity-fault-management delay-statistics` command
 - `show oam ethernet connectivity-fault-management interfaces (detail | extensive)` command
 - `show oam ethernet connectivity-fault-management mep-database` command

Configuring Routers to Support an ETH-DM Session

This topic contains the following tasks:

- Configuring MEP Interfaces on page 724
- Ensuring that Distributed pppd Is Not Disabled on page 725
- Enabling the Hardware-Assisted Timestamping Option on page 726

Configuring MEP Interfaces

Before you can start an Ethernet frame delay measurement session across an Ethernet service, you must configure two MX Series routers to support ETH-DM.

To configure an Ethernet interface on a DPC in MX Series router to support ETH-DM:

1. On each router, configure two physical or logical Ethernet interfaces connected by a VLAN. The following configuration is typical for single-tagged logical interfaces:

```
[edit interfaces]
interface {
  ethernet-interface-name {
    vlan-tagging;
    unit logical-unit-number {
      vlan-id vlan-id; # Both interfaces on this VLAN
    }
  }
}
```

Both interfaces will use the same VLAN ID.

2. On each router, attach peer MEPs to the two interfaces. The following configuration is typical:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      maintenance-domain md-name { # On both routers
        level number;
      }
    }
  }
}
```

```

maintenance-association ma-name { # On both routers
    continuity-check {
        interval 100ms;
        hold-interval 1;
    }
    mep mep-id { # Attach to VLAN interface
        auto-discovery;
        direction (up | down);
        interface interface-name;
        priority number;
    }
}
}
}
}
}
}

```

Ensuring that Distributed *ppmd* Is Not Disabled

By default, the router's period packet management process (*ppmd*) runs sessions distributed to the Packet Forwarding Engine in addition to the Routing Engine. This process is responsible for periodic transmission of packets on behalf of its various client processes, such as Bidirectional Forwarding Detection (BFD), and it also receives packets on behalf of client processes.

In addition, *ppmd* handles time-sensitive periodic processing and performs such processes as sending process-specific packets and gathering statistics. With *ppmd* processes running distributed on both the Routing Engine and the Packet Forwarding Engine, you can run such processes as BFD on the Packet Forwarding Engine.

Distributed *ppmd* Required for ETH-DM

Ethernet frame delay measurement requires that *ppmd* remains distributed to the Packet Forwarding Engine. If *ppmd* is not distributed to the Packet Forwarding Engines of both routers, ETH-DM PDU frame timestamps and ETH-DM statistics are not valid.

Before you start ETH-DM, you must verify that the following configuration statement is *NOT* present:

```

[edit]
routing-options {
    ppm {
        no-delegate-processing;
    }
}

```

If distributed *ppmd* processing is disabled (as shown in the stanza above) on either router, you must re-enable it in order to use the ETH-DM feature.

Procedure to Ensure that Distributed *ppmd* Is Not Disabled

To ensure that distributed *ppmd* is not disabled on a router:

1. Display the packet processing management (PPM) configuration to determine whether distributed *ppmd* has been disabled.
 - In the following example, distributed *ppmd* is enabled on the router. In this case, you do not need to modify the router configuration:

```
[edit]
user@host# show routing-options
ppm;
```

- In the following example, distributed **ppmd** is disabled on the router. In this case, you must proceed to Step 2 to modify the router configuration:

```
[edit]
user@host# show routing-options
ppm {
    no-delegate-processing;
}
```

2. Modify the router configuration to re-enable distributed **ppmd** and restart the Ethernet OAM Connectivity Fault Management process *ONLY IF* distributed **ppmd** is disabled (as determined in the previous step).

- a. Before continuing, make any necessary preparations for the possible loss of connectivity on the router.

Restarting the **ethernet-connectivity-fault-management** process has the following effect on your network:

- All connectivity fault management (CFM) sessions re-establish.
- All ETH-DM requests on the router terminate.
- All ETH-DM statistics and frame counts reset to 0.

- b. Modify the router configuration to re-enable distributed **ppmd**. For example:

```
[edit]
user@host# delete routing-options ppm no-delegate-processing
```

- c. Commit the updated router configuration. For example:

```
[edit]
user@host# commit and-quit
commit complete
exiting configuration mode
```

- d. To restart the Ethernet OAM Connectivity-Fault-Management process, enter the **restart ethernet-connectivity-fault-management** *<gracefully | immediately | soft>* operational mode command. For example:

```
user@host> restart ethernet-connectivity-fault-management
Connectivity fault management process started, pid 9893
```

Enabling the Hardware-Assisted Timestamping Option

By default, Ethernet frame delay measurement uses software for timestamping transmitted and received ETH-DM frames. For Ethernet interfaces on Enhanced Dense Port Concentrators (DPCs) and Enhanced Queuing DPCs only, you can

optionally use hardware timing to assist in the timestamping of received ETH-DM frames to increase the accuracy of delay measurements.

Enabling hardware-assisted timestamping of received frames can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction.

To enable Ethernet frame delay measurement hardware assistance on the reception path, include the `hardware-assisted-timestamping` statement at the `[edit protocols oam ethernet connectivity-fault-management performance-monitoring]` hierarchy level:

```
[edit protocols]
oam {
  ethernet {
    connectivity-fault-management {
      performance-monitoring {
        hardware-assisted-timestamping;
      }
    }
  }
}
```

Starting an ETH-DM Session

This topic contains the following information:

- Using the `monitor ethernet delay-measurement` Command on page 727
- Starting a One-Way ETH-DM Session on page 728
- Starting a Two-Way ETH-DM Session on page 728

Using the `monitor ethernet delay-measurement` Command

After you have configured two MX Series routers to support ITU-T Y.1731 Ethernet frame delay measurement (ETH-DM), you can initiate a one-way or two-way Ethernet frame delay measurement session from the CFM maintenance association end point (MEP) on one of the routers to the peer MEP on the other router.

To start an ETH-DM session between the specified local MEP and the specified remote MEP, enter the `monitor ethernet delay-measurement` command at operational mode. The syntax of the command is as follows:

```
monitor ethernet delay-measurement
(one-way | two-way)
maintenance-domain md-name
maintenance-association ma-name
(remote-mac-address | mep remote-mep-id)
<count frame-count>
<wait interval-seconds>
```

For a one-way frame delay measurement, the command displays a runtime display of the number of 1 DM frames sent from the initiator MEP during that ETH-DM session. One-way frame delay and frame delay variation measurements from an ETH-DM session are collected in a CFM database at the router that contains the receiver MEP. You can retrieve ETH-DM statistics from a CFM database at a later time.

For a two-way frame delay measurement, the command displays two-way frame delay and frame delay variation values for each round-trip frame exchange during that ETH-DM session, as well as a runtime display of useful summary information about the session: average delay, average delay variation, best-case delay, and worst-case delay. Two-way frame delay and frame delay variation values measurements from an ETH-DM session are collected in a CFM database at the router that contains the initiator MEP. You can retrieve ETH-DM statistics from a CFM database at a later time.



NOTE: Although you can trigger frame delay collection for up to 65,535 ETH-DM requests at a time, a router stores only the last 100 frame delay statistics per CFM session (pair of peer MEPs).

For a complete description of the `monitor ethernet delay-measurement` operational command, see the *JUNOS System Basics and Services Command Reference*.

Starting a One-Way ETH-DM Session

To start a one-way Ethernet frame delay measurement session, enter the `monitor ethernet delay-measurement one-way` command from operational mode, and specify the peer MEP by its MAC address or by its MEP identifier.

For example:

```
user@host> monitor ethernet delay-measurement one-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
One-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA
```



NOTE: If you attempt to monitor delays to a nonexistent MAC address, you must type **Ctrl + C** to explicitly quit the `monitor ethernet delay-measurement` command and return to the CLI command prompt.

Starting a Two-Way ETH-DM Session

To start a two-way Ethernet frame delay measurement session, enter the `monitor ethernet delay-measurement two-way` command from operational mode, and specify the peer MEP by its MAC address or by its MEP identifier.

For example:

```
user@host> monitor ethernet delay-measurement two-way 00:05:85:73:39:4a
maintenance-domain md6 maintenance-association ma6 count 10
Two-way ETH-DM request to 00:05:85:73:39:4a, Interface xe-5/0/0.0
DMR received from 00:05:85:73:39:4a Delay: 100 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 8 usec
```



```
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 111 usec Delay variation: 19 usec
DMR received from 00:05:85:73:39:4a Delay: 110 usec Delay variation: 1 usec
DMR received from 00:05:85:73:39:4a Delay: 119 usec Delay variation: 9 usec
DMR received from 00:05:85:73:39:4a Delay: 122 usec Delay variation: 3 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 30 usec
DMR received from 00:05:85:73:39:4a Delay: 92 usec Delay variation: 0 usec
DMR received from 00:05:85:73:39:4a Delay: 108 usec Delay variation: 16 usec
```

```
--- Delay measurement statistics ---
```

```
Packets transmitted: 10, Valid packets received: 10
```

```
Average delay: 103 usec, Average delay variation: 8 usec
```

```
Best case delay: 92 usec, Worst case delay: 122 usec
```



NOTE: If you attempt to monitor delays to a nonexistent MAC address, you must type **Ctrl + C** to explicitly quit the `monitor ethernet delay-measurement` command and return to the CLI command prompt.

-
- Related Topics**
- Ethernet Frame Delay Measurements Overview on page 711
 - Guidelines for Starting an ETH-DM Session on page 718
 - `monitor ethernet delay-measurement` command
 - Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 720
 - Managing ETH-DM Statistics and ETH-DM Frame Counts on page 729

Managing ETH-DM Statistics and ETH-DM Frame Counts

This topic contains the following information:

- Displaying ETH-DM Statistics Only on page 729
- Displaying ETH-DM Statistics and Frame Counts on page 730
- Displaying ETH-DM Frame Counts for MEPs by Enclosing CFM Entity on page 730
- Displaying ETH-DM Frame Counts for MEPs by Interface or Domain Level on page 731
- Clearing ETH-DM Statistics and Frame Counts on page 732

Displaying ETH-DM Statistics Only

Purpose Display ETH-DM statistics.

By default, the `show oam ethernet connectivity-fault-management delay-statistics` command displays ETH-DM statistics for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

- Action**
- To display the ETH-DM statistics collected for MEPs belonging to MA `ma1` and within MD `md1`:

```
user@host> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain ma1 maintenance-association ma1
```

- To display the ETH-DM statistics collected for ETH-DM sessions for the local MEP 201 belonging to MA ma2 and within MD md2:

```
user@host> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md2 maintenance-association ma2 local-mep 201
```

- To display the ETH-DM statistics collected for ETH-DM sessions from local MEPs belonging to MA ma3 and within MD md3 to remote MEP 302:

```
user@host> show oam ethernet connectivity-fault-management delay-statistics
maintenance-domain md3 maintenance-association ma3 remote-mep 302
```

Displaying ETH-DM Statistics and Frame Counts

Purpose Display ETH-DM statistics and ETH-DM frame counts.

By default, the `show oam ethernet connectivity-fault-management mep-statistics` command displays ETH-DM statistics and frame counts for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).

- Action**
- To display the ETH-DM statistics and ETH-DM frame counts for MEPs in MA ma1 and within MD md1:

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md1 maintenance-association ma1
```

- To display the ETH-DM statistics and ETH-DM frame counts for the local MEP 201 in MA ma2 and within MD md2:

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md2 maintenance-association ma2 local-mep 201
```

- To display the ETH-DM statistics and ETH-DM frame counts for the local MEP in MD md3 and within MA ma3 that participates in an ETH-DM session with the remote MEP 302:

```
user@host> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain ma3 maintenance-association ma3 remote-mep 302
```

Displaying ETH-DM Frame Counts for MEPs by Enclosing CFM Entity

Purpose Display ETH-DM frame counts for CFM maintenance association end points (MEPs).

By default, the `show oam ethernet connectivity-fault-management mep-database` command displays CFM database information for MEPs in the specified CFM maintenance association (MA) within the specified CFM maintenance domain (MD).



NOTE: At the router attached to the initiator MEP for a one-way session, or at the router attached to the receiver MEP for a two-way session, you can only display ETH-DM frame counts.

- Action** ■ To display CFM database information (including ETH-DM frame counts) for all MEPs in MA `ma1` within MD `md1`:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain ma1 maintenance-association ma1
```

- To display CFM database information (including ETH-DM frame counts) only for local MEP `201` in MA `ma1` within MD `md1`:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md2 maintenance-association ma2 local-mep 201
```

- To display CFM database information (including ETH-DM frame counts) only for remote MEP `302` in MD `md3` within MA `ma3`:

```
user@host> show oam ethernet connectivity-fault-management mep-database
maintenance-domain ma3 maintenance-association ma3 remote-mep 302
```

Displaying ETH-DM Frame Counts for MEPs by Interface or Domain Level

Purpose Display ETH-DM frame counts for CFM maintenance association end points (MEPs).

By default, the `show oam ethernet connectivity-fault-management interfaces` command displays CFM database information for MEPs attached to CFM-enabled Ethernet interfaces on the router or at a maintenance domain level. For Ethernet interfaces that support ETH-DM, any frame counts are also displayed when you specify the `detail` or `extensive` command option.



NOTE: At the router attached to the initiator MEP for a one-way session, or at the router attached to the receiver MEP for a two-way session, you can only display ETH-DM frame counts.

- Action** ■ To display CFM database information (including ETH-DM frame counts) for all MEPs attached to CFM-enabled Ethernet interfaces on the router:

```
user@host> show oam ethernet connectivity-fault-management interfaces
detail
```

- To display CFM database information (including ETH-DM frame counts) only for the MEPs attached to CFM-enabled router interface `ge-5/2/9.0`:

```
user@host> show oam ethernet connectivity-fault-management interfaces
ge-5/2/9.0 detail
```

- To display CFM database information (including ETH-DM frame counts) only for MEPs enclosed within CFM maintenance domains (MDs) at level 6:

```
user@host> show oam ethernet connectivity-fault-management interfaces
level 6 detail
```

Clearing ETH-DM Statistics and Frame Counts

Purpose Clear the ETH-DM statistics and ETH-DM frame counts.

By default, statistics and frame counts are deleted for all MEPs attached to CFM-enabled interfaces on the router. However, you can filter the scope of the command by specifying an interface name.

- Action**
- To clear the ETH-DM statistics and ETH-DM frame counts for all MEPs attached to CFM-enabled interfaces on the router:

```
user@host> clear oam ethernet connectivity-fault-management statistics
```

- To clear the ETH-DM statistics and ETH-DM frame counts only for MEPs attached to the logical interface `ge-0/5/9.0`:

```
user@host> clear oam ethernet connectivity-fault-management statistics
ge-0/5/9.0
```

- Related Topics**
- `clear oam ethernet connectivity-fault-management statistic` command
 - `show oam ethernet connectivity-fault-management delay-statistics` command
 - `show oam ethernet connectivity-fault-management interfaces (detail | extensive)` command
 - `show oam ethernet connectivity-fault-management mep-statistics` command
 - `show oam ethernet connectivity-fault-management mep-database` command

Example: One-Way Ethernet Frame Delay Measurement

This topic contains the following information:

- Description of the Example One-Way Frame Delay Measurement on page 733
- Steps for the Example One-Way Frame Delay Measurement on page 734

Description of the Example One-Way Frame Delay Measurement

This example shows how you can configure two MX Series routers (MX-PE1 and MX-PE2) to support an ETH-DM session between two peer MEPs (MEP 201 and MEP 101), initiate a one-way ETH-DM session (from MEP 101 to MEP 201), and then display the ETH-DM statistics and frame counts collected. To increase the accuracy of the ETH-DM statistics, enable optional hardware-assisted timestamping of received ETH-DM frames on the router that contains the receiver MEP.

Routers Used in This Example

To support one-way ETH-DM with optional hardware timestamping of frames on the reception path, the routers used in this example are configured as follows:

- Routers MX-PE1 and MX-PE2 are MX Series routers.
- The system clocks of routers MX-PE1 and MX-PE2 are closely synchronized.
- On router MX-PE1, interface `ge-5/2/9` is an Ethernet port on an Enhanced or Enhanced Queuing Dense Port Concentrator (DPC). The traffic load received on this DPC is heavy.
- On router MX-PE2, interface `ge-0/2/5` is an Ethernet port on a DPC.

ETH-DM Frame Counts for this Example

Both routers count the number of ETH-DM frames sent and received by the peer MEPs in the session and store the frame counts in the CFM databases as follows:

- At router MX-PE2, which contains the initiator MEP 101, the CFM database stores the ETH-DM frame counts for a one-way ETH-DM initiator (the count of 1DM frames sent).
- At router MX-PE1, which contains the receiver MEP 201, the CFM database stores the ETH-DM frame counts for a one-way ETH-DM receiver (the count of valid 1DM frames received and the count of invalid 1DM frames received).

ETH-DM Statistics for this Example

For a one-way frame delay measurement, only the router that contains the receiver MEP measures and stores frame delay statistics. In this example, ETH-DM statistics collected for the session are available only at router MX-PE1.

Steps for the Example One-Way Frame Delay Measurement

The following steps describe an example one-way Ethernet frame delay measurement:

1. At router **MX-PE1**, configure MEP **201** as a CFM maintenance association end point in CFM maintenance domain **md6** as follows:
 - a. Define the maintenance domain **md6** by associating it with maintenance domain level **6** and maintenance association identifier **ma6**.
 - b. Configure the maintenance association by specifying continuity protocol options and specifying MEP identifier **201**.
 - c. Configure MEP **201** by attaching it to logical interface **ge-5/2/9.0**, which is a single-tag interface on VLAN **512**.

The following configuration is only a partial example of a complete and functional router configuration:

```
[edit]
interfaces { # Configure a single-tag logical interface on VLAN 512
  ge-5/2/9 { # Interface must be on a DPC
    vlan-tagging;
    unit 0 {
      vlan-id 512;
    }
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        traceoptions {
          file eoam_cfm.log size 1g files 2 world-readable;
          flag all;
        }
        maintenance-domain md6 { # Define MD 'md6' on router MX-PE1
          level 6;
          maintenance-association ma6 { # Configure MA 'ma6' on router MX-PE1
            continuity-check {
              interval 100ms;
              hold-interval 1;
            }
            mep 201 { # Configure MEP 201 on router MX-PE1
              interface ge-5/2/9.0; # Attach to logical interface on VLAN 512
              direction down;
              auto-discovery;
            }
          }
        }
      }
    }
  }
}
```

2. At router **MX-PE2**, configure MEP 101 as a CFM maintenance association end point in CFM maintenance domain **md6** as follows:
 - a. Define the maintenance domain **md6** by associating it with maintenance domain level 6 and maintenance association identifier **ma6**.
 - b. Configure the maintenance association by specifying continuity protocol options and specifying MEP identifier **101**.
 - c. Configure MEP **101** by attaching it to logical interface **ge-0/2/5.0**, which is a single-tag interface on VLAN **512**.

The following configuration is only a partial example of a complete and functional configuration for router **MX-PE2**:

```
[edit]
interfaces { # Configure a single-tag logical interface on VLAN 512
  ge-0/2/5 { # Interface must be on a DPC
    vlan-tagging;
    unit 0 {
      vlan-id 512;
    }
  }
}
protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        traceoptions {
          file eoam_cfm.log size 1g files 2 world-readable;
          flag all;
        }
        maintenance-domain md6 { # Define MD 'md6' on router MX-PE2
          level 6;
          maintenance-association ma6 { # Configure MA 'ma6' on router MX-PE2
            continuity-check {
              interval 100ms;
              hold-interval 1;
            }
            mep 101 { # Configure MEP 101 on router MX-PE2
              interface ge-0/2/5.0; # Attach to logical interface on VLAN 512
              direction down;
              auto-discovery;
            }
          }
        }
      }
    }
  }
}
```

3. (Optional) To increase the accuracy of the ETH-DM statistics, modify the configuration of router **MX-PE1**, which contains the receiver MEP, by enabling hardware-assisted timestamping of 1DM frames received on the router.

```
[edit protocols]
oam {
```

```

ethernet {
    connectivity-fault-management {
        performance-monitoring {
            hardware-assisted-timestamping;
        }
    }
}

```



NOTE: The hardware-assisted timestamping option for ETH-DM is available for Ethernet interfaces on Enhanced or Enhanced Queuing DPCs only.

4. At router **MX-PE2**, start a one-way frame delay measurement session from local MEP 101 to remote MEP 201 on router **MX-PE1**:

```

user@MX-PE2> monitor ethernet delay-measurement one-way mep 201
maintenance-domain md6 maintenance-association ma6 count 10

One-way ETH-DM request to 00:90:69:0a:43:94, Interface ge-0/2/5.0
1DM Frames sent : 10
--- Delay measurement statistics ---
Packets transmitted: 10
Average delay: NA, Average delay variation: NA
Best case delay: NA, Worst case delay: NA

```

5. At router **MX-PE2**, which contains the initiator MEP, only the ETH-DM frame counts are available. Furthermore, the only frame count tallied for the initiator of a one-way frame delay measurement is the count of 1DM frames transmitted.

ETH-DM frame counts (the number of 1DM, DMM, and DMR frames exchanged during an ETH-DM session) are stored in the CFM database of both the initiator and receiver MEPs. When you display CFM database information, you can also display the ETH-DM frame counts. You can display CFM database information for all interfaces on the router, or you can limit the output to MEPs associated with certain CFM MDs and MAs.

- To display CFM database information for MEPs specified by enclosing CFM entities, use the **mep-database** form of the **show oam ethernet connectivity-fault-management** command. A CFM database also stores any ETH-DM frame counts.

In the example configuration for router **MX-PE2**, MEP 101 is the only MEP defined in MA **ma6** within MD **md6**. Therefore, the **show oam ethernet connectivity-fault management mep-database** command output displays CFM database information for MEP 101 only, even though you do not filter the command output by including the **local-mep** or **remote-mep** command options.

```

user@MX-PE2> show oam ethernet connectivity-fault-management mep-database
maintenance-domain md6 maintenance-association ma6

```

```

Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3

```



```

frames
MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
Auto-discovery: enabled, Priority: 0
Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                       : no
  Cross-connect CCM received                  : no
  RDI sent by some MEP                        : no
Statistics:
  CCMs sent                                   : 1590
  CCMs received out of sequence               : 0
  LBMs sent                                   : 0
  Valid in-order LBRs received                : 0
  Valid out-of-order LBRs received            : 0
  LBRs received with corrupted data           : 0
  LBRs sent                                   : 0
  LTMs sent                                   : 0
  LTMs received                              : 0
  LTRs sent                                   : 0
  LTRs received                              : 0
  Sequence number of next LTM request         : 0
  1DMs sent                                   : 10
  Valid 1DMs received                         : 0
  Invalid 1DMs received                       : 0
  DMMs sent                                   : 0
  DMRs sent                                   : 0
  Valid DMRs received                        : 0
  Invalid DMRs received                      : 0
Remote MEP count: 1
  Identifier  MAC address  State  Interface
    201      00:90:69:0a:43:94  ok    ge-0/2/5.0

```

- To display CFM database information for MEPs specified by interface name, use the `interfaces detail` form of the `show oam ethernet connectivity-fault-management` command. A CFM database also stores any ETH-DM frame counts.

In the example configuration for router **MX-PE2**, MEP 101 is the only MEP assigned to an interface on the router. Therefore, the `show oam ethernet connectivity-fault-management interfaces (detail | extensive)` command output displays CFM database information for MEP 101 only, even though you do not filter the command output by including the *ethernet-interface-name* or level *md-level* command options.

```

user@MX-PE2> show oam ethernet connectivity-fault-management interfaces
detail

```

```

Interface name: ge-0/2/5.0, Interface status: Active, Link status: Up
Maintenance domain name: md6, Format: string, Level: 6
Maintenance association name: ma6, Format: string
Continuity-check status: enabled, Interval: 100ms, Loss-threshold: 3
frames
MEP identifier: 101, Direction: down, MAC address: 00:90:69:0a:48:57
MEP status: running
Defects:
  Remote MEP not receiving CCM                : no
  Erroneous CCM received                       : no

```

```

Cross-connect CCM received          : no
RDI sent by some MEP               : no
Statistics:
CCMs sent                          : 1590
CCMs received out of sequence      : 0
LBMs sent                          : 0
Valid in-order LBRs received       : 0
Valid out-of-order LBRs received   : 0
LBRs received with corrupted data  : 0
LBRs sent                          : 0
LTMs sent                          : 0
LTMs received                      : 0
LTRs sent                          : 0
LTRs received                      : 0
Sequence number of next LTM request : 0
1DMs sent                          : 10
Valid 1DMs received                : 0
Invalid 1DMs received              : 0
DMMs sent                          : 0
DMRs sent                          : 0
Valid DMRs received                : 0
Invalid DMRs received              : 0
Remote MEP count: 1
Identifier  MAC address  State  Interface
   201      00:90:69:0a:43:94  ok    ge-0/2/5.0

```



NOTE: You can use these same commands—`show oam ethernet connectivity-fault-management mep-database` and `show oam ethernet connectivity-fault-management interfaces (detail | extensive)`—at router MX-PE1 to display the CFM database information (which includes any ETH-DM frame counts) for receiver MEP 201.

6. At router MX-PE1, which contains the receiver MEP, you can use two different `show oam ethernet connectivity-fault-management` commands to display ETH-DM statistics and ETH-DM frame counts.
 - a. To display only the delay statistics, use the `delay-statistics` form of the `show oam ethernet connectivity-fault-management` command:

```
user@MX-PE1> show oam ethernet connectivity-fault-management
delay-statistics maintenance-domain md6
```

```
MEP identifier: 201, MAC address: 00:90:69:0a:43:94
Remote MEP count: 1
```

```
Remote MAC address: 00:90:69:0a:48:57
Delay measurement statistics:
Index  One-way delay  Two-way delay
      (usec)         (usec)
  1          370
  2          357
  3          344
  4          332
  5          319
  6          306
  7          294
  8          281
  9          269
 10          255
Average one-way delay          : 312 usec
Average one-way delay variation: 11 usec
Best case one-way delay        : 255 usec
Worst case one-way delay       : 370 usec
```

- b. To display both the ETH-DM statistics and the CFM database information (which includes any ETH-DM frame counts), use the `mep-statistics` form of the `show oam ethernet connectivity-fault-management` command:

```
user@MX-PE1> show oam ethernet connectivity-fault-management mep-statistics
maintenance-domain md6
```

```
MEP identifier: 201, MAC address: 00:90:69:0a:43:94
Remote MEP count: 1
CCMs sent                      : 3240
CCMs received out of sequence  : 0
LBRs sent                      : 0
Valid in-order LBRs received   : 0
Valid out-of-order LBRs received : 0
LBRs received with corrupted data : 0
LBRs sent                     : 0
LTMs sent                     : 0
```

```

LTMs received           : 0
LTRs sent               : 0
LTRs received           : 0
Sequence number of next LTM request : 0
1DMs sent               : 0
Valid 1DMs received     : 10
Invalid 1DMs received   : 0
DMRs sent               : 0
DMRs sent               : 0
Valid DMRs received     : 0
Invalid DMRs received   : 0

```

```

Remote MEP identifier: 101
Remote MAC address: 00:90:69:0a:48:57

```

Delay measurement statistics:

Index	One-way delay (usec)	Two-way delay (usec)
1	370	
2	357	
3	344	
4	332	
5	319	
6	306	
7	294	
8	281	
9	269	
10	255	

```

Average one-way delay      : 312 usec
Average one-way delay variation: 11 usec
Best case one-way delay    : 255 usec
Worst case one-way delay   : 370 usec

```

- Related Topics**
- Guidelines for Configuring Routers to Support an ETH-DM Session on page 717
 - Guidelines for Starting an ETH-DM Session on page 718
 - Guidelines for Managing ETH-DM Statistics and ETH-DM Frame Counts on page 720

Chapter 45

Configuring IEEE 802.1x Port-Based Network Access Control

This section contains the following topics:

- IEEE 802.1x Port-Based Network Access Control Overview on page 741
- Administrative State of the Authenticator Port on page 742
- Administrative Mode of the Authenticator Port on page 742
- Configuring the Authenticator on page 742
- Viewing the dot1x Configuration on page 743

IEEE 802.1x Port-Based Network Access Control Overview

MX Series routers support the IEEE 802.1x Port-Based Network Access Control (dot1x) protocol on Ethernet interfaces for validation of client and user credentials to prevent unauthorized access to a specified router port. Before authentication is complete, only 802.1x control packets are allowed and forwarded to the router control plane for processing. All other packets are dropped.

Authentication methods used must be 802.1x compliant. Authentication using RADIUS and Microsoft Active Directory servers is supported. The following user/client authentication methods are allowed:

- EAP-MD5 (RFC 3748)
- EAP-TTLS requires a server certificate (RFC 2716)
- EAP-TLS requires a client and server certificate
- PEAP requires only a server certificate

You can use both client and server certificates in all types of authentication except EAP-MD5.



NOTE: On the MX Series router, 802.1x can be enabled on bridged ports only and not on routed ports.

Dynamic changes to a user session are supported to allow the router administrator to terminate an already authenticated session by using the “RADIUS disconnect” message defined in RFC 3576.

Administrative State of the Authenticator Port

The administrative state of an authenticator port can take any of the following three states:

- Force authorized—Allows network access to all users of the port without requiring them to be authenticated. This is equivalent to not having any authentication enabled on the port.
- Force unauthorized—Denies network access to all users of the port. This is equivalent to disabling the port.
- Automatic—This is the default mode where the authentication server response determines if the port is opened for traffic or not. Only the successfully authenticated clients are allowed access, all others are denied.

In JUNOS Software, the default mode is “automatic”. The “force authorized” and “force unauthorized” admin modes are not supported. You can achieve the functionality of “force authorized” mode by disabling **dot1x** on the required port. You can achieve the functionality of “force unauthorized” mode by disabling the port itself.

Administrative Mode of the Authenticator Port

JUNOS Software supports the supplicant mode of “Single” and not the “Single Secure” or “Multiple” modes. The “Single” mode option authenticates only the first client that connects to a port. All other clients that connect later (802.1x compliant or noncompliant) are allowed free access on that port without any further authentication. If the first authenticated client logs out, all other users are locked out until a client authenticates again.

Configuring the Authenticator

To configure IEEE 802.1x Port-Based Network Access Control protocol on Ethernet interfaces you must configure the **authenticator** statement at the **[edit protocols dot1x]** hierarchy level. Use the **authentication-profile-name** *access-profile-name* statement to specify the authenticating RADIUS server, and use the **interface** statement to specify and configure the Gigabit Ethernet or Fast Ethernet interface on the router specifically for IEEE 802.1x protocol use; both at the **[edit protocols dot1x authenticator]** hierarchy level.

```
[edit protocols dot1x]
authenticator {
  authentication-profile-name access-profile-name;
  interface (xe-fpc/pic/port | ge-fpc/pic/port | fe-fpc/pic/port) {
    maximum-requests seconds;
    quiet-period seconds;
    reauthentication (disable | interval seconds);
```

```
    retries integer;  
    server-timeout seconds;  
    supplicant (single);  
    supplicant-timeout seconds;  
    transmit-period seconds;  
  }  
}
```

Viewing the dot1x Configuration

To view all dot1x configurations, use the **show dot1x interface** operational mode command. To view a dot1x configuration for a specific interface, use the **show dot1x interface** (*xe-fpc/pic/port* | *ge-fpc/pic/port* | *fe-fpc/pic/port*) **detail** operational mode command. See the *Network Interfaces Command Reference* for more information about this command.

Chapter 46

Configuring IEEE 802.3ah OAM Link-Fault Management

This section includes the following topics:

- IEEE 802.3ah OAM Link-Fault Management Overview on page 745
- Configuring IEEE 802.3ah OAM Link-Fault Management on page 746
- Enabling IEEE 802.3ah OAM Support on page 746
- Configuring Link Discovery on page 746
- Configuring the OAM PDU Interval on page 747
- Configuring the OAM PDU Threshold on page 747
- Configuring Threshold Values for Local Fault Events on an Interface on page 747
- Disabling the Sending of Link Event TLVs on page 748
- Detecting Remote Faults on page 748
- Configuring an OAM Action Profile on page 748
- Specifying the Actions to Be Taken for Link-Fault Management Events on page 749
- Monitoring the Loss of Link Adjacency on page 750
- Monitoring Protocol Status on page 750
- Configuring Threshold Values for Fault Events in an Action Profile on page 750
- Applying an Action Profile on page 751
- Setting a Remote Interface into Loopback Mode on page 751
- Enabling Remote Loopback Support on the Local Interface on page 751
- Example: Configuring IEEE 802.3ah OAM Support on an Interface on page 752

IEEE 802.3ah OAM Link-Fault Management Overview

Ethernet interfaces capable of running at 100 Mbps or faster on MX Series, M Series (except M10 and M7 routers), and T Series routers support the IEEE 802.3ah standard for Operation, Administration, and Management (OAM). You can configure IEEE 802.3ah OAM on Ethernet point-to-point direct links or links across Ethernet repeaters. The IEEE 802.3ah standard meets the requirement for OAM capabilities as Ethernet moves from being solely an enterprise technology to being a WAN and access technology, as well as being backward-compatible with existing Ethernet technology. JUNOS Software supports IEEE 802.3ah link-fault management.

The features of link-fault management are:

- Discovery
- Link monitoring
- Remote fault detection
- Remote loopback



NOTE: Ethernet running on top of a Layer 2 protocol, such as Ethernet over ATM, is not supported in OAM configurations.

Configuring IEEE 802.3ah OAM Link-Fault Management

You can configure threshold values for fault events that trigger the sending of link event TLVs when the values exceed the threshold. To set threshold values for fault events on an interface, include the `event-thresholds` statement at the `[edit protocols oam ethernet link-fault-management interface]` hierarchy level.

You can also configure OAM threshold values within an action profile and apply the action profile to multiple interfaces. To create an action profile, include the `action-profile` statement at the `[edit protocols oam ethernet link-fault-management]` hierarchy level.

To view OAM statistics, use the `show oam ethernet link-fault-management operational` mode command. To clear OAM statistics, use the `clear oam ethernet link-fault-management statistics` operational mode command. To clear link-fault management state information and restart the link discovery process on Ethernet interfaces, use the operational `clear oam ethernet link-fault-management state` mode command. For more information about these commands, see the *JUNOS Interfaces Command Reference*.

Enabling IEEE 802.3ah OAM Support

To enable IEEE 802.3ah OAM support, include the `interface` statement at the `[edit protocols oam ethernet link-fault-management]` hierarchy level:

```
[edit protocols oam ethernet link-fault-management interface interface-name]
```

When you enable IEEE 802.3ah OAM on a physical interface, the discovery process is automatically triggered.

Configuring Link Discovery

When the IEEE 802.3ah OAM protocol is enabled on a physical interface, the discovery process is automatically triggered. The discovery process permits Ethernet interfaces to discover and monitor the peer on the link if it also supports the IEEE 802.3ah standard.

You can specify the discovery mode used for IEEE 802.3ah OAM support. The discovery process is triggered automatically when OAM IEEE 802.3ah functionality is enabled on a port. Link monitoring is done when the interface sends periodic OAM PDUs.

To configure the discovery mode, include the `link-discovery` statement at the `[edit protocol oam ethernet link-fault-management interface interface-name]` hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]  
link-discovery (active | passive);
```

In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. After the discovery process has been initiated, both sides participate in discovery.

Configuring the OAM PDU Interval

Periodic OAM PDUs are sent to perform link monitoring.

You can specify the periodic OAM PDU sending interval for fault detection.

To configure the sending interval, include the `pdu-interval` statement at the `[edit protocol oam ethernet link-fault-management interface interface-name]` hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]  
pdu-interval interval;
```

The periodic OAM PDU interval range is from 100 through 1000 milliseconds. The default sending interval is 1000 milliseconds.

Configuring the OAM PDU Threshold

You can specify the number of OAM PDUs that an interface can miss before the link between peers is considered down.

To configure the number of PDUs that can be missed from the peer, include the `pdu-threshold` statement at the `[edit protocol oam ethernet link-fault-management interface interface-name]` hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]  
pdu-threshold threshold-value;
```

The threshold value range is from 3 through 10. The default is three PDUs.

Configuring Threshold Values for Local Fault Events on an Interface

You can configure threshold values on an interface for the local errors that trigger the sending of link event TLVs.

To set the error threshold values for sending event TLVs, include the `frame-error`, `frame-period`, `frame-period-summary`, and `symbol-period` statements at the `[edit protocols`

oam ethernet link-fault-management interface *interface-name* event-thresholds] hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]
event-thresholds {
    frame-error count;
    frame-period count;
    frame-period-summary count;
    symbol-period count;
}
```

Disabling the Sending of Link Event TLVs

You can disable the sending of link event TLVs.

To disable the monitoring and sending of PDUs containing link event TLVs in periodic PDUs, include the `no-allow-link-events` statement at the [edit protocols oam ethernet link-fault-management interface *interface-name* negotiation-options] hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name
negotiation-options]
no-allow-link-events;
```

Detecting Remote Faults

Fault detection is either based on flags or fault event type length values (TLVs) received in OAM protocol data units (PDUs). Flags that trigger a link fault are:

- Critical Event
- Dying Gasp
- Link Fault

The link event TLVs are sent by the remote DTE by means of event notification PDUs. Link event TLVs are:

- Errored Symbol Period Event
- Errored Frame Event
- Errored Frame Period Event
- Errored Frame Seconds Summary Event

Configuring an OAM Action Profile

You can create an action profile to define event fault flags and thresholds and the action to be taken. You can then apply the action profile to one or more interfaces.

To configure an action profile, include the `action-profile` statement at the [edit protocols oam ethernet link-fault-management] hierarchy level:

```
action-profile profile-name {
```

```

    action {
        syslog;
        link-down;
        send-critical-event;
    }
    event {
        link-adjacency-loss;
        link-event-rate {
            frame-error count;
            frame-period count;
            frame-period-summary count;
            symbol-period count;
        }
        protocol-down;
    }
}

```

Specifying the Actions to Be Taken for Link-Fault Management Events

You can specify the action to be taken by the system when the configured link-fault event occurs. Multiple action profiles can be applied to a single interface. For each action-profile, at least one event and one action must be specified. The actions are taken only when all of the events in the action profile are true. If more than one action is specified, all the actions are executed.

You might want to set a lower threshold for a specific action such as logging the error and set a higher threshold for another action such as sending a critical event TLV.

To specify the action, include the `action` statement at the `[edit protocols oam ethernet link-fault-management action-profile profile-name]` hierarchy level:

```

[edit protocol oam ethernet link-fault-management action-profile profile-name]
event {
    link-adjacency-loss;
    protocol-down;
}
action {
    syslog;
    link-down;
    send-critical-event;
}

```

To create a system log entry when the link-fault event occurs, include the `syslog` statement.

To administratively disable the link when the link-fault event occurs, include the `link-down` statement.

To send IEEE 802.3ah link event TLVs in the OAM PDU when a link-fault event occurs, include the `send-critical-event` statement.



NOTE: If multiple actions are specified in the action profile, all of the actions are executed in no particular order.

Monitoring the Loss of Link Adjacency

You can specify actions be taken when link adjacency is lost. When link adjacency is lost, the system takes the action defined in the **action** statement of the action profile.

To configure the system to take action when link adjacency is lost, include the **link-adjacency-loss** statement at the [edit protocols oam ethernet link-fault-management action-profile *profile-name* event] hierarchy level:

```
[edit protocol oam ethernet link-fault-management action-profile profile-name]
link-adjacency-loss;
```

Monitoring Protocol Status

When a higher level protocol has signaled a down status to the IEEE 802.3ah protocol, the system takes the action defined in the **action** statement of the action profile.

To monitor the IEEE 802.3ah protocol, include the **protocol-down** statement at the [edit protocols oam ethernet link-fault-management action-profile *profile-name* event] hierarchy level:

```
[edit protocol oam ethernet link-fault-management action-profile profile-name]
protocol-down;
```



NOTE: If multiple events are specified in the action profile, all the events must occur before the specified action is taken.

Configuring Threshold Values for Fault Events in an Action Profile

You can configure link event thresholds for received error events that trigger the action specified in the **action** statement. You can then apply the action profile to one or more interfaces.

To configure link event thresholds, include the **link-event-rate** statement at the [edit protocols oam ethernet link-fault-management action-profile *profile-name* event] hierarchy level:

```
link-event-rate {
  frame-error count;
  frame-period count;
  frame-period-summary count;
  symbol-period count;
}
```

Applying an Action Profile

You can apply an action profile to one or more interfaces.

To apply an action profile to an interface, include the `apply-action-profile` statement at the `[edit protocols oam ethernet link-fault-management action-profile interface interface-name]` hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]  
  apply-action-profile profile-name;
```

Setting a Remote Interface into Loopback Mode

You can configure the software to set the remote DTE into loopback mode on the following interfaces:

- IQ2 and IQ2-E Gigabit Ethernet interfaces
- Ethernet interfaces on the MX Series routers

JUNOS Software can place a remote DTE into loopback mode (if remote-loopback mode is supported by the remote DTE). When you place a remote DTE into loopback mode, the interface receives the remote-loopback request and puts the interface into remote-loopback mode. When the interface is in remote-loopback mode, all frames except OAM PDUs are looped back without any changes made to the frames. OAM PDUs continue to be sent to the management plane and processed.

To configure remote loopback, include the `remote-loopback` statement at the `[edit protocol oam ethernet link-fault-management interface interface-name]` hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name]  
  remote-loopback;
```

To take the remote DTE out of loopback mode, remove the `remote-loopback` statement from the configuration.

Enabling Remote Loopback Support on the Local Interface

You can allow a remote DTE to set a local interface into remote loopback mode on IQ2 and IQ2-E Gigabit Ethernet interfaces and all Ethernet interfaces on the MX Series routers. When a remote-loopback request is sent by a remote DTE, the JUNOS Software places the local interface into loopback mode. When an interface is in loopback mode, all frames except OAM PDUs are looped back without any changes to the frames. OAM PDUs continue to be sent to the management plane and processed. By default, the remote loopback feature is not enabled.

To enable remote loopback, include the `allow-remote-loopback` statement at the `[edit protocol oam ethernet link-fault-management interface interface-name negotiation-options]` hierarchy level:

```
[edit protocol oam ethernet link-fault-management interface interface-name
 negotiation-options]
allow-remote-loopback;
```



NOTE: Activation of OAM remote loopback may result in data frame loss.

Example: Configuring IEEE 802.3ah OAM Support on an Interface

Configure 802.3ah OAM support on an MX Series 10-Gigabit Ethernet interface:

```
[edit]
protocols {
  oam {
    ethernet {
      link-fault-management {
        interface xe-0/0/0 {
          link-discovery active;
          pdu-interval 800;
          pdu-threshold 4;
          remote-loopback;
          negotiation-options {
            allow-remote-loopback;
          }
          event-thresholds {
            frame-error 30;
            frame-period 50;
            frame-period summary 40;
            symbol-period 20;
          }
        }
      }
    }
  }
}
```


Chapter 47

Configuring VRRP and VRRP for IPv6

This section contains the following topics:

- VRRP and VRRP for IPv6 Overview on page 753
- Configuring VRRP and VRRP for IPv6 on page 753

VRRP and VRRP for IPv6 Overview

For Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, 10-Gigabit Ethernet, and Ethernet logical interfaces, you can configure the Virtual Router Redundancy Protocol (VRRP) and VRRP for IPv6. VRRP and VRRP for IPv6 allow hosts on a LAN to make use of redundant routers on that LAN without requiring more than the static configuration of a single default route on the hosts. The VRRP routers share the IP address corresponding to the default route configured on the hosts. At any time, one of the VRRP routers is the master (active) and the others are backups. If the master fails, one of the backup routers becomes the new master router, thus always providing a virtual default router and allowing traffic on the LAN to be routed without relying on a single router.

VRRP is defined in RFC 3768, *Virtual Router Redundancy Protocol*.

For VRRP and VRRP for IPv6 overview information, configuration guidelines, and statement summaries, see the *JUNOS High Availability Configuration Guide*.

Configuring VRRP and VRRP for IPv6

To configure VRRP or VRRP for IPv6, include the `vrrp-group` or `vrrp-inet6-group` statement, respectively. These statements are available at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* family inet address *address*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family inet address *address*]

The VRRP and VRRP IPv6 configuration statements are as follows:

```
(vrrp-group | vrrp-inet6-group) group-number {  
  (accept-data | no-accept-data);  
  advertise-interval seconds;  
  authentication-key key;
```

```

authentication-type authentication;
fast-interval milliseconds;
(preempt | no-preempt) {
    hold-time seconds;
}
priority-number number;
track {
    priority-hold-time;
    interface interface-name {
        priority-cost priority;
        bandwidth-threshold bits-per-second {
            priority-cost;
        }
    }
}
virtual-address [ addresses ];
}

```

To trace VRRP and VRRP for IPv6 operations, include the `traceoptions` statement at the `[edit protocols vrrp]` hierarchy level:

```

[edit protocols vrrp]
traceoptions {
    file <filename> <files number <match regular-expression <microsecond-stamp>
    <size size> <world-readable | no-world-readable>;
    flag flag;
    no-remote-trace;
}

```

When there are multiple VRRP groups, there is a few seconds delay between the time the first gratuitous ARP is sent out and the rest of the gratuitous ARP are sent. Configuring `failover-delay` compensates for this delay. To configure the failover delay from 500 to 2000 milliseconds for VRRP and VRRP for IPv6 operations, use the `set failover-delay milliseconds` statement at the `[edit protocols vrrp]` hierarchy level:

```

[edit protocols vrrp]
set failover-delay milliseconds;

```

To configure the startup period for VRRP and VRRP for IPv6 operations, include the `startup-silent-period` statement at the `[edit protocols vrrp]` hierarchy level:

```

[edit protocols vrrp]
startup-silent-period seconds;

```

Chapter 48

Configuring Gigabit Ethernet Accounting and Policing

This section contains the following topics:

- Gigabit Ethernet Accounting and Policing Overview on page 755
- Configuring Gigabit Ethernet Policers on page 757
- Configuring Gigabit Ethernet Two-Color and Tricolor Policers on page 763
- Configuring MAC Address Accounting on page 766

Gigabit Ethernet Accounting and Policing Overview

For Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can configure granular per-VLAN class-of-service (CoS) capabilities and extensive instrumentation and diagnostics on a per-VLAN and per-MAC address basis.

VLAN rewrite, tagging, and deleting enables you to use VLAN address space to support more customers and services.

VPLS allows you to provide a point-to-multipoint LAN between a set of sites in a VPN. Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) are combined with VPLS to deliver metro Ethernet service.

For Gigabit Ethernet IQ2 and IQ2-E and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, you can apply Layer 2 policing to logical interfaces in the egress or ingress direction. Policers are configured at the [edit firewall] hierarchy level.

Table 67 on page 755 lists the capabilities of Gigabit Ethernet IQ PICs and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

Table 67: Capabilities of Gigabit Ethernet IQ and Gigabit Ethernet with SFPs

Capability	Gigabit Ethernet IQ (SFP)	Gigabit Ethernet (SFP)
Layer 2		
802.3ad link aggregation	Yes	Yes

Table 67: Capabilities of Gigabit Ethernet IQ and Gigabit Ethernet with SFPs (continued)

Capability	Gigabit Ethernet IQ (SFP)	Gigabit Ethernet (SFP)
Maximum VLANs per port	384	1023
Maximum transmission unit (MTU) size	9192	9192
MAC learning	Yes	Yes
MAC accounting	Yes	Yes
MAC filtering	Yes	Yes
Destinations per port	960	960
Sources per port	64	64
Hierarchical MAC policers	Yes, premium and aggregate	No, aggregate only
Multiple TPID support and IP service for nonstandard TPIDs	Yes	Yes
Multiple Ethernet encapsulations	Yes	Yes
Dual VLAN tags	Yes	No
VLAN rewrite	Yes	No
Layer 2 VPNs		
VLAN CCC	Yes	Yes
Port-based CCC	Yes	Yes
Extended VLAN CCC Virtual Metropolitan Area Network (VMAN) Tag Protocol	Yes	Yes
CoS		
PIC-based egress queues	Yes	Yes
Queued VLANs	Yes	No
VPLS	Yes	Yes

For more information about configuring VPLS, see the *JUNOS VPNs Configuration Guide* and the *JUNOS Feature Guide*.

You can also configure CoS on logical IQ interfaces. For more information, see the *JUNOS Class of Service Configuration Guide*.

Configuring Gigabit Ethernet Policers

On Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can define rate limits for premium and aggregate traffic received on the interface. These policers allow you to perform simple traffic policing without configuring a firewall filter. First you configure the Ethernet policer profile, next you classify ingress and egress traffic, then you can apply the policer to a logical interface.

For Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), the policer rates you configure can be different than the rates on the Packet Forward Engine. The difference results from Layer 2 overhead. The PIC accounts for this difference.



NOTE:

On MX Series routers with Gigabit Ethernet or Fast Ethernet PICs, the following considerations apply:

- Interface counters do not count the 7-byte preamble and 1-byte frame delimiter in Ethernet frames.
- In MAC statistics, the frame size includes MAC header and CRC before any VLAN rewrite/imposition rules are applied.
- In traffic statistics, the frame size encompasses the L2 header without CRC after any VLAN rewrite/imposition rule.

For information on understanding Ethernet frame statistics, see the *MX Series Layer 2 Configuration Guide*.

This section contains the following topics:

- Configuring a Policer on page 757
- Specifying an Input Priority Map on page 758
- Specifying an Output Priority Map on page 759
- Applying a Policer on page 759
- Configuring MAC Address Filtering on page 761
- Example: Configuring Gigabit Ethernet Policers on page 762

Configuring a Policer

To configure an Ethernet policer profile, include the `ethernet-policer-profile` statement at the [edit interfaces *interface-name* *gigether-options* *ethernet-switch-profile*] hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-switch-profile]
ethernet-policer-profile {
```

```

    policer cos-policer-name {
        aggregate {
            bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
            burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
        }
        premium {
            bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
            burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
        }
    }
}

```

In the Ethernet policer profile, the aggregate-priority policer is mandatory; the premium-priority policer is optional.

For aggregate and premium policers, you specify the bandwidth limit in bits per second. You can specify the value as a complete decimal number or as a decimal number followed by the abbreviation **k** (1000), **m** (1,000,000), or **g** (1,000,000,000). There is no absolute minimum value for bandwidth limit, but any value below 61,040 bps will result in an effective rate of 30,520 bps. The maximum bandwidth limit is 4.29 Gbps.

The maximum burst size controls the amount of traffic bursting allowed. To determine the burst-size limit, you can multiply the bandwidth of the interface on which you are applying the filter by the amount of time you allow a burst of traffic at that bandwidth to occur:

$$\text{burst size} = \text{bandwidth} \times \text{allowable time for burst traffic}$$

If you do not know the interface bandwidth, you can multiply the maximum MTU of the traffic on the interface by 10 to obtain a value. For example, the burst size for an MTU of 4700 would be 47,000 bytes. The burst size should be at least 10 interface MTUs. The maximum value for the burst-size limit is 100 MB.

Specifying an Input Priority Map

An input priority map identifies ingress traffic with specified IEEE 802.1p priority values, and classifies that traffic as premium.

If you include a premium-priority policer, you can specify an input priority map by including the `ieee802.1p premium` statement at the `[edit interfaces interface-name gigether-options ethernet-policer-profile input-priority-map]` hierarchy level:

```

[edit interfaces interface-name gigether-options ethernet-policer-profile input-priority-map]
ieee802.1p premium [ values ];

```

The priority values can be from 0 through 7. The remaining traffic is classified as nonpremium (or aggregate). For a configuration example, see “Example: Configuring Gigabit Ethernet Policers” on page 762.



NOTE: On IQ2 and IQ2-E interfaces and MX Series interfaces, when a VLAN tag is pushed, the inner VLAN IEEE 802.1p bits are copied to the IEEE bits of the VLAN or VLANs being pushed. If the original packet is untagged, the IEEE bits of the VLAN or VLANs being pushed are set to 0.

Specifying an Output Priority Map

An output priority map identifies egress traffic with specified queue classification and packet loss priority (PLP), and classifies that traffic as premium.

If you include a premium-priority policer, you can specify an output priority map by including the `classifier` statement at the `[edit interfaces interface-name gigether-options ethernet-policer-profile output-priority-map]` hierarchy level:

```
[edit interfaces interface-name gigether-options ethernet-policer-profile
 output-priority-map]
classifier {
  premium {
    forwarding-class class-name {
      loss-priority (high | low);
    }
  }
}
```

You can define a forwarding class, or you can use a predefined forwarding class. Table 68 on page 759 shows the predefined forwarding classes and their associated queue assignments.

Table 68: Default Forwarding Classes

Forwarding Class Name	Queue
best-effort	Queue 0
expedited-forwarding	Queue 1
assured-forwarding	Queue 2
network-control	Queue 3

For more information about CoS forwarding classes, see the *JUNOS Class of Service Configuration Guide*. For a configuration example, see “Example: Configuring Gigabit Ethernet Policers” on page 762.

Applying a Policer

On all MX Series Router interfaces, Gigabit Ethernet IQ, IQ2, and IQ2-E PICs, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can apply input and output policers that define rate limits for premium and aggregate traffic received on the

logical interface. Aggregate policers are supported on Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router).

These policers allow you to perform simple traffic policing without configuring a firewall filter. For information about defining these policers, see “Configuring Gigabit Ethernet Policers” on page 757.

To apply policers to specific source MAC addresses, include the `accept-source-mac` statement:

```
accept-source-mac {
  mac-address mac-address {
    policer {
      input cos-policer-name;
      output cos-policer-name;
    }
  }
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

You can specify the MAC address as *nn:nn:nn:nn:nn:nn* or *nnnn.nnnn.nnnn*, where *n* is a hexadecimal number. You can configure up to 64 source addresses. To specify more than one address, include multiple `mac-address` statements in the logical interface configuration.



NOTE: On untagged Gigabit Ethernet interfaces you should not configure the `source-address-filter` statement at the [edit interfaces *ge-fpc/pic/port* *gigether-options*] hierarchy level and the `accept-source-mac` statement at the [edit interfaces *ge-fpc/pic/port* *gigether-options* unit *logical-unit-number*] hierarchy level simultaneously. If these statements are configured for the same interfaces at the same time, an error message is displayed.

On tagged Gigabit Ethernet interfaces you should not configure the `source-address-filter` statement at the [edit interfaces *ge-fpc/pic/port* *gigether-options*] hierarchy level and the `accept-source-mac` statement at the [edit interfaces *ge-fpc/pic/port* *gigether-options* unit *logical-unit-number*] hierarchy level with an identical MAC address specified in both filters. If these statements are configured for the same interfaces with an identical MAC address specified, an error message is displayed.



NOTE: If the remote Ethernet card is changed, the interface does not accept traffic from the new card because the new card has a different MAC address.

The MAC addresses you include in the configuration are entered into the router's MAC database. To view the router's MAC database, enter the **show interfaces mac-database *interface-name*** command:

```
user@host> show interfaces mac-database interface-name
```

In the **input** statement, list the name of one policer template to be evaluated when packets are received on the interface.

In the **output** statement, list the name of one policer template to be evaluated when packets are transmitted on the interface.



NOTE: On IQ2 and IQ2-E PIC interfaces, the default value for maximum retention of entries in the MAC address table has changed, for cases in which the table is not full. The new holding time is 12 hours. The previous retention time of 3 minutes is still in effect when the table is full.

You can use the same policer one or more times.

If you apply both policers and firewall filters to an interface, input policers are evaluated before input firewall filters, and output policers are evaluated after output firewall filters.

Configuring MAC Address Filtering

You cannot explicitly define traffic with specific source MAC addresses to be rejected; however, for Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can block all incoming packets that do not have a source address specified in the **accept-source-mac** statement. For more information about the **accept-source-mac** statement, see “Applying a Policer” on page 759.

To enable this blocking, include the **source-filtering** statement at the [edit interfaces *interface-name* *gigether-options*] hierarchy level:

```
[edit interfaces interface-name gigether-options]  
source-filtering;
```

For more information about the **source-filtering** statement, see “Enabling Ethernet MAC Address Filtering” on page 591.

To accept traffic even though it does not have a source address specified in the **accept-source-mac** statement, include the **no-source-filtering** statement at the [edit interfaces *interface-name* *gigether-options*] hierarchy level:

```
[edit interfaces interface-name gigether-options]  
no-source-filtering;
```

For more information about the **accept-source-mac** statement, see “Applying a Policer” on page 759.

Example: Configuring Gigabit Ethernet Policers

Configure interface `ge-6/0/0` to treat priority values 2 and 3 as premium. On ingress, this means that IEEE 802.1p priority values 2 and 3 are treated as premium. On egress, it means traffic that is classified into queue 0 or 1 with PLP of low and queue 2 or 3 with PLP of high, is treated as premium.

Define a policer that limits the premium bandwidth to 100 Mbps and burst size to 3 k, and the aggregate bandwidth to 200 Mbps and burst size to 3 k.

Specify that frames received from the MAC address `00:01:02:03:04:05` and the VLAN ID 600 are subject to the policer on input and output. On input, this means frames received with the source MAC address `00:01:02:03:04:05` and the VLAN ID 600 are subject to the policer. On output, this means frames transmitted from the router with the destination MAC address `00:01:02:03:04:05` and the VLAN ID 600 are subject to the policer.

```
[edit interfaces]
ge-6/0/0 {
  gigether-options {
    ether-switch-profile {
      ether-policer-profile {
        input-priority-map {
          ieee-802.1p {
            premium [ 2 3 ];
          }
        }
        output-priority-map {
          classifier {
            premium {
              forwarding-class best-effort {
                loss-priority low;
              }
              forwarding-class expedited-forwarding {
                loss-priority low;
              }
              forwarding-class assured-forwarding {
                loss-priority high;
              }
              forwarding-class network-control {
                loss-priority high;
              }
            }
          }
        }
      }
    }
  }
  policer policer-1 {
    premium {
      bandwidth-limit 100m;
      burst-size-limit 3k;
    }
    aggregate {
      bandwidth-limit 200m;
      burst-size-limit 3k;
    }
  }
}
```

```

    }
  }
}
unit 0 {
  accept-source-mac {
    mac-address 00:01:02:03:04:05 {
      policer {
        input policer-1;
        output policer-1;
      }
    }
  }
}
}

```

Configuring Gigabit Ethernet Two-Color and Tricolor Policers

For Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers, you can configure two-color and tricolor marking policers and apply them to logical interfaces to prevent traffic on the interface from consuming bandwidth inappropriately.

Networks police traffic by limiting the input or output transmission rate of a class of traffic on the basis of user-defined criteria. Policing traffic allows you to control the maximum rate of traffic sent or received on an interface and to partition a network into multiple priority levels or classes of service.

Policers require you to apply a burst size and bandwidth limit to the traffic flow, and set a consequence for packets that exceed these limits—usually a higher loss priority, so that packets exceeding the policer limits are discarded first.

Juniper Networks router architectures support three types of policer:

- Two-color policer—A two-color policer (or “policer” when used without qualification) meters the traffic stream and classifies packets into two categories of packet loss priority (PLP) according to a configured bandwidth and burst-size limit. You can mark packets that exceed the bandwidth and burst-size limit in some way, or simply discard them. A policer is most useful for metering traffic at the port (physical interface) level.
- Single-rate tricolor marking (srTCM)—A single-rate tricolor marking policer is defined in RFC 2697, *A Single Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services (DiffServ) environment. This type of policer meters traffic based on the configured committed information rate (CIR), committed burst size (CBS), and excess burst size (EBS). Traffic is marked as belonging to one of three categories (green, yellow, or red) based on whether the packets arriving are below the CBS (green), exceed the CBS (yellow) but not the EBS, or exceed the EBS (red). Single-rate TCM is most useful when a service is structured according to packet length and not peak arrival rate.
- Two-rate Tricolor Marking (trTCM)—This type of policer is defined in RFC 2698, *A Two Rate Three Color Marker*, as part of an assured forwarding (AF) per-hop-behavior (PHB) classification system for a Differentiated Services

(DiffServ) environment. This type of policer meters traffic based on the configured CIR and peak information rate (PIR), along with their associated burst sizes, the CBS and EBS. Traffic is marked as belonging to one of three categories (green, yellow, or red) based on whether the packets arriving are below the CIR (green), exceed the CIR (yellow) but not the PIR, or exceed the PIR (red). Two-rate TCM is most useful when a service is structured according to arrival rates and not necessarily packet length.

Unlike policing (described in “Configuring Gigabit Ethernet Policers” on page 757), configuring two-color policers and tricolor marking policers requires that you configure a firewall filter.

This section contains the following topics:

- Configuring a Policer on page 764
- Applying a Policer on page 765
- Example: Configuring and Applying a Policer on page 765

Configuring a Policer

Two-color and tricolor marking policers are configured at the [edit firewall] hierarchy level.

A tricolor marking policer polices traffic on the basis of metering rates, including the CIR, the PIR, their associated burst sizes, and any policing actions configured for the traffic.

To configure tricolor policer marking, include the `three-color-policer` statement with options at the [edit firewall] hierarchy level:

```
[edit firewall]
three-color-policer name {
  action {
    loss-priority high {
      then discard;
    }
  }
  single-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    excess-burst-size bytes;
  }
  two-rate {
    (color-aware | color-blind);
    committed-information-rate bps;
    committed-burst-size bytes;
    peak-information-rate bps;
    peak-burst-size bytes;
  }
}
```

For more information about configuring tricolor policer markings, see the *JUNOS Policy Framework Configuration Guide* and the *JUNOS Class of Service Configuration Guide*.

Applying a Policer

Apply a two-color policer or tricolor policer to a logical interface to prevent traffic on the interface from consuming bandwidth inappropriately. To apply two-color or tricolor policers, include the `layer2-policer` statement:

```
layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    policer-name;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Use the `input-policer` statement to apply a two-color policer to received packets on a logical interface and the `input-three-color` statement to apply a tricolor policer. Use the `output-policer` statement to apply a two-color policer to transmitted packets on a logical interface and the `output-three-color` statement to apply a tricolor policer. The specified policers must be configured at the [edit firewall] hierarchy level. For each interface, you can configure a three-color policer or two-color input policer or output policers—you cannot configure both a three-color policer and a two-color policer.

Example: Configuring and Applying a Policer

Configure tricolor policers and apply them to an interface:

```
[edit firewall]
three-color-policer three-color-policer-color-blind {
    logical-interface-policer;
    two-rate {
        color-blind;
        committed-information-rate 1500000;
        committed-burst-size 150;
        peak-information-rate 3;
        peak-burst-size 300;
    }
}
three-color-policer three-color-policer-color-aware {
    logical-interface-policer;
    two-rate {
        color-aware;
        committed-information-rate 1500000;
        committed-burst-size 150;
        peak-information-rate 3;
    }
}
```

```

        peak-burst-size 300;
    }
}
[edit interfaces ge-1/1/0]
unit 1 {
    layer2-policer {
        input-three-color three-color-policer-color-blind;
        output-three-color three-color-policer-color-aware;
    }
}

```

Configure a two-color policer and apply it to an interface:

```

[edit firewall]
policer two-color-policer {
    logical-interface-policer;
    if-exceeding {
        bandwidth-percent 90;
        burst-size-limit 300;
    }
    then loss-priority-high;
}
[edit interfaces ge-1/1/0]
unit 2 {
    layer2-policer {
        input-policer two-color-policer;
        output-policer two-color-policer;
    }
}

```

Configuring MAC Address Accounting

For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), you can configure whether source and destination MAC addresses are dynamically learned. To configure MAC address accounting, include the `mac-learn-enable` statement at the `[edit interfaces interface-name gigether-options ethernet-switch-profile]` hierarchy level:

```

[edit interfaces interface-name gigether-options ethernet-switch-profile]
mac-learn-enable;

```

To prohibit the interface from dynamically learning source and destination MAC addresses, include the `no-mac-learn-enable` statement at the `[edit interfaces interface-name gigether-options ethernet-switch-profile]` hierarchy level:

```

[edit interfaces interface-name gigether-options ethernet-switch-profile]
no-mac-learn-enable;

```

MAC address learning is based on source addresses. You can start accounting for traffic after it has been sent from the MAC address. Once the MAC address is learned, the frames and bytes transmitted to or received from the MAC address can be tracked.

Chapter 49

Configuring Gigabit Ethernet Autonegotiation

This section contains the following topics:

- Gigabit Ethernet Autonegotiation Overview on page 767
- Configuring Gigabit Ethernet Autonegotiation on page 767

Gigabit Ethernet Autonegotiation Overview

Autonegotiation is enabled by default on all Gigabit Ethernet and Tri-Rate Ethernet copper interfaces. However, you can explicitly enable autonegotiation to configure remote fault options manually.



NOTE: For Gigabit Ethernet interfaces installed in J4350 and J6350 Services Routers, when you manually configure either the link mode or speed settings, the system ignores the configuration and generates a system log message. When autonegotiation is enabled and you specify the link mode and speed, the link autonegotiates with the manually configured settings. When autonegotiation is disabled and you configure both the link mode and speed, the link operates with the manually configured settings. If you disable autonegotiation and do not manually configure the link mode and speed, the link operates at 1000 Mbps full duplex.



NOTE: When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.

Configuring Gigabit Ethernet Autonegotiation

Configuring Gigabit Ethernet Autonegotiation with Remote Fault

To configure explicit autonegotiation and remote fault, include the `auto-negotiation` statement and the `remote-fault` option at the `[edit interfaces ge-fpc/pic/port gigerether-options]` hierarchy level.

```
[edit interfaces ge-fpc/pic/port gigerether-options]
```

(auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online | local-interface-offline>

Configuring Flow Control

To enable flow control, include the **flow-control** statement at the [edit interfaces *ge-fpc/pic*]/*port* **gigether-options**] hierarchy level. For more information, see “Configuring Flow Control” on page 594

Configuring Autonegotiation Speed on MX Series Routers

MX Series routers with Combo Line Rate DPCs and Tri-Rate Copper SFPs support autonegotiation of speed. The autonegotiation specified interface speed is propagated to CoS, routing protocols, and other system components. Half-duplex mode is not supported.

To specify the autonegotiation speed, use the **speed** (**auto** | **1Gbps** | **100Mbps** | **10Mbps**) statement at the [edit interfaces *ge-fpc/pic/port*] hierarchy level.

To set port speed negotiation to a specific rate, set the port speed to **1Gbps**, **100Mbps**, or **10Mbps**. If the negotiated speed and the interface speed do not match, the link will not be brought up.

If you set the autonegotiation speed **auto** option, then the port speed is negotiated.

You can disable auto MDI/MDIX using the **no-auto-mdix** statement at the [edit interfaces *ge-fpc/pic/port* **gigether-options**] hierarchy level.

Use the **show interfaces *ge-fpc/pic/port* brief** command to display the auto negotiation of speed and auto MDI/MDIX states.

Displaying Autonegotiation Status

To display Gigabit Ethernet interface details, including the autonegotiation status, use the operational mode command **show interfaces *ge-fpc/pic/port* extensive**.

Table 69 on page 768 and Table 70 on page 770 provide information about the autonegotiation status on local and remote routers with fiber interfaces. The status of the link and LED can vary depending on the level of autonegotiation set and the transmit and receive fiber status.

Table 69: Mode and Autonegotiation Status (Local)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	ON	Default	Green	UP	Complete
ON	OFF	Default	Red	DOWN	
OFF	ON	Default	Red	DOWN	

Table 69: Mode and Autonegotiation Status (Local) *(continued)*

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
OFF	OFF	Default	Red	DOWN	
ON	ON	Default	Red	DOWN	
ON	ON	Default	Green	UP	No-autonegotiation
ON	OFF	Default	Red	DOWN	
OFF	OFF	Default	Red	DOWN	
ON	ON	Default	Green	UP	
ON	ON	Default	Red	DOWN	
ON	ON	No-autonegotiation	Green	UP	Incomplete
ON	OFF	No-autonegotiation	Red	DOWN	
OFF	ON	No-autonegotiation	Green	UP	
OFF	OFF	No-autonegotiation	Red	DOWN	
ON	ON	No-autonegotiation	Red	DOWN	
ON	ON	Explicit	Green	UP	Complete
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Red	DOWN	
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit	Red	DOWN	
ON	ON	Explicit	Green	UP	No-autonegotiation
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Green	UP	
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit	Red	DOWN	
ON	ON	Explicit + RFI-Offline	Green	UP	Complete
OFF	ON	Explicit + RFI-Offline	Red	DOWN	
OFF	OFF	Explicit + RFI-Offline	Red	DOWN	
ON	ON	Explicit + RFI-Offline	Red	DOWN	
ON	ON	Explicit + RFI-Offline	Green	UP	No-autonegotiation

Table 69: Mode and Autonegotiation Status (Local) *(continued)*

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	OFF	Explicit + RFI-Offline	Red	DOWN	
OFF	ON	Explicit + RFI-Offline	Green	UP	
OFF	OFF	Explicit + RFI-Offline	Red	DOWN	
ON	ON	Explicit + RFI-Offline	Red	DOWN	
ON	ON	Explicit + RFI-Offline	Red	DOWN	Complete
ON	OFF	Explicit + RFI-Offline	Red	DOWN	
OFF	ON	Explicit + RFI-Online	Red	DOWN	
OFF	OFF	Explicit + RFI-Online	Red	DOWN	
ON	ON	Explicit + RFI-Online	Red	DOWN	
ON	ON	Explicit + RFI-Online	Green	UP	No-autonegotiation*
ON	OFF	Explicit + RFI-Online	Red	DOWN	
OFF	ON	Explicit + RFI-Online	Green	UP	
OFF	OFF	Explicit + RFI-Online	Red	DOWN	
ON	ON	Explicit + RFI-Online	Green	UP	
ON	ON	Explicit + RFI-Online	Red	DOWN	
ON	ON	Explicit + RFI-Online	Red	DOWN	Complete
ON	OFF	Explicit + RFI-Online	Red	DOWN	
OFF	ON	Explicit + RFI-Online	Red	DOWN	
OFF	OFF	Explicit + RFI-Online	Red	DOWN	
ON	ON	Explicit + RFI-Online	Red	DOWN	
ON	ON	Explicit + RFI-Online	Green	UP	Complete

Table 70: Mode and Autonegotiation Status (Remote)

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	ON	Default	Green	UP	Complete

Table 70: Mode and Autonegotiation Status (Remote) *(continued)*

Transmit	Receive	Mode	LED	Link	Autonegotiation Status
ON	ON	Default	Red	DOWN	
ON	OFF	Default	Red	DOWN	
OFF	ON	Default	Red	DOWN	
OFF	OFF	Default	Red	DOWN	
ON	ON	No-autonegotiation	Green	UP	Incomplete
ON	ON	No-autonegotiation	Red	DOWN	
ON	OFF	No-autonegotiation	Red	DOWN	
OFF	ON	No-autonegotiation	Green	UP	
OFF	OFF	No-autonegotiation	Red	DOWN	
ON	ON	Explicit	Green	UP	Complete
ON	ON	Explicit	Red	DOWN	
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Red	DOWN	
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit	Red	DOWN	Complete
ON	OFF	Explicit	Red	DOWN	
OFF	ON	Explicit	Red	DOWN	
OFF	OFF	Explicit	Red	DOWN	
ON	ON	Explicit + RFI-Offline	Red	DOWN	Complete
ON	OFF	Explicit + RFI-Offline	Red	DOWN	
OFF	ON	Explicit + RFI-Offline	Red	DOWN	
OFF	OFF	Explicit + RFI-Offline	Red	DOWN	
ON	ON	Explicit + RFI-Online	Green	UP	Complete
ON	ON	Explicit + RFI-Online	Red	DOWN	
ON	OFF	Explicit + RFI-Online	Red	DOWN	
OFF	ON	Explicit + RFI-Online	Red	DOWN	
OFF	OFF	Explicit + RFI-Online	Red	DOWN	

Chapter 50

Configuring Gigabit Ethernet OTN Options

This section contains the following topics:

- Gigabit Ethernet OTN Options Configuration Overview on page 773
- Gigabit Ethernet OTN Options on page 773

Gigabit Ethernet OTN Options Configuration Overview

M120, M320, T320, T640, and T1600 router platforms support Optical Transport Network (OTN) interfaces, including the 10-Gigabit Ethernet DWDM OTN PIC, and provide ITU-G.709 support. Use the `set otn-options` statement at the `[edit interfaces if-fpc/pic/port]` hierarchy level to configure the OTN options.

Gigabit Ethernet OTN Options

The following example shows the configuration settings for Gigabit Ethernet OTN options:

```
[edit interfaces ge-fpc/pic/port]
otn-options {
  fec (efec | gfec | none);
  (laser-enable | no-laser-enable);
  (line-loopback | no-line-loopback);
  pass-thru;
  rate (fixed-stuff-bytes | no-fixed-stuff-bytes | pass-thru);
  trigger (oc-lof | oc-lom | oc-los | oc-wavelength-lock | odu-ais | odu-bbe-th | odu-bdi |
    odu-es-th | odu-lck | odu-oci | odu-sd | odu-ses-th | odu-ttim | odu-uas-th | opu-ptm
    | otu-ais | otu-bbe-th | otu-bdi | otu-es-th | otu-fec-deg | otu-fec-exe | otu-iae | otu-sd
    | otu-ses-th | otu-ttim | otu-uas-th);
  tti;
}
```



NOTE: The Gigabit Ethernet interface and the XENPAK interface support the read/write overhead bytes only for the APS/PPC (bytes 0 through 3).

You can use the following show commands to view the OTN configuration:

- `show interfaces extensive`—See the *JUNOS Interfaces Command Reference* for command details.

- **show chassis hardware**—See the *JUNOS System Basics and Services Command Reference* for command details.
- **show chassis pic**—See the *JUNOS System Basics and Services Command Reference* for command details.

Chapter 51

Configuring the Management Ethernet Interface

This section contains the following topics:

- Management Ethernet Interface Overview on page 775
- Configuring a Consistent Management IP Address on page 775
- Configuring the MAC Address on the Management Ethernet Interface on page 777

Management Ethernet Interface Overview

The router's management Ethernet interface, `fxp0` or `em0`, is an out-of-band management interface that needs to be configured only if you want to connect to the router through the management port on the front of the router. You can configure an IP address and prefix length for this interface, which you commonly do when you first install the JUNOS Software:

```
[edit]
user@host# set interfaces (fxp0 | em0) unit 0 family inet address/prefix-length
[edit]
user@host# show
interfaces {
  (fxp0 | em0) {
    unit 0 {
      family inet {
        address/prefix-length;
      }
    }
  }
}
```

Configuring a Consistent Management IP Address

On routers with multiple Routing Engines, each Routing Engine is configured with a separate IP address for the management Ethernet interface. To access the master Routing Engine, you must know which Routing Engine is active and use the appropriate IP address.

Optionally, for consistent access to the master Routing Engine, you can configure an additional IP address and use this address for the management interface regardless of which Routing Engine is active. This additional IP address is active only on the

management Ethernet interface for the master Routing Engine. During switchover, the address moves to the new master Routing Engine.



NOTE: For M Series, MX Series, and most T Series routers, the management Ethernet interface is **fxp0**. For TX Matrix Plus routers and T1 600 routers configured in a routing matrix, the management Ethernet interface is **em0**.



NOTE: Automated scripts that you have developed for standalone T1 600 routers (T1 600 routers that are not in a routing matrix) might contain references to the **fxp0** management Ethernet interface. Before reusing the scripts on T1 600 routers in a routing matrix, edit the command lines that reference the **fxp0** management Ethernet interface so that the commands reference the **em0** management Ethernet interface instead.

To configure an additional IP address for the management Ethernet interface, include the **master-only** statement at the **[edit groups]** hierarchy level.

In the following example, IP address **10.17.40.131** is configured for both Routing Engines and includes a **master-only** statement. With this configuration, the **10.17.40.131** address is active only on the master Routing Engine. The address remains consistent regardless of which Routing Engine is active. IP address **10.17.40.132** is assigned to **fxp0** on **re0**, and address **10.17.40.133** is assigned to **fxp0** on **re1**.

```
[edit groups re0 interfaces fxp0]
unit 0 {
  family inet {
    address 10.17.40.131/25 {
      master-only;
    }
    address 10.17.40.132/25;
  }
}
[edit groups re1 interfaces fxp0]
unit 0 {
  family inet {
    address 10.17.40.131/25 {
      master-only;
    }
    address 10.17.40.133/25;
  }
}
```

This feature is available on all routers that include dual Routing Engines. On the TX Matrix router, this feature is applicable to the switch-card chassis (SCC) only.

Configuring the MAC Address on the Management Ethernet Interface

By default, the router's management Ethernet interface uses as its MAC address the MAC address that is burned into the Ethernet card.



NOTE: For M Series, MX Series, and most T Series routers, the management Ethernet interface is `fxp0`. For TX Matrix Plus routers and T1600 routers configured in a routing matrix, the management Ethernet interface is `em0`.



NOTE: Automated scripts that you have developed for standalone T1600 routers (T1600 routers that are not in a routing matrix) might contain references to the `fxp0` management Ethernet interface. Before reusing the scripts on T1600 routers in a routing matrix, edit the command lines that reference the `fxp0` management Ethernet interface so that the commands reference the `em0` management Ethernet interface instead.

To display the MAC address used by the router's management Ethernet interface, enter the `show interface fxp0` or `show interface em0` operational mode command.

To change the management Ethernet interface's MAC address, include the `mac` statement at the `[edit interfaces fxp0]` or `[edit interfaces em0]` hierarchy level:

```
[edit interfaces (fxp0 | em0)]
mac mac-address;
```

Specify the MAC address as six hexadecimal bytes in one of the following formats: `nnnn.nnnn.nnnn` (for example, `0011.2233.4455`) or `nn:nn:nn:nn:nn:nn` (for example, `00:11:22:33:44:55`).



NOTE: If you integrate a standalone T640 router into a routing matrix, the PIC MAC addresses for the integrated T640 router are derived from a pool of MAC addresses maintained by the TX Matrix router. For each MAC address you specify in the configuration of a formerly standalone T640 router, you must specify the same MAC address in the configuration of the TX Matrix router.

Similarly, if you integrate a standalone T1600 router into a routing matrix, the PIC MAC addresses for the integrated T1600 router are derived from a pool of MAC addresses maintained by the TX Matrix Plus router. For each MAC address you specify in the configuration of a formerly standalone T1600 router, you must specify the same MAC address in the configuration of the TX Matrix Plus router.

Chapter 52

Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength

This section contains the following topics:

- 10-Gigabit Ethernet DWDM Interface Wavelength Overview on page 779
- Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength on page 779

10-Gigabit Ethernet DWDM Interface Wavelength Overview

For M320, M120, T320, and T640 routers, the 10-Gigabit Ethernet DWDM PIC enables you to configure 10-Gigabit Ethernet DWDM interfaces with full C-band International Telecommunication Union (ITU)-Grid tunable optics, as defined in the following specifications:

- *Intel TXN13600 Optical Transceiver I2C Interface and Customer EEPROM Preliminary Specification*, July 2004.
- *I2C Reference Document for 300 Pin MSA 10G and 40G Transponder*, Edition 4, August 04, 2003.

By default, the wavelength is 1550.12 nanometers (nm), which corresponds to 193.40 terahertz (THz).

Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength

To configure the wavelength on a 10-Gigabit Ethernet DWDM interface, include the `wavelength` statement at the [edit interfaces *ge-fpc/pic/port* optics-options] hierarchy level:

```
[edit interfaces ge-0/0/0 optics-options]
wavelength nm;
```

For interface diagnostics, you can issue the `show interfaces diagnostics optics ge-fpc/pic/port` operational mode command.

Table 71 on page 780 shows configurable wavelengths and the corresponding frequency for each configurable wavelength.

Table 71: Wavelength-to-Frequency Conversion Matrix

Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)	Wavelength (nm)	Frequency (THz)
1528.77	196.10	1540.56	194.60	1552.52	193.10
1529.55	196.00	1541.35	194.50	1553.33	193.00
1530.33	195.90	1542.14	194.40	1554.13	192.90
1531.12	195.80	1542.94	194.30	1554.94	192.80
1531.90	195.70	1543.73	194.20	1555.75	192.70
1532.68	195.60	1544.53	194.10	1556.56	192.60
1533.47	195.50	1545.32	194.00	1557.36	192.50
1534.25	195.40	1546.12	193.90	1558.17	192.40
1535.04	195.30	1546.92	193.80	1558.98	192.30
1535.82	195.20	1547.72	193.70	1559.79	192.20
1536.61	195.10	1548.52	193.60	1560.61	192.10
1537.40	195.00	1549.32	193.50	1561.42	192.00
1538.19	194.90	1550.12	193.40	1562.23	191.90
1538.98	194.80	1550.92	193.30	1563.05	191.80
1539.77	194.70	1551.72	193.20	1563.86	191.70

Chapter 53

Configuring 10-Gigabit Ethernet Framing

This section contains the following topics:

- 10-Gigabit Ethernet Framing Overview on page 781
- Configuring 10-Gigabit Ethernet Framing on page 781

10-Gigabit Ethernet Framing Overview

The 10-Gigabit Ethernet IQ2 and IQ2-E PIC for the M120, M320, and T Series routers operates with Type 3 FPCs. The 10-Gigabit Ethernet IQ2 and IQ2-E PIC supports all features of the IQ2 and IQ2-E family PICs. Additionally, it provides one external interface running at 10 Gbps that operates in two modes:

- 10GBASE-R, LAN Physical Layer Device (LAN PHY)
- 10GBASE-W, WAN Physical Layer Device (WAN PHY)

When the external interface is running in LAN PHY mode, it bypasses the WIS sublayer to directly stream block-encoded Ethernet frames on a 10-Gigabit Ethernet serial interface. When the external interface is running in WAN PHY mode, it uses the WIS sublayer to transport 10-Gigabit Ethernet frames in an OC192c SONET payload.

Although the external interface provides a lower throughput when running in WAN PHY mode because of the extra SONET overhead, it can interoperate with SONET section or line level repeaters. This creates an advantage when the interface is used for long-distance, point-to-point 10-Gigabit Ethernet links. When the external interface is running in WAN PHY mode, some SONET options are supported. For information about SONET options supported on this interface, see “Configuring SONET Options for 10-Gigabit Ethernet Interfaces” on page 872.

Configuring 10-Gigabit Ethernet Framing

The 10-Gigabit Ethernet IQ2 and IQ2-E PIC uses the interface type *xe-fpc/pic/port*. On this single-port PIC, the port number is always zero.

The *xe-fpc/pic/port* interface inherits all the configuration commands that are used for gigabit Ethernet (*ge-fpc/pic/port*) interfaces.

To configure LAN PHY or WAN PHY operating mode, include the *framing* statement with the *lan-phy* or *wan-phy* option at the [edit interfaces *xe-fpc /pic/0*] hierarchy level.

```
[edit interfaces xe-fpc/pic/0 framing]  
framing (lan-phy | wan-phy);
```

To display interface information, use the operational mode command **show interfaces xe-fpc/pic/port extensive**.



NOTE: If you configure the WAN PHY mode on an aggregated Ethernet interface, you must set the aggregated Ethernet link speed to OC192.

Chapter 54

Configuring 10-Gigabit Ethernet Notification of Link Down Alarm

This section contains the following topics:

- 10-Gigabit Ethernet Notification of Link Down Alarm Overview on page 783
- Configuring 10-Gigabit Ethernet Notification of Link Down Alarm on page 783

10-Gigabit Ethernet Notification of Link Down Alarm Overview

Notification of link down alarm generation and transfer is supported for all 10-Gigabit Ethernet PIC interfaces, M120, M320, and T Series routers.

Configuring 10-Gigabit Ethernet Notification of Link Down Alarm

To configure this option, include the `asynchronous-notification` statement at the [edit interfaces *ge-fpc/pic/port* *gigether-options*] hierarchy level:

```
[edit interfaces]
ge-fpc/pic/port {
  gigether-options {
    asynchronous-notification;
  }
}
```


Chapter 55

Configuring Point-to-Point Protocol over Ethernet

This chapter includes the following topics:

- PPPoE Overview on page 785
- Configuring PPPoE on page 788
- Disabling the Sending of PPPoE Keepalive Messages on page 796
- Verifying a PPPoE Configuration on page 796

PPPoE Overview

The Point-to-Point Protocol over Ethernet (PPPoE) connects multiple hosts on an Ethernet LAN to a remote site through a single customer premises equipment (CPE) device. Hosts share a common digital subscriber line (DSL), a cable modem, or a wireless connection to the Internet.

A J Series router can be configured as the CPE device for PPPoE connections. To use PPPoE, you must configure the router as a PPPoE client, encapsulate PPP packets over Ethernet, and initiate a PPPoE session.



NOTE: J4300 and J6300 routers with asymmetrical DSL (ADSL) Physical Interface Modules (PIMs) and symmetrical high-speed DSL (SHDSL) PIMs can use PPPoE over Asynchronous Transfer Mode (ATM) to connect through DSL lines only, not for direct ATM connections. For information about configuring ADSL and SHDSL interfaces, see “Configuring ATM-over-ADSL Interfaces” on page 355 and “Configuring ATM-over-SHDSL Interfaces” on page 361.

M120 and M320 Internet routers can be configured as a PPPoE access concentrator server. To configure a PPPoE server on an M120 or M320 Ethernet logical interface, specify PPPoE encapsulation, include the **pp0** statement for the pseudo PPPoE physical interface, and include the **server** statement in the PPPoE options under the logical interface.



NOTE: PPPoE encapsulation is not supported on M120 or M320 routers on an ATM2 IQ interface.

On the J Series router, PPPoE establishes a point-to-point connection between the client (the Services Router) and the server, also called an access concentrator. Multiple hosts can be connected to the Services Router, and their data can be authenticated, encrypted, and compressed before the traffic is sent to the PPPoE session on the Services Router's Fast Ethernet or ATM-over-ADSL interface. PPPoE is easy to configure and enables services to be managed on a per-user basis rather than on a per-site basis.

This overview contains the following topics:

- PPPoE Interfaces on page 786
- PPPoE Stages on page 786
- Optional CHAP Authentication on page 788

PPPoE Interfaces

The PPPoE interface to the access concentrator can be a Fast Ethernet interface on any Services Router, a Gigabit Ethernet interface on J4350 and J6350 Services Routers, an ATM-over-ADSL or ATM-over-SHDSL interface on all J Series Services Routers except the J2300, or an ATM-over-SHDSL interface on a J2300 Services Router. The PPPoE configuration is the same for both interfaces. The only difference is the encapsulation for the underlying interface to the access concentrator:


- If the interface is Fast Ethernet, use a PPPoE encapsulation.
- If the interface is ATM over ADSL, use a PPPoE over ATM encapsulation.

The PPPoE interface on M120 or M320 routers acting as a access concentrator can be a Gigabit Ethernet or 10-Gigabit Ethernet interface.

Ethernet Interface

The Services Router encapsulates each PPP frame in an Ethernet frame and transports the frames over an Ethernet loop. Figure 69 on page 786 shows a typical PPPoE session between a Services Router and an access concentrator on the Ethernet loop.

Figure 69: PPPoE Session on an Ethernet Loop



PPPoE Stages

PPPoE has two stages, the discovery stage and the PPPoE session stage. In the discovery stage, the client discovers the access concentrator by identifying the

Ethernet media access control (MAC) address of the access concentrator and establishing a PPPoE session ID. In the PPPoE session stage, the client and the access concentrator build a point-to-point connection over Ethernet, based on the information collected in the discovery stage.

PPPoE Discovery Stage

A Services Router initiates the PPPoE discovery stage by broadcasting a PPPoE active discovery initiation (PADI) packet. To provide a point-to-point connection over Ethernet, each PPPoE session must learn the Ethernet MAC address of the access concentrator and establish a session with a unique session ID. Because the network might have more than one access concentrator, the discovery stage allows the client to communicate with all of them and select one.



NOTE: A Services Router cannot receive PPPoE packets from two different access concentrators on the same physical interface.

The PPPoE discovery stage consists of the following steps:

1. PPPoE active discovery initiation (PADI)—The client initiates a session by broadcasting a PADI packet on the LAN to request a service.
2. PPPoE active discovery offer (PADO)—Any access concentrator that can provide the service requested by the client in the PADI packet replies with a PADO packet that contains its own name, the unicast address of the client, and the service requested. An access concentrator can also use the PADO packet to offer other services to the client.
3. PPPoE active discovery request (PADR)—From the PADOs it receives, the client selects one access concentrator based on its name or the services offered and sends it a PADR packet to indicate the service or services needed.
4. PPPoE active discovery session-Confirmation (PADS)—When the selected access concentrator receives the PADR packet, it accepts or rejects the PPPoE session.
 - To accept the session, the access concentrator sends the client a PADS packet with a unique session ID for a PPPoE session and a service name that identifies the service under which it accepts the session.
 - To reject the session, the access concentrator sends the client a PADS packet with a service name error and resets the session ID to zero.

PPPoE Session Stage

The PPPoE session stage starts after the PPPoE discovery stage is over. The access concentrator can start the PPPoE session after it sends the PADS packet to the client, or the client can start the PPPoE session after it receives a PADS packet from the access concentrator. A Services Router supports multiple PPPoE sessions on each interface, but no more than 256 PPPoE sessions on all interfaces on the Services Router.

Each PPPoE session is uniquely identified by the Ethernet address of the peer and the session ID. After the PPPoE session is established, data is sent as in any other PPP encapsulation. The PPPoE information is encapsulated within an Ethernet frame and is sent to a unicast address. In this stage, both the client and the server must allocate resources for the PPPoE logical interface.

After a session is established, the client or the access concentrator can send a PPPoE active discovery termination (PADT) packet anytime to terminate the session. The PADT packet contains the destination address of the peer and the session ID of the session to be terminated. After this packet is sent, the session is closed to PPPoE traffic.

Optional CHAP Authentication

For interfaces with PPPoE encapsulation, you can configure interfaces to support the PPP Challenge Handshake Authentication Protocol (CHAP). When you enable CHAP on an interface, the interface can authenticate its peer and be authenticated by its peer.

If you configure an interface to handle incoming CHAP packets only (by including the **passive** statement at the [edit interfaces *interface-name* ppp-options chap] hierarchy level), the interface does not challenge its peer. However, if the interface is challenged, it responds to the challenge. If you do not include the **passive** statement, the interface always challenges its peer.

For more information about CHAP, see “Configuring the PPP Challenge Handshake Authentication Protocol” on page 112.

Configuring PPPoE

To configure PPPoE on a J Series Services Router, perform the following tasks:

1. Configure PPPoE encapsulation for an Ethernet interface or Ethernet over ATM encapsulation for an ATM-over-ADSL interface.
2. If you are configuring ATM over ADSL, configure LLC encapsulation on the logical interface.
3. Specify the logical Ethernet interface or the logical ATM interface as the underlying interface for the PPPoE session.
4. Configure the operational mode as client.
5. Identify the access concentrator by a unique name.
6. Optionally, specify how many seconds to wait before attempting to reconnect.
7. Provide a name for the type of service provided by the access concentrator.
8. Optionally, configure the maximum transmission unit (MTU) of the interface.
9. Configure the PPPoE interface address.
10. Configure the destination PPPoE interface address.
11. Optionally, configure the MTU size for the protocol family.
12. Optionally, disable the sending of keepalive messages on the logical interface.

To configure PPPoE on an M120 or M320 Internet Router operating as an access concentrator, perform the following tasks:

1. Configure PPPoE encapsulation for an Ethernet interface.
2. Specify the logical Ethernet interface as the underlying interface for the PPPoE session.
3. Optionally, configure the maximum transmission unit (MTU) of the interface.
4. Configure the operational mode as server.
5. Configure the PPPoE interface address.
6. Configure the destination PPPoE interface address.
7. Optionally, configure the MTU size for the protocol family.

Setting the Appropriate Encapsulation on the PPPoE Interface

For PPPoE on an Ethernet interface, you must configure encapsulation on the logical interface and use PPP over Ethernet encapsulation.

For PPPoE on an ATM-over-ADSL interface, you must configure encapsulation on both the physical and logical interfaces. To configure encapsulation on an ATM-over-ADSL physical interface, use Ethernet over ATM encapsulation. To configure encapsulation on an ATM-over-ADSL logical interface, use PPPoE over AAL5 LLC encapsulation. LLC encapsulation allows a single ATM virtual connection to transport multiple protocols.



NOTE: PPPoE encapsulation is not supported on an M120 or M320 router on an ATM2 IQ interface.

When you configure a point-to-point encapsulation such as PPP on a physical interface, the physical interface can have only one logical interface (only one unit statement) associated with it.

To configure physical interface properties, include the **encapsulation** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]  
encapsulation ethernet-over-atm;
```

To configure logical interface encapsulation properties, include the **encapsulation** statement:

```
encapsulation ppp-over-ether;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Perform the task appropriate for the interface on which you are using PPPoE:

- Configuring PPPoE Encapsulation on an Ethernet Interface on page 790
- Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface on page 790

Configuring PPPoE Encapsulation on an Ethernet Interface

Both the client and the server must be configured to support PPPoE. To configure PPPoE encapsulation on an Ethernet interface, include the `encapsulation` statement:

```
encapsulation ppp-over-ether;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces `pp0` unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces `pp0` unit *logical-unit-number*]

Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface

To configure the PPPoE encapsulation on a ATM-over-ADSL interface, perform the following steps:

1. Include the `encapsulation` statement at the [edit interfaces *interface-name*] hierarchy level, and specify `ethernet-over-atm`:

```
[edit interfaces pp0]
encapsulation ethernet-over-atm;
```

2. Configure LLC encapsulation on the logical interface by including the `encapsulation` statement and specifying `ppp-over-ether-over-atm-llc`:

```
encapsulation ppp-over-ether-over-atm-llc;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces `pp0` unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces `pp0` unit *logical-unit-number*]

Configuring a PPPoE Interface

The following sections describe how to configure a PPPoE interface:

- Configuring the PPPoE Underlying Interface on page 791
- Identifying the Access Concentrator on page 791
- Configuring the PPPoE Automatic Reconnect Wait Timer on page 792
- Configuring the PPPoE Service Name on page 792
- Configuring the PPPoE Server Mode on page 793
- Configuring the PPPoE Client Mode on page 793

- Configuring the PPPoE Source and Destination Addresses on page 793
- Deriving the PPPoE Source Address From a Specified Interface on page 794
- Configuring the PPPoE IP Address by Negotiation on page 794
- Configuring the Protocol MTU PPPoE on page 794
- Example: Configuring a PPPoE Client Interface on a J Series Services Router on page 795
- Example: Configuring a PPPoE Server Interface on an M120 or M320 Router on page 796

Configuring the PPPoE Underlying Interface

To configure the underlying Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or ATM interface, include the `underlying-interface` statement at the `[edit interfaces pp0 unit logical-unit-number pppoe-options]` hierarchy level:

```
[edit interfaces pp0]
unit logical-unit-number {
  pppoe-options {
    underlying-interface interface-name;
  }
}
```

You can include this statement at the following hierarchy levels:

- `[edit interfaces pp0 unit logical-unit-number pppoe-options]`
- `[edit logical-systems logical-system-name interfaces pp0 unit logical-unit-number pppoe-options]`

Specify the logical Ethernet, Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or ATM interface as the underlying interface—for example, `at-0/0/1.0` (ATM VC), `fe-1/0/1.0` (Fast Ethernet interface), or `ge-2/0/0` (Gigabit Ethernet interface).

Identifying the Access Concentrator

When configuring a PPPoE client, identify the access concentrator by a unique name by including the `access-concentrator` statement at the `[edit interfaces interface-name unit logical-unit-number pppoe-options]` hierarchy level:

```
[edit interfaces pp0]
unit logical-unit-number{
  pppoe-options {
    access-concentrator name;
  }
}
```

Specify the access concentrator name.

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* pppoe-options]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-options]

Configuring the PPPoE Automatic Reconnect Wait Timer

By default, after a PPPoE session is terminated, the session attempts to reconnect immediately. When configuring a PPPoE client, you can specify how many seconds to wait before attempting to reconnect, by including the **auto-reconnect** statement at the [edit interfaces *interface-name* unit *logical-unit-number* pppoe-options] hierarchy level:

```
[edit interfaces pp0]
unit logical-unit-number {
  pppoe-options {
    auto-reconnect seconds;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* pppoe-options]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-options]

You can configure the reconnection attempt to occur in 0 through 4,294,967,295 seconds after the session terminates.

Configuring the PPPoE Service Name

When configuring a PPPoE client, identify the type of service provided by the access concentrator—such as the name of the Internet service provider (ISP), class, or quality of service—by including the **service-name** statement at the [edit interfaces *interface-name* unit *logical-unit-number* pppoe-options] hierarchy level:

```
[edit interfaces pp0]
unit logical-unit-number {
  pppoe-options {
    service-name name;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* pppoe-options]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-options]

Configuring the PPPoE Server Mode

When configuring a PPPoE server, identify the mode by including the **server** statement at the [edit interfaces *interface-name* unit *logical-unit-number* pppoe-options] hierarchy level:

```
[edit interfaces pp0]
unit logical-unit-number {
  pppoe-options {
    server;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* pppoe-options]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-options]

Configuring the PPPoE Client Mode

When configuring a PPPoE client, identify the mode by including the **client** statement at the [edit interfaces *interface-name* unit *logical-unit-number* pppoe-options] hierarchy level:

```
[edit interfaces pp0]
unit logical-unit-number {
  pppoe-options {
    client;
  }
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number* pppoe-options]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* pppoe-options]

Configuring the PPPoE Source and Destination Addresses

When configuring a PPPoE client or server, assign source and destination addresses—for example, 192.168.1.1/32 and 192.168.1.2. To assign the source and destination address, include the **address** and **destination** statements:

```
address address {
  destination address;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces pp0.0 family inet]

- [edit logical-systems *logical-system-name* interfaces pp0.0 family inet]

Deriving the PPPoE Source Address From a Specified Interface

For a router supporting PPPoE, you can derive the source address from a specified interface—for example, the loopback interface, `lo0.0`—and assign a destination address—for example, `192.168.1.2`. The specified interface must include a logical unit number and have a configured IP address. To derive the source address and assign the destination address, include the `unnumbered-address` and `destination` statements:

```
unnumbered-address interface-name destination address;  
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces pp0.0 family inet]
- [edit logical-systems *logical-system-name* interfaces pp0.0 family inet]

Configuring the PPPoE IP Address by Negotiation

You can have the PPPoE client router obtain an IP address by negotiation with the remote end. This method might require the access concentrator to use a RADIUS authentication server. To obtain an IP address from the remote end by negotiation, include the `negotiate-address` statement:

```
negotiate-address;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces pp0.0 family (inet | inet6 | mpls)]
- [edit logical-systems *logical-system-name* interfaces pp0.0 family (inet | inet6 | mpls)]

Configuring the Protocol MTU PPPoE

You can configure the maximum transmission unit (MTU) size for the protocol family. Specify a range from 0 through 5012 bytes. Ensure that the size of the media MTU is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. To set the MTU, include the `mtu` statement:

```
mtu bytes;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces pp0.0 family (inet | inet6 | mpls)]
- [edit logical-systems *logical-system-name* interfaces pp0.0 family (inet | inet6 | mpls)]

You can modify the MTU size of the interface by including the `mtu bytes` statement at the [edit interfaces pp0] hierarchy level:

```
[edit interfaces pp0]
mtu bytes;
```

The default media MTU size used and the range of available sizes on a physical interface depends on the encapsulation used on that interface.

Example: Configuring a PPPoE Client Interface on a J Series Services Router

Configure a PPPoE over ATM-over-ADSL interface:

```
[edit interfaces]
at-2/0/0 {
  encapsulation ethernet-over-atm;
  atm-options {
    vpi 0;
  }
  dsl-options {
    operating-mode auto;
  }
  unit 0 {
    encapsulation ppp-over-ether-over-atm-llc;
    vci 0.120;
  }
}
pp0 {
  mtu 1492;
  unit 0 {
    ppp-options {
      chap {
        access-profile A-ppp-client;
        local-name A-at-2/0/0.0;
      }
    }
    pppoe-options {
      underlying-interface at-2/0/0;
      client;
      access-concentrator ispl.com;
      service-name "video@ispl.com";
      auto-reconnect 100;
    }
    no-keepalives;
    family inet {
      negotiate-address;
      mtu 100;
    }
    family inet6 {
      negotiate-address;
      mtu 200;
    }
    family mpls {
      negotiate-address;
      mtu 300;
    }
  }
}
```

```
}
```

Example: Configuring a PPPoE Server Interface on an M120 or M320 Router

Configure a PPPoE server over a Gigabit Ethernet interface:

```
[edit interfaces]
ge-1/0/0 {
  vlan-tagging;
  unit 1 {
    encapsulation ppp-over-ether;
    vlan-id 10;
  }
}
pp0 {
  unit 0 {
    pppoe-options {
      underlying-interface ge-1/0/0.0;
      server;
    }
    ppp-options {
    }
    family inet {
      address 22.2.2.1/32 {
        destination 22.2.2.2;
      }
    }
  }
}
```

Disabling the Sending of PPPoE Keepalive Messages

When configuring the client, you can disable the sending of keepalive messages on a logical interface by including the `no-keepalives` statement:

```
no-keepalives;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces pp0.0]
- [edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*]

Verifying a PPPoE Configuration

To verify a PPPoE configuration, you can issue the following operational mode commands:

- show interfaces at-*fpc/pic/port* extensive
- show interfaces pp0
- show pppoe interfaces

- show pppoe version
- show pppoe statistics

For more information about these operational mode commands, see the *J-series Services Router Basic LAN and WAN Access Configuration Guide* and the *JUNOS Interfaces Command Reference*.

Chapter 56

Configuring Ethernet Ring Protection Switching

- Ethernet Ring Protection Switching Overview on page 799
- Ethernet Ring Protection Switching Functionality on page 800
- Configuring Ethernet Ring Protection Switching on page 804
- Ethernet Ring Protection Switching Configuration Example on page 805

Ethernet Ring Protection Switching Overview

MX Series routers support *Ethernet ring protection switching*, which helps achieve high reliability and network stability. Links in the ring will never form loops that fatally affect the network operation and services availability. The basic idea of an Ethernet ring is to use one specific link to protect the whole ring. This special link is called a *ring protection link (RPL)*. If no failure happens in other links of the ring, the RPL blocks the traffic and is not used. RPL is controlled by a special node called an *RPL owner*. There is only one RPL owner in a ring. The RPL owner is responsible for blocking traffic over the RPL. Under ring failure conditions, the RPL owner is responsible for unblocking traffic over the RPL. A ring failure results in protection switching of the RPL traffic. An APS protocol is used to coordinate the protection actions over the ring. Protection switching blocks traffic on the failed link and unblocks the traffic on the RPL. When the failure clears, revertive protection switching blocks traffic over the RPL and unblocks traffic on the link on which the failure is cleared.

The following standards provide detailed information on Ethernet ring protection switching:

- IEEE 802.1Q - 1998
- IEEE 802.1D - 2004
- IEEE 802.1Q - 2003
- Draft ITU-T Recommendation G.8032/Y.1344, *Ethernet Ring protection switching*
- ITU-T Y.1731, *OAM functions and mechanisms for Ethernet-based networks*

For additional information on configuring Ethernet ring protection switching on MX Series routers, see the *MX Solutions Guide* for a complete example of Ethernet rings, and the *Layer 2 Configuration Guide* for information about STP loop avoidance and prevention.

Ethernet Ring Protection Switching Functionality

This section includes the following topics:

- Acronyms on page 800
- Ring Nodes on page 800
- Ring Node States on page 801
- Failure Detection on page 801
- Logical Ring on page 801
- FDB Flush on page 801
- Traffic Blocking and Forwarding on page 801
- RAPS Message Blocking and Forwarding on page 802
- Dedicated Signaling Control Channel on page 803
- RAPS Message Termination on page 803
- Manual Switch on page 803
- Non-Revertive Switch on page 803
- Multiple Rings on page 803
- Node ID on page 804
- Bridge Domains with the Ring Port on page 804

Acronyms

The following acronyms are used in this section:

- MA—maintenance association
- MEP—maintenance association end point
- OAM—connectivity fault management daemon
- FDB—MAC forwarding database
- STP—Spanning Tree Protocol
- RAPS—ring automatic protection switching
- WTR—wait to restore
- RPL—ring protection link

Ring Nodes

Multiple nodes are used to form a ring. For each ring node. There are two different node types:

- Normal node—The node has no special role on the ring.
- RPL owner node—The node owns the RPL and blocks or unblocks traffic over the RPL. This node also initiates the RAPS message.

Ring Node States

There are three different states for each node of a specific ring:

- **init**—Not a participant of a specific ring.
- **idle**—No failure on the ring, the node is performing normally. For normal node, traffic is unblocked on both ring ports. For the RPL owner, traffic is blocked on the ring port that connects to the RPL and unblocked on the other ring port.
- **protection**—A failure occurred on the ring. For normal node, traffic is blocked on the ring port that connects to the failing link and unblocked on working ring ports. For the RPL owner, traffic is unblocked on both ring ports if they connect to non-failure links.

There can only one RPL owner for each ring. The user configuration must guarantee this, because the APS protocol cannot check this.

Failure Detection

Ethernet ring operation depends on quick and accurate failure detection. The failure condition *signal failure (SF)* is supported. For SF detection, an Ethernet continuity check MEP must be configured for each ring link. For fast protection switching, a 10 ms transmission period for this MEP group is supported. OAM monitors the MEP group's MA and reports SF or SF clear events to the Ethernet ring control module. For this MEP group, the action profile must be configured to update the interface device IFF_LINKDOWN flag. OAM updates the IFF_LINKDOWN flag to notify the Ethernet ring control module.

Logical Ring

This feature currently supports only the physical ring, which means that two adjacent nodes of a ring must be physically connected and the ring operates on the physical interface, not the VLAN.

FDB Flush

When ring protection switching occurs, normally an *FDB flush* should be executed. The Ethernet ring control module should use the same mechanism as the STP to trigger the FDB flush. The Ethernet ring control module controls the ring port physical interface's default STP index to execute the FDB flush.

Traffic Blocking and Forwarding

The Ethernet ring control module uses the same mechanism as the STP to control forwarding or discarding of user traffic. The Ethernet ring control module sets the ring port physical interface default STP index state to forwarding or discarding in order to control user traffic.

RAPS Message Blocking and Forwarding

The router treats the RAPS message the same as user traffic for forwarding RAPS messages between two ring ports. The ring port physical interface default STP index state also controls forwarding RAPS messages between the two ring ports. Other than forwarding RAPS between the two ring ports, as shown in Figure 70 on page 802, the system also needs to forward the RAPS message between the CPU (Ethernet ring control module) and the ring port. This type of forwarding does not depend on the ring port physical interfaces STP index state. The RAPS message is always sent by the router through the ring ports, as shown in Figure 71 on page 802. An RAPS message received from a discarding ring port is sent to the Ethernet ring control module but is not sent to the other ring port.

Figure 70: Protocol Packets from the Network to the Router





Figure 71: Protocol Packets from the Router to the Network



Juniper Networks routers use an implicit filter to achieve these routes. Each implicit filter binds to a bridge domain. Therefore, the east ring port control channel and the west ring port control channel of a particular ring instance must be configured to the same bridge domain. For each ring port control channel, a filter term is generated to control RAPS message forwarding. The filter number is the same as the number of bridge domains that contain the ring control channels. If a bridge domain contains control channels from multiple rings, the filter related to this bridge domain will have multiple terms and each term will relate to a control channel. The filter has command parts and control-channel related parts, as follows:

- Common terms:
 - term 1: if [Ethernet type is not OAM Ethernet type (0x8902)]
 { accept packet }
 - term 2: if [source MAC address belongs to this bridge]
 { drop packet, our packet loop through the ring and come back to home }
 - term 3: if [destination is the RAPS PDU multicast address(0x01,0x19,0xa7, 0x00,0x00,0x01) AND[ring port STP status is DISCARDING]
 { send to CPU }

- Control channel related terms:
 - if [destination is the RAPS PDU multicast address(0x01,0x19,0xa7,0x00,0x00,0x01] AND[ring port STP status is FORWARDING] AND [Incoming interface IFL equal to control channel IFL] { send packet to CPU and send to the other ring port } default term: accept packet.

Dedicated Signaling Control Channel

For each ring port, a dedicated signaling control channel with a dedicated VLAN ID must be configured. In Ethernet ring configuration, only this control logical interface is configured and the underlying physical interface is the physical ring port. Each ring requires that two control physical interfaces be configured. These two logical interfaces must be configured in a bridge domain in order to forward RAPS PDUs between the two ring control physical interfaces. If the control channel logical interface is not a trunk port, only control logical interfaces will be configured in ring port configuration. If this control channel logical interface is a trunk port, in addition to the control channel logical interfaces, a dedicated VLAN ID must be configured.

RAPS Message Termination

The RAPS message starts from the originating node, travels through the entire ring, and terminates in the originating node unless a failure is present in the ring. The originating node must drop the RAPS message if the source MAC address in the RAPS message belongs to itself. The source MAC address is the node's node ID.

Manual Switch

Manual switch is not supported in this release.

Non-Revertive Switch

In revertive operation, once the condition causing a switch has cleared, traffic is blocked on the RPL and restored to the working transport entity. In non-revertive operation, traffic is allowed to use the RPL if it has not failed, even after a switch condition has cleared. Non-revertive switching is not supported in this release.

Multiple Rings

The Ethernet ring control module supports multiple rings in each node, (two logical interfaces are part of each ring). However, interconnection of multiple rings is not supported in this release. The interconnection of two rings means that two rings may share the same link or share the same node.

Node ID

For each node in the ring, a unique *node ID* identifies each node. The node ID is the node's MAC address. You can configure this node ID when configuring the ring on the node or automatically select an ID such as STP. In most cases, you will not configure this and the router will select a node ID, like STP does. It should be the manufacturing MAC address. The ring node ID should not be changed, even if you change the manufacturing MAC address. Any MAC address can be used if you make sure each node in the ring has a different node ID.

Bridge Domains with the Ring Port

From the router point of view, the protection group is seen as an abstract logical port that can be configured to any bridge domain. Therefore, if you configure one ring port or its logical interface in a bridge domain; you must configure the other related ring port or its logical interface to the same bridge domain. The bridge domain that includes the ring port acts as any other bridge domain and supports the IRB layer 3 interface.

Configuring Ethernet Ring Protection Switching

The inheritance model follows:

```

protection-group {
  ethernet-ring ring-name {
    east-interface {
      control-channel channel-name {
        vlan number;
      }
    }
    guard-interval number;
    node-id mac-address;
    restore-interval number;
    ring-protection-link-owner;
    west-interface {
      control-channel channel-name {
        vlan number;
      }
    }
  }
}

```

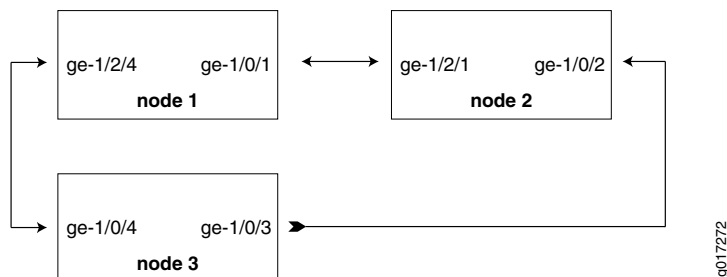
For each ring, a protection group must be configured. There may be several rings in each node, so there should be multiple protection groups corresponding to the related Ethernet rings.

Three interval parameters (*restore-interval*, *guard-interval* and *hold-interval*) can be configured at the protection group level. These configurations are global configurations and apply to all Ethernet rings if the Ethernet ring doesn't have a more specific configuration for these values. If no parameter is configured at the protection group level, the global configuration of this parameter uses the default value.

Ethernet Ring Protection Switching Configuration Example

In this section, a configuration example of a three node ring is given. The ring topology is shown in Figure 72 on page 805.

Figure 72: Example of a Three Node Ring



The configuration in this section is only for the RAPS channel. The bridge domain for user traffic is the same as the normal bridge domain. The only exception is if a bridge domain includes a ring port; then it must also include the other ring port of the same ring.

Configuration for Node 1

```

interfaces {
  ge-1/0/1 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 100;
    }
  }
  ge-1/2/4 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 100;
    }
  }
}
bridge-domains {
  bd1 {
    domain-type bridge;
    interface ge-1/2/4.1;
    interface ge-1/0/1.1;
  }
}
protocols {
  protection-group {
    ethernet-ring pg101 {
      node-id 00:01:01:00:00:01;
      ring-protection-link-owner;
      east-interface {

```

```

        control-channel ge-1/0/1.1;
        ring-protection-link-end;
    }
    west-interface {
        control-channel ge-1/2/4.1;
    }
}
}
}
protocols {
    oam {
        ethernet {
            connectivity-fault-management {
                action-profile rmep-defaults {
                    default-action {
                        interface-down;
                    }
                }
            }
            maintenance-domain d1 {
                level 0;
                maintenance-association 100 {
                    mep 1 {
                        interface ge-1/0/1;
                        remote-mep 2 {
                            action-profile rmep-defaults;
                        }
                    }
                }
            }
            maintenance-domain d2 {
                level 0;
                maintenance-association 100 {
                    mep 1 {
                        interface ge-1/2/4;
                        remote-mep 2 {
                            action-profile rmep-defaults;
                        }
                    }
                }
            }
        }
    }
}
}
}
}
}
```

Configuration for Node 2

```

interfaces {
  ge-1/0/2 {
    vlan-tagging;
    encapsulation flexible-ethernet-services;
    unit 1 {
      encapsulation vlan-bridge;
      vlan-id 100;
    }
  }
}

```

```

ge-1/2/1 {
  vlan-tagging;
  encapsulation flexible-ethernet-services;
  unit 1 {
    encapsulation vlan-bridge;
    vlan-id 100;
  }
}

bridge-domains {
  bd1 {
    domain-type bridge;
    interface ge-1/2/1.1;
    interface ge-1/0/2.1;
  }
}

protocols {
  protection-group {
    ethernet-ring pg102 {
      east-interface {
        control-channel ge-1/0/2.1;
      }
      west-interface {
        control-channel ge-1/2/1.1;
      }
    }
  }
}

protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        action-profile rmep-defaults {
          default-action {
            interface-down;
          }
        }
        maintenance-domain d1 {
          level 0;
          maintenance-association 100 {
            mep 2 {
              interface ge-1/2/1;
              remote-mep 1 {
                action-profile rmep-defaults;
              }
            }
          }
        }
      }
    }
    maintenance-domain d3 {
      level 0;
      maintenance-association 100 {

```

```

        mep 1 {
            interface ge-1/0/2;
            remote-mep 2 {
                action-profile rmep-defaults;
            }
        }
    }
}

```

Configuration for Node 3

```

interfaces {
    ge-1/0/4 {
        vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 1 {
            encapsulation vlan-bridge;
            vlan-id 100;
        }
    }

    ge-1/0/3 {
        vlan-tagging;
        encapsulation flexible-ethernet-services;
        unit 1 {
            encapsulation vlan-bridge;
            vlan-id 100;
        }
    }
}

bridge-domains {
    bd1 {
        domain-type bridge;
        interface ge-1/0/4.1;
        interface ge-1/0/3.1;
    }
}

protocols {
    protection-group {
        ethernet-ring pg103 {
            east-interface {
                control-channel ge-1/0/3.1;
            }
            west-interface {
                control-channel ge-1/0/4.1;
            }
        }
    }
}

```



```

protocols {
  oam {
    ethernet {
      connectivity-fault-management {
        action-profile rmep-defaults {
          default-action {
            interface-down;
          }
        }
      }
      maintenance-domain d2 {
        level 0;
        maintenance-association 100 {
          mep 2 {
            interface ge-1/0/4;
            remote-mep 1 {
              action-profile rmep-defaults;
            }
          }
        }
      }
      maintenance-domain d3 {
        level 0;
        maintenance-association 100 {
          mep 2 {
            interface ge-1/0/3;
            remote-mep 1 {
              action-profile rmep-defaults;
            }
          }
        }
      }
    }
  }
}

```

Examples: Ethernet RPS Output

This section provides output examples based on the configuration shown in “Ethernet Ring Protection Switching Configuration Example” on page 805. The show commands used in these examples can help verify configuration and correct operation. The following situations are shown:

- Normal Situation on page 809
- Failure Situation on page 811

Normal Situation

RPL Owner Node If ring has no failure, the show command will have the following output for node 1:

```
user@node1> show protection-group ethernet-ring aps
```

```
Ethernet Ring Name Request/state No Flush Ring Protection Link Blocked
```

```

pg101                NR                No                Yes

Originator Remote Node ID
Yes
user@node1> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101

Interface    Control Channel Forward State Ring Protection Link End
ge-1/0/1     ge-1/0/1.1      discarding   Yes
ge-1/2/4     ge-1/2/4.1      forwarding   No

Signal Failure Admin State
Clear         IFF ready
Clear         IFF ready
user@node1> show protection-group ethernet-ring node-state
Ethernet ring APS State Event Ring Protection Link Owner
pg101         idle NR-RB Yes

Restore Timer Quard Timer Operation state
disabled      disabled operational
user@node1> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent : 1
RAPS received : 0
Local SF happened: : 0
Remote SF happened: : 0
NR event happened: : 0
NR-RB event happened: : 1

```

Other Nodes For Node 2 and Node 3, the outputs should be same:

```

user@node2> show protection-group ethernet-ring aps
Ethernet Ring Name Request/state No Flush Ring Protection Link Blocked
pg102             NR             No      Yes

Originator Remote Node ID
No          00:01:01:00:00:01
user@node2> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102

Interface    Control Channel Forward State Ring Protection Link End
ge-1/2/1     ge-1/2/1.1      forwarding   No
ge-1/0/2     ge-1/0/2.1      forwarding   No

Signal Failure Admin State
Clear         IFF ready
Clear         IFF ready
user@node2> show protection-group ethernet-ring node-state
Ethernet ring APS State Event Ring Protection Link Owner
pg102         idle NR-RB No

Restore Timer Quard Timer Operation state
disabled      disabled operational
user@node2> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg101
RAPS sent : 0
RAPS received : 1
Local SF happened: : 0
Remote SF happened: : 0

```

```
NR event happened:           : 0
NR-RB event happened:        : 1
```

Failure Situation

RPL Owner Node If ring has a link failure between Node2 and Node 3, the **show** command will have the following outputs for Node 1:

```
user@node1> show protection-group ethernet-ring aps
Ethernet Ring Name Request/state No Flush Ring Protection Link Blocked
pg101              SF           NO      No

Originator Remote Node ID
No          00:01:02:00:00:01
user@node1> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg101

Interface Control Channel Forward State Ring Protection Link End
ge-1/0/1   ge-1/0/1.1      forwarding Yes
ge-1/2/4   ge-1/2/4.1      forwarding No

Signal Failure Admin State
Clear          IFF ready
Clear          IFF ready
user@node1> show protection-group ethernet-ring node-state
Ethernet ring APS State Event Ring Protection Link Owner
pg101          protected SF Yes

Restore Timer Quard Timer Operation state
disabled      disabled operational
user@node1> show protection-group ethernet-ring statistics group-name pg101
Ethernet Ring statistics for PG pg101
RAPS sent           : 1
RAPS received       : 1
Local SF happened:   : 0
Remote SF happened:  : 1
NR event happened:   : 0
NR-RB event happened: : 1
```

Other Nodes For Node 2 and Node 3, the outputs should be same:

```
user@node2> show protection-group ethernet-ring aps
Ethernet Ring Name Request/state No Flush Ring Protection Link Blocked
pg102              SF           No      No

Originator Remote Node ID
Yes          00:00:00:00:00:00
user@node2> show protection-group ethernet-ring interface
Ethernet ring port parameters for protection group pg102

Interface Control Channel Forward State Ring Protection Link End
ge-1/2/1   ge-1/2/1.1      forwarding No
ge-1/0/2   ge-1/0/2.1      discarding No

Signal Failure Admin State
Clear          IFF ready
set            IFF ready
user@node2> show protection-group ethernet-ring node-state
```

```
Ethernet ring    APS State    Event           Ring Protection Link Owner
pg102           idle         NR-RB          No

Restore Timer   Quard Timer   Operation state
disabled        disabled     operational
user@node2> show protection-group ethernet-ring statistics group-name pg102
Ethernet Ring statistics for PG pg101
RAPS sent                      : 1
RAPS received                  : 1
Local SF happened:              : 1
Remote SF happened:             : 0
NR event happened:              : 0
NR-RB event happened:           : 1
```

Chapter 57

Example Ethernet Configurations

This section contains the following example configurations for Ethernet interfaces:

- Example: Configuring Fast Ethernet Interfaces on page 813
- Example: Configuring Gigabit Ethernet Interfaces on page 813
- Example: Configuring Aggregated Ethernet Interfaces on page 814
- Example: Configuring Aggregated Ethernet Link Protection on page 815

Example: Configuring Fast Ethernet Interfaces

The following configuration is sufficient to get a Fast Ethernet interface up and running. By default, IPv4 Fast Ethernet interfaces use Ethernet Version 2 encapsulation.

```
[edit]
user@host# set interfaces fe-5/2/1 unit 0 family inet address local-address
user@host# show
interfaces {
  fe-5/2/1 {
    unit 0 {
      family inet {
        address local-address;
      }
    }
  }
}
```

Example: Configuring Gigabit Ethernet Interfaces

The following configuration is sufficient to get a Gigabit Ethernet, Tri-Rate Ethernet copper, or 10-Gigabit Ethernet interface up and running. By default, IPv4 Gigabit Ethernet interfaces on MX Series, M Series, and T Series routers use 802.3 encapsulation. J Series Gigabit Ethernet interfaces do not support 802.3 encapsulation.

```
[edit]
user@host# set interfaces ge-2/0/1 unit 0 family inet address local-address
user@host# show
interfaces {
  ge-2/0/1 {
    unit 0 {
      family inet {
```

```

        address local-address;
    }
}
}

```

The M160, M320, M120, T320, and T640 2-port Gigabit Ethernet PIC supports two independent Gigabit Ethernet links.

Each of the two interfaces on the PIC is named:

```
ge-fpc/pic/[0.1]
```

Each of these interfaces has functionality identical to the Gigabit Ethernet interface supported on the single-port PIC.

Example: Configuring Aggregated Ethernet Interfaces

Aggregated Ethernet interfaces can use interfaces from different FPCs, DPCs, or PICs. The following configuration is sufficient to get an aggregated Gigabit Ethernet interface up and running.

```

[edit chassis]
aggregated-devices {
  ethernet {
    device-count 15;
  }
}

[edit interfaces]
ge-1/3/0 {
  gigether-options {
    802.3ad ae0;
  }
}
ge-2/0/1 {
  gigether-options {
    802.3ad ae0;
  }
}
ae0 {
  aggregated-ether-options {
    link-speed 1g;
    minimum-links 1;
  }
}
vlan-tagging;
unit 0 {
  vlan-id 1;
  family inet {
    address 14.0.100.50/24;
  }
}
unit 1 {
  vlan-id 1024;
}

```

```

        family inet {
            address 14.0.101.50/24;
        }
    }
    unit 2 {
        vlan-id 1025;
        family inet {
            address 14.0.102.50/24;
        }
    }
    unit 3 {
        vlan-id 4094;
        family inet {
            address 14.0.103.50/24;
        }
    }
}

```

Example: Configuring Aggregated Ethernet Link Protection

The following configuration enables link protection on the `ae0` interface, and specifies the `ge-1/0/0` interface as the primary link and `ge-1/0/1` as the secondary link.

```

[edit interfaces]
ae0 {
    aggregated-ether-options {
        link protection;
    }
}
[edit interfaces]
ge-1/0/0 {
    gigether-options {
        802.3ad ae0 primary;
    }
}
[edit interfaces]
ge-1/0/1 {
    gigether-options {
        802.3ad ae0 backup;
    }
}

```


Part 11

Configuring ISDN Interfaces

- Configuring ISDN Interfaces on page 819

Chapter 58

Configuring ISDN Interfaces

This section contains the following topics:

- ISDN Interfaces Overview on page 819
- Configuring ISDN Services Physical and Logical Interface Properties on page 820
- Configuring ISDN Physical Interface Properties on page 821
- Configuring ISDN Logical Interface Properties on page 823
- Disabling ISDN Processes on page 840

ISDN Interfaces Overview

ISDN is a set of standards for digital transmission over different media created by the Consultative Committee for International Telegraphy and Telephony (CCITT) and the International Telecommunication Union (ITU). ISDN is a dial-on-demand service that provides fast call setup, low latency, and the ability to carry high-quality voice, data, and video transmissions. ISDN is also a circuit-switched service that can be used on both multipoint and point-to-point connections.

You configure two types of interfaces for ISDN service: a physical interface and a logical interface called the *dialer interface*.

Four types of Physical Interface Modules (PIMs) provide ISDN connectivity on J Series Services Routers:

- 1-port S/T interface supporting ITU-T I.430, ETSI TS 101080, and GR-1089-Core Type III
- 1-port U interface supporting ANSI T.601 and GR-1089-Core
- 4-port S/T interface supporting ITU-T I.430, ETSI TS 101080, and GR-1089-Core Type III as a field-replaceable unit (FRU)
- 4-port U interface supporting ANSI T.601 and GR-1089-Core

For information about installing hardware, see the *J-series Services Router Getting Started Guide*.

For information about installing ISDN service over an ISDN line, contact your telecommunications service provider.

Configuring ISDN Services Physical and Logical Interface Properties

To configure ISDN services physical interface properties, include the `isdn-options` statement at the `[edit interfaces br-pim/0/port]` hierarchy level:

```
[edit interfaces br-pim/0/port]
isdn-options {
  calling-number number;
  incoming-called-number number <reject>;
  spid1 spid-string;
  spid2 spid-string;
  static-tei-val value;
  switch-type (att5e | etsi | ni1 | ntdms100 | ntt);
  t310 seconds;
  tei-option (first-call | power-up);
}
dialer-options {
  pool pool-name <priority priority>;
}
```

To configure ISDN services logical interface properties, include the following statements:

```
[edit interfaces dln unit logical-unit-number]
dialer-options {
  activation-delay seconds;
  callback;
  callback-wait-period time;
  deactivation-delay seconds;
  dial-string dial-string-numbers;
  idle-timeout seconds;
  incoming-map {
    caller (caller-id | accept-all);
    initial-route-check seconds;
    load-interval seconds;
    load-threshold percent;
    pool pool-name;
    redial-delay time;
    watch-list {
      [ routes ];
    }
  }
  encapsulation [
    (cisco-hdlc | multilink-ppp | ppp);
  ]
}
```

To configure a primary interface to use an ISDN logical interface as a backup or “failover” interface when the primary connection experiences interruptions in Internet connectivity, include the `backup-options` statement to specify the ISDN interface at the `[edit interfaces interface-name unit logical-unit-number]` hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
```

```
backup-options {
  interface dln.0;
}
```

You can include these statements at the following hierarchy levels:

- [edit interfaces]
- [edit logical-systems *logical-system-name* interfaces]

To configure the Services Router to reject incoming ISDN calls (supported when dial-in is configured), include the **reject-incoming** statement at the [edit system processes isdn-signaling] hierarchy level:

```
[edit system processes isdn-signaling]
reject-incoming;
}
```

To disable ISDN, include the **disable** statement at the [edit system processes isdn-signaling] hierarchy level:

```
[edit system processes isdn-signaling]
disable;
}
```

To disable the dial-out on demand process, include the **disable** statement at the [edit system processes dialer-services] hierarchy level:

```
[edit system processes dialer-services]
disable;
}
```

Configuring ISDN Physical Interface Properties

You specify the physical ISDN interface in the form **br-*pim*/0/*port***. *pim* is the slot in which the PIM is installed. The second number is always 0. *port* is the configured port number.

You specify the B-channel in the form **bc-*pim*/0/*port*:*n***. *n* is the B-channel ID and can be 1 or 2. You specify the D-channel in the form **dc-*pim*/0/*port*:0**.



NOTE: The B- and D-channel interfaces do not have any configurable parameters. However, when interface statistics are displayed, B- and D-channel interfaces have statistical values.

To enable ISDN interfaces installed on your Services Router to work properly, you must configure the interface properties. To configure physical interface properties, include the **isdn-options** statement at the [edit interfaces **br-*pim*/0/*port***] hierarchy level:

```
[edit interfaces br-pim/0/port]
isdn-options {
```

```

calling-number number;
incoming-called-number number <reject>;
spid1 spid-string;
spid2 spid-string;
static-tei-val value;
switch-type (att5e | etsi | ni1 | ntdms100 | ntt);
t310 seconds;
tei-option (first-call | power-up);
}
dialer-options {
    pool pool-name <priority priority>;
}

```

You can configure the following ISDN options:

- **calling-number**—The calling number included in outgoing calls.
- **incoming-called-number**—Screening of incoming calls. If the incoming number of the incoming call is configured, the call is accepted. If the reject option is specified with the number, the call is rejected. If no numbers are configured, all calls are accepted. See “Configuring an ISDN Interface to Screen Incoming Calls” on page 823.
- **pool**—The dial pool for logical and physical ISDN interfaces. The dial pool allows logical (dialer) and physical (**br-pim/0/port**) interfaces to be bound together dynamically on a per-call basis. On a dialer interface, pool directs the dialer interface to a dial pool. On a **br-pim/0/port** interface, pool defines the pool to which the interface belongs. Specify a priority value from 0 (lowest) to 255 (highest) for the interface.
- **spid1**—The Service Profile Identifier (SPID). *spid-string* is a numeric value. If your service provider requires SPIDs, you cannot place calls until the interface sends a valid, assigned SPID to the service provider when accessing the ISDN connection. A single SPID must be configured as **spid1**.
- **spid2**—A second SPID, used for DMS-100 and NI1 switch types.
- **static-tei-val**—A static Terminal Endpoint Identifier (TEI) value. The TEI value represents any ISDN-capable device attached to an ISDN network that is the terminal endpoint. TEIs are used to distinguish between different devices using the same ISDN links. Specify a value from 0 through 63. You cannot configure a TEI value with multiple SPIDs—dynamic TEI assignment is required.



NOTE: TEI assignment is usually done dynamically instead of statically using the TEI management protocol. When the TEI management protocol is used, values 64-126 are assigned to terminal endpoints. TEI value 127 is used for group assignment.

- **switch-type**—The ISDN switch type. The following switches are compatible:
 - **att5e**—AT&T 5ESS
 - **etsi**—NET3 for United Kingdom and Europe
 - **ni1**—National ISDN-1

- **ntdms100**—Northern Telecom DMS-100
- **ntt**—NTT Group switch for Japan
- **tei-option**—When the Terminal Endpoint Identifier (TEI) negotiates with the ISDN provider. Specify first-call (activation does not occur until the call setup is sent) or power-up (activation occurs when the Services Router is powered on). The default value is power-up.
- **t310**—Q.931-specific timer for T310, in seconds. Specify the number of seconds from 1 through 65536. The default value is 10 seconds.

Configuring an ISDN Interface to Screen Incoming Calls

By default, an ISDN interface is configured to accept all incoming calls. If multiple devices are connected to the same ISDN line, you can configure an ISDN interface to screen incoming calls based on the incoming called number.

You can specify the incoming called numbers that an ISDN interface accepts. You can use the **reject** option to specify a number that the ISDN interface can ignore. The **reject** option is useful when an incoming called number is specified on one device connected to an ISDN line, and you want the incoming called number rejected on a second ISDN device connected to the same ISDN line. For example, if the first ISDN device has the called number 4085321901, you can configure the called number 4085321901 with the **reject** option on the second ISDN device.

When an incoming ISDN call is received, the Services Router matches the incoming called number against the called numbers configured on its ISDN interfaces. If an exact match is not found, or if the called number is configured with the **reject** option, the incoming call is ignored. Each ISDN interface accepts only the calls whose called number are configured on it.

To specify that an incoming called number be rejected by the interface, include the **incoming-called-number** statement with the **reject** option at the [edit interfaces *br-pim/0/port* isdn-options] hierarchy level:

```
[edit interfaces br-pim/0/port]
 isdn-options {
   incoming-called-number number reject;
 }
```

You can configure up to 30 incoming called numbers.

Configuring ISDN Logical Interface Properties

You configure ISDN services interface properties at the logical unit level.

The dialer interface, **dln**, is a logical interface for configuring dialing properties for a backup ISDN connection. The interface can be configured in two modes:

- Multilink mode using **multilink-ppp** encapsulation. This mode is used when the router supports B-channel bundling (two B-channels connected to provide a 128-Kbps connection) and runs Multilink Point-to-Point Protocol (MLPPP). When

the dialer interface (*dl*n**) is in multilink mode, the value of *n* is from 0 through 149. However, you can only configure one dialer interface with **multilink-ppp** encapsulation. For example, you cannot have both **dl1** and **dl2** as multilink dialers simultaneously. If you need to have multiple multilink dialers, then the values should be **dl*n*.1**, **dl*n*.2**, and so forth.

- Normal mode using **ppp** or **cisco-hdlc** encapsulation. This mode is used when the router is using one B-channel. When the dialer interface (*dl*n**) is in normal mode, the value of *n* is always from 0 through 149.



NOTE: Ensure that the same IP subnet address is not configured on different dialer interfaces. Configuring the same IP subnet address on different dialer interfaces can result in inconsistency in the route and packet loss. Packets may be routed through any of the dialer interfaces that have the same IP subnet address, instead of being routed through the dialer interface to which the ISDN call is connected.

You can configure the following ISDN services logical interface properties:

- Configuring an ISDN Dialer Interface as a Backup Interface on page 826
- Applying the Dial-on-Demand Dialer Filter to the Dialer Interfaces on page 828
- Configuring Bandwidth on Demand on page 829
- Configuring Dial-In and Callback on page 832
- Configuring Dialer Watch on page 835

The dialer interface cannot be configured:

- As a backup interface and as a dialer filter simultaneously.
- As a backup interface and as a dialer watch simultaneously.
- As a dialer watch interface and as a dialer filter simultaneously.
- As a backup interface for more than one primary interface.

For specific ISDN configuration information for dial-on-demand routing (DDR) and adding Open Shortest Path First (OSPF) demand circuits to a Services Router, see the *JUNOS Routing Protocols Configuration Guide*.

For general information about logical unit properties, see “Configuring Logical Interface Properties” on page 143. For general information about **family inet** properties, see “Configuring Protocol Family and Interface Address Properties” on page 169.

To configure logical interface properties, include the **encapsulation** statement at the **[edit interfaces *dl*n**]** hierarchy level and the **dialer-options** statement at the **[edit interfaces *dl*n** unit *logical-unit-number*]** hierarchy level:

```
[edit interfaces dln]
encapsulation (cisco-hdlc | multilink-ppp | ppp);
[edit interfaces dln unit logical-unit-number]
dialer-options {
    activation-delay seconds;
```



```

callback;
callback-wait-period time;
deactivation-delay seconds;
dial-string dial-string-numbers;
idle-timeout seconds;
incoming-map {
    caller (caller-id | accept-all);
    initial-route-check seconds;
    load-interval seconds;
    load-threshold percent;
    pool pool-name;
    redial-delay time;
    watch-list {
        [ routes ];
    }
}
}
}

```

You can configure the following options:

- **activation-delay**—ISDN activation delay, in seconds. Specify a number from 1 through 4294967295.
- **callback**—Configure the dialer to terminate the incoming call and call back the originator after the callback wait period.
- **callback-wait-period**—For interfaces configured for ISDN with callback, specify the amount of time the dialer waits before calling back the caller. The default is 5 seconds.
- **caller**—Specify the dialer to accept a specified caller number or accept all incoming calls.
- **deactivation-delay**—ISDN deactivation delay, in seconds. Specify from 1 through 4294967295.
- **encapsulation**—Logical link-layer encapsulation type. For normal mode, specify **cisco-hdlc** for Cisco-compatible High-Level Data Link control (HDLC) or **ppp** for Point-to-Point Protocol. For multilink mode, specify **multilink-ppp**.
- **dial-string**—Phone number to be dialed. Do not include hyphens in number.
- **idle-timeout**—Number of seconds the link is idle before losing connectivity. The default is 120 seconds.
- **incoming-map**—Specify the dialer to accept incoming calls. This statement is required at one end of the ISDN connection.



CAUTION: Changing the caller incoming map when a call is connected can create inconsistencies in the route and prevent traffic on a subnet from being transmitted. This is seen when two dialer interfaces are configured and the association of the caller incoming-map from one interface to the other is changed when a call is connected on one of the interfaces. The cause of the inconsistency is that dialer interfaces are pseudo interfaces that are always up, even if not actually connected.

- **initial-route-check**—Allows the router to check whether the primary route is up after the initial startup of the router is complete and the timer expires.
- **load-interval**—Interval used to calculate the average load on the network. By default, the average interface load is calculated every 60 seconds. You can specify an interval from 20 through 180 seconds, configurable in intervals of 10 seconds. For more information about the load interval, see “Configuring Bandwidth on Demand” on page 829.
- **load-threshold**—Bandwidth threshold percentage used for adding interfaces. Another link is added to the multilink bundle when the bandwidth reaches the threshold value you set. Specify a percentage between 0 and 100. When the value is set to 0, all available channels are dialed. The default value is 100.
- **pool**—For logical and physical ISDN interfaces, specify the dial pool. The dial pool allows logical (dialer) and physical (**br-pim/0/port**) interfaces to be bound together dynamically on a per-call basis. On a dialer interface, **pool** directs the dialer interface which dial pool to use. On a **br-pim/0/port** interface, **pool** defines the pool to which the interface belongs.
- **redial-delay**—Specify the delay (in seconds) between two successive calls made by the dialer (for dialout). The default is 3 seconds.
- **watch-list**—IP prefix of one or more routes. The primary route is considered up if there is at least one valid route for any of the addresses in the watch list to an interface other than the backup interface.

Changing the caller incoming map when a call is connected can create inconsistencies in the route and prevent traffic on a subnet from being transmitted. This is seen when two dialer interfaces are configured and the association of the caller incoming-map from one interface to the other is changed when a call is connected on one of the interfaces.

The cause of the inconsistency is that dialer interfaces are pseudo interfaces that are always up, even if not actually connected.

Configuring an ISDN Dialer Interface as a Backup Interface

Configuring the ISDN interface as a backup interface ensures continuous network connectivity. The Services Router can be configured to fail over to the ISDN interface if the primary connection experiences interruptions in Internet connectivity.

To configure an ISDN interface as the backup interface, include the **backup-options** statement at the [edit interfaces *interface-name* unit *logical-unit-number*] hierarchy level:

```
[edit interfaces interface-name unit logical-unit-number]
backup-options {
  interface dln.0;
}
```


interface-name is the primary interface. The backup interface is specified as *dl**n*.

Example: Configuring an ISDN Interface as the Backup Interface

The following example illustrates a backup configuration using a primary serial interface, two dialer interfaces, and a physical ISDN interface.

See Figure 73 on page 827 for the topology used for this example.

Figure 73: ISDN Backup Topology



Configure dialer interface d10 as the backup interface on the primary serial interface t1-4/0/1:

Configuration on the Primary Serial Interface

```
[edit interfaces]
t1-4/0/1 {
  encapsulation ppp;
  unit 0 {
    backup-options {
      interface d10.0;
    }
    family inet {
      address 192.168.10.2/30;
    }
  }
}
```

Configuration on the Dialer Interface

```
[edit interfaces]
d10 {
  encapsulation ppp;
  unit 0 {
    dialer-options {
      pool 10;
      dial-string 5552222;
      activation-delay 10;
      deactivation-delay 10;
      incoming-map
        caller 5552222 accept-all
    }
    family inet {
      address 172.16.10.2/32 {
        destination 172.16.10.1;
      }
    }
  }
}
```

```
}

```

**Configuration on the
Physical ISDN Interface**

```
[edit interfaces]
br-6/0/0 {
  isdn-options {
    calling-number 5558888;
    spid1 51255511110101 5551111;
    spid2 51255511120101 5551112;
    switch-type ni1;
    t310 70;
  }
  dialer-options {
    pool 10 priority 3;
    pool 2 priority 25;
  }
}
```

Applying the Dial-on-Demand Dialer Filter to the Dialer Interfaces

Dial-on-demand routing (DDR) links two sites over a public network and provides bandwidth. An ISDN connection allows an ISDN line to be activated only when there is network traffic configured as an “interesting” packet. An interesting packet is defined using the firewall filter feature of the Services Router.

To configure DDR, you configure the dialer interface as a passive static route with a lower priority than dynamic routes. If the dynamic route is lost, and a packet destined for that IP address is received, the dialer interface initiates an ISDN connection and sends the packet over it. When no new packets are sent to the destination, the dialer interface initiates an inactivity timer. The ISDN connection is terminated when the timer expires.

To configure dial-on-demand connectivity, perform the following steps:

- Define the dialer filter.
- Configure the firewall rule.
- Apply the dialer filter to the dialer interface.

To define the filter, include the **dialer-filter interesting-traffic** statement at the **[edit firewall family inet]** hierarchy level.

To configure the firewall rule, include the **term** and **from** statements at the **[edit firewall family inet dialer-filter *filter-name*]** hierarchy level.


To apply the filter to the dialer interface, include the **filter dialer** statement at the **[edit interfaces dln unit *logical-unit-number* family *family*]** hierarchy level.

Example: Applying the Dialer Filter

The following example illustrates a dialer filter configuration configured at the **[edit firewall family inet]** hierarchy level and applied to a physical interface and a dialer interface.

See Figure 74 on page 829 for the topology used for this example.

Figure 74: Dialer Filter Topology



**Configuration for the
Dialer Filter**

```
[edit firewall family inet]
dialer-filter interesting-traffic {
  term 1 {
    from {
      destination-address {
        10.2.1.1/30;
      }
    }
    then note;
  }
}
```

**Configuration on the
Dialer Interface**

```
[edit interfaces]
d10 {
  encapsulation ppp;
  unit 0 {
    dialer-options {
      pool 1;
      dial-string 350100;
    }
  }
  family inet {
    filter {
      dialer interesting-traffic;
    }
    address 50.2.0.1/24;
  }
}
```

Configuring Bandwidth on Demand

You can define a bandwidth threshold for network traffic on the Services Router using the dialer interface and ISDN interfaces. Initially, only one ISDN link is active and all packets are sent through this interface. When a predefined bandwidth threshold is reached on this interface, the dialer interface activates another ISDN link and initiates a data connection.

To configure bandwidth on demand, perform the steps in the following sections to configure the dialer interface and the physical ISDN interfaces:

- Configuring the Dialer Interface on page 830
- Configuring the ISDN Interface on page 831
- Example: Configuring Bandwidth on Demand on page 831

Configuring the Dialer Interface

To configure the dialer interface for bandwidth on demand, include the `encapsulation multilink-ppp` statement at the `[edit interfaces dln]` hierarchy level:

```
[edit interfaces]
dln {
    encapsulation multilink-ppp;
}
```

To configure dialer options, include the `dialer-options` statement at the `[edit interfaces dln unit logical-unit-number]` hierarchy level:

```
[edit interfaces dln unit logical-unit-number]
dialer-options {
    dial-string dial-string-numbers;
    load-interval seconds;
    load-threshold percent;
    pool pool-name;
}
```

To configure unit properties, include the `unit logical-unit-number` statement at the `[edit interfaces dln]` hierarchy level:

```
[edit interfaces dln unit logical-unit-number]
family family {
    mtu bytes;
    negotiate-address;
    filter {
        filter-name;
        fragment-threshold bytes;
        mrru bytes;
        ppp-options {
            chap {
                access-profile name;
            }
        }
    }
}
```

You can configure the following unit properties:

- **family**—Protocol family information for the logical interface. For *family*, specify *inet* (for Internet Protocol version 4 [IPv4]) suite.
- **filter**—Dialer filter name. The dialer filter applied here is configured at the [edit firewall family inet] hierarchy level and also applied to the physical ISDN interface.
- **fragment-threshold**—Maximum size, in bytes, for multilink packet fragments. Any nonzero value must be a multiple of 64 bytes. The value can be between 128 and 16320. The default is 0 bytes (no fragmentation).
- **mrru**—Maximum received reconstructed unit (MRRU), in bytes. The value can be between 1500 and 4500. The default is 1500 bytes.
- **negotiate-address**—For interfaces with Point-to-Point Protocol (PPP) encapsulation, enable the interface to be assigned an IP address by the remote end.

Configuring the ISDN Interface

To configure the ISDN interface for bandwidth on demand, include the **pool** statement at the [edit interfaces *br-pim/0/port dialer-options*] hierarchy level:


```
[edit interfaces br-pim/0/port]
dialer-options {
  pool pool-name;
}
```

Each ISDN interface must use the same dialer pool name to participate in the bandwidth-on-demand configuration.

Example: Configuring Bandwidth on Demand

Figure 75 on page 831 illustrates a bandwidth-on-demand configuration using multiple physical ISDN interfaces.

Figure 75: Bandwidth-on-Demand Topology



Configuration for the Dialer Interface

```
[edit interfaces]
d10 {
  encapsulation multilink-ppp;
  unit 0 {
    dialer-options {
      pool 10;
      dial-string 5552222; #Phone number to be dialed
      load-threshold 95;#Dial more ISDN if load exceeds 95% of
      #current capacity
    }
  }
}
```

```

        fragment-threshold 1024; #Allowed only when dialer is in multilink mode
        mrru 1500; #Allowed only when dialer is in multilink mode
        encapsulation multilink-ppp;
        rtp {
            f-max-period 100;
            queues q3;
        }
    }
    family inet {
        negotiate-address;
    }
}

```

**Configuration for the
First Physical ISDN
Interface**

```

[edit interfaces]
br-3/0/0 {
    isdn-options {
        switch-type ni1;
    }
    dialer-options {
        pool 10;
    }
}

```

**Configuration for the
Second Physical ISDN
Interface**

```

[edit interfaces]
br-3/0/1 {
    isdn-options {
        switch-type ni1;
    }
    dialer-options {
        pool dialer-pool10;
    }
}

```

Configuring Dial-In and Callback

You can configure dial-in on the dialer interface to permit incoming calls. Using dial-in, all incoming calls on a BRI interface are mapped to a dialer interface based on a caller ID. The incoming call's caller ID is compared against all caller IDs configured on all dialers to find the valid match. Multiple caller IDs can be configured on a dialer interface. The same caller IDs cannot be configured on different dialers.

Instead of accepting incoming calls, you can configure the dialer interface to call back the caller. When callback is configured, the call is rejected, and after a brief delay the caller is called back using the dial-string configured on the dialer interface. Multiple dial-strings cannot be configured on a dialer when callback is configured.

To configure dial-in or callback, perform the steps in the following sections to configure the dialer interface and the physical ISDN interfaces:

- Configuring Dial-In on page 833
- Disabling Dial-In on page 833

- Configuring Callback on page 834
- Example: Configuring Dial-In and Callback on page 834

Configuring Dial-In

To configure the dialer interface for dial-in operation, include the `incoming-map` statement with options at the `[edit interfaces dln unit logical-unit-number dialer-options]` hierarchy level:

```
[edit interfaces dln unit logical-unit-number]
dialer-options {
  incoming-map {
    caller (caller-id | accept-all);
  }
}
```



NOTE: The `incoming-map` statement is mandatory for the router to accept any incoming ISDN calls.

Include the option `accept-all` to accept all incoming calls. You can configure the `accept-all` option for only one of the dialer interfaces associated with an ISDN physical interface. The dialer interface with the `accept-all` option configured will be used only if the incoming call's caller ID does not match against the caller IDs configured on other dialer interfaces.

Include the `caller caller-id` statement to configure the dialer interface to accept calls from a specific caller ID. You can configure a maximum of 15 caller IDs per dialer interface.

The same caller ID cannot be configured on different dialer interfaces. However, you can configure a subset of the caller ID configured on another dialer interface. For example, you can configure the caller IDs 14085551515, 4085551515, and 5551515 on different dialer interfaces.

Disabling Dial-In

When dial-in is configured on the Services Router, incoming ISDN calls are accepted by the Services Router. However, you can configure the Services Router to reject all incoming ISDN calls when dial-in is configured.

To configure the Services Router to reject incoming ISDN calls, include the `reject-incoming` statement at the `[edit system processes isdn-signaling]` hierarchy level:

```
[edit system processes isdn-signaling]
reject-incoming;
```

For more information about disabling dial-in, see the *JUNOS System Basics Configuration Guide* and the *J-series Services Router Basic LAN and WAN Access Configuration Guide*.

Configuring Callback

To configure the dialer interface to call back a specific caller, include the caller *caller-id* statement and the *callback* statement at the [edit interfaces *dlIn* unit *logical-unit-number* *dialer-options*] hierarchy level:

```
[edit interfaces dlIn unit logical-unit-number]
dialer-options {
  incoming-map {
    caller caller-id;
    callback;
    callback-wait-period time;
  }
}
```

Include the optional *callback-wait-period* statement to change the time at which the dialer interface calls back the caller. The default period is 5 seconds.

Before configuring the callback on a dialer interface, ensure that:

- The dialer interface is not configured as a backup for a primary interface.
- The dialer interface does not have a watch list configured.
- Only one dial string is configured for the dialer interface.
- Dial-in is configured on the dialer interface of the remote router that is dialing in.

Example: Configuring Dial-In and Callback

The following illustrates configurations for dial-in and callback operations.

Configuration to Accept All Incoming Calls

```
[edit interfaces]
dl0 {
  encapsulation ppp;
  unit 0 {
    dialer-options {
      dial-string 7031231282;
      incoming-map;
      accept-all;
    }
    pool 2;
    family inet {
      address 10.1.1.2;
    }
  }
}
```

Configuration to Accept Calls from a Specific Caller ID

```
[edit interfaces]
dl0 {
  encapsulation ppp;
  unit 0 {
    dialer-options {
      incoming-map {
```

```

        caller 14082711234;
    }
    pool 1;
    family inet {
        address 10.2.1.1;
    }
}
}
}

```

**Configuration to Call
Back Calls from a
Specific Caller ID**

```

[edit interfaces]
dIO {
    encapsulation ppp;
    unit 0 {
        dialer-options {
            incoming-map {
                caller 14082711234;
            }
            callback;
            callback-wait-period 2;
            pool 1;
            family inet {
                address 10.2.1.1;
            }
        }
    }
}
}

```

Configuring Dialer Watch

Dialer watch is a feature that integrates backup dialing with routing capabilities and provides reliable connectivity without relying on “interesting” packets to trigger outgoing ISDN connections. With dialer watch, the Services Router monitors the existence of a specified route and if the route fails, the dialer interface initiates the ISDN connection as a backup connection.

To configure dialer watch, perform the steps in the following sections to configure the dialer interface and the physical ISDN interface:

- Configuring the Dialer Interface on page 835
- Configuring the Physical Interface on page 836
- Example: Configuring Dialer Watch on page 836
- Example: Complete ISDN Called-Calling Router Configuration on page 837

Configuring the Dialer Interface

To configure the dialer interface for dialer watch, include the following statements at the [edit interfaces *dln*] and the [edit interfaces *dln* unit *logical-unit-number*] hierarchy levels:

```

[edit interfaces]
dln {

```

```

encapsulation (cisco-hdlc | multilink-ppp | ppp);
hold-time (up | down) milliseconds;
unit logical-unit-number {
    dialer-options {
        activation-delay seconds;
        deactivation-delay seconds;
        dial-string dial-string-numbers;
        hold-time seconds;
        initial-route-check seconds;
        pool pool-name;
        watch-list {
            [ routes ];
        }
        family family {
            ip-address;
        }
    }
}

```

Configuring the Physical Interface

To configure the physical interface for dialer watch, include the **pool** statement at the [edit interfaces *br-pim*/O/*port* dialer-options] hierarchy level:

```

[edit interfaces]
br-pim/O/port {
    dialer-options {
        pool name;
    }
}

```


Each physical interface must use the same pool to participate in dialer watch.

Example: Configuring Dialer Watch

The following example illustrates a dialer watch configuration using one physical interface and one dialer interface.

See Figure 76 on page 836 for the topology used in this example.

Figure 76: Dialer Watch Topology



Configuration for the Physical Interface

```

[edit interfaces]
br-2/0/0 {
    isdn-options {
        switch-type ntdms100;
    }
}

```

```

        dialer-options {
            pool 1 priority 1;
        }
    }

Configuration for the  
Dialer Interface
[edit interfaces]
dlo {
    unit 0 {
        dialer-options {
            pool 1;
            dial-string 384030;
            watch-list {
                2.2.2.2/24;
                3.3.3.3/24;
            }
        }
        family inet {
            address 40.0.0.1/8;
        }
    }
}

```

Example: Complete ISDN Called-Calling Router Configuration

This example configures the calling J Series router (R1) and the calling J Series router (R2). The routers are both directly connected to an ISDN switch.

```

Configuration of Calling  
Router (R1)
[edit]
system {
    login {
        user isdn {
            uid 2000;
            class super-user;
            authentication {
                encrypted-password "$1$IS8Vkg3V$tzySvfBSZh1vYHSZQ6fM1";
                ## SECRET-DATA
            }
        }
    }
    services {
        web-management {
            http;
        }
    }
}

interfaces {
    fe-0/0/0 {
        unit 0 {
            family inet {
                address 192.168.1.1/24;
            }
        }
    }
}

```

```

br-3/0/0 {
  traceoptions {
    flag q921;
    flag q931;
    file {
      isdn_logg;
    }
  }
  isdn-options {
    switch-type etsi;
    spid1 116;
  }
  dialer-options {
    pool 100;
  }
}
dl100 {
  encapsulation ppp;
  unit 0 {
    dialer-options {
      pool 100;
      dial-string 119;
    }
    family inet {
      filter {
        dialer nss;
      }
      address 10.1.1.1/24;
    }
  }
}

firewall {
  family inet {
    dialer-filter nss {
      term 1 {
        from {
          destination-address {
            10.1.1.0/24;
          }
        }
        then note;
      }
    }
  }
}

access {
  profile isdn {
    client isdn chap-secret "$9$Lpax7VsYoGUHwsP5F39C"; ## SECRET-DATA
  }
}

```

**Configuration of Called
Router (R1)**

```

[edit]
system {
  root-authentication {
    encrypted-password "$1$UfcFhjcm$ftfgaLjMgRvFhrT3obrHu."; ## SECRET-DATA
  }
  services {
    web-management {
      http {
        interface [ fe-0/0/0.0 fe-0/0/1.0 ];
      }
    }
  }
  syslog {
    user * {
      any emergency;
    }
    file messages {
      any any;
      authorization info;
    }
    file interactive-commands {
      interactive-commands any;
    }
  }
}

interfaces {
  br-0/0/4 {
    isdn-options {
      switch-type etsi;
      spid1 119;
      tei-option power-up;
    }
    dialer-options {
      pool 100;
    }
  }
  dl100 {
    encapsulation ppp;
    unit 0 {
      dialer-options {
        pool 100;
        dial-string 116;
        incoming-map {
          caller 116;
        }
      }
      family inet {
        filter {
          dialer nss;
        }
        address 10.1.1.2/24;
      }
    }
  }
}

```

```

}

firewall {
  family inet {
    dialer-filter nss {
      term 1 {
        from {
          address {
            10.1.1.0/24;
          }
        }
        then note;
      }
    }
  }
}

```

Disabling ISDN Processes

You can disable ISDN entirely or disable certain processes at the system process level.

To disable ISDN entirely, include the **disable** statement at the **[edit system processes isdn-signaling]** hierarchy level:

```

[edit system processes isdn-signaling]
disable;

```

To disable the dial-out on demand process, include the **disable** statement at the **[edit system processes dialer-services]** hierarchy level:

```

[edit system processes dialer-services]
disable;

```

To disable dial-in and force the Services Router to reject incoming ISDN calls, include the **reject-incoming** statement at the **[edit system processes isdn-signaling]** hierarchy level:

```

[edit system processes isdn-signaling]
reject-incoming;

```


Part 12

Configuring SONET Interfaces

- Configuring SONET/SDH Interfaces on page 843

Chapter 59

Configuring SONET/SDH Interfaces

This chapter discusses configuration of the SONET/SDH interface properties and provides configuration examples in the following sections:

- SONET/SDH Interfaces Overview on page 843
- Configuring SONET/SDH Physical Interface Properties on page 844
- Configuring the Media MTU on SONET/SDH Interfaces on page 873
- Enabling Passive Monitoring on SONET/SDH Interfaces on page 874
- Configuring the Clock Source on SONET/SDH Interfaces on page 875
- Configuring Receive and Transmit Leaky Bucket Properties on SONET/SDH Interfaces on page 876
- Damping Interface Transitions on SONET/SDH Interfaces on page 877
- Configuring Interface Encapsulation on SONET/SDH Interfaces on page 878
- Example: Configuring SONET/SDH Interfaces on page 881
- Configuring Aggregated SONET/SDH Interfaces on page 881

SONET/SDH Interfaces Overview

Synchronous Digital Hierarchy (SDH) is a CCITT standard for a hierarchy of optical transmission rates. Synchronous Optical Network (SONET) is a USA standard that is largely equivalent to SDH. Both are widely used methods for very high speed transmission of voice and data signals across the numerous world-wide fiber-optic networks.

SDH and SONET use light-emitting diodes or lasers to transmit a binary stream of light-on and light-off sequences at a constant rate. At the far end optical sensors convert the pulses of light back to electrical representations of the binary information.

In wavelength-division multiplexing (WDM), light at several different wavelengths (colors to a human eye) is transmitted on the same fiber segment, greatly increasing the throughput of each fiber cable.

In dense wavelength-division multiplexing (DWDM), many optical data streams at different wavelengths are combined into one fiber.

The basic building block of the SONET/SDH hierarchy in the optical domain is an OC1; in the electrical domain, it is an STS-1. An OC1 operates at 51.840 Mbps. OC3 operates at 155.520 Mbps.

A SONET/SDH stream can consist of discrete lower-rate traffic flows that have been combined using time-division multiplexing (TDM) techniques. This method is useful, but a portion of the total bandwidth is consumed by the TDM overhead. When a SONET/SDH stream consists of only a single, very high speed payload, it is referred to as operating in concatenated mode. A SONET/SDH interface operating in this mode has a “c” added to the rate descriptor. For example, a concatenated OC48 interface is referred to as OC48c.

SONET and SDH traffic streams exhibit very few differences in behavior that are significant to Juniper Networks SONET/SDH interfaces; in general, this chapter uses *SONET/SDH* to indicate behavior that is identical for the two standards. However, there is one important difference that requires you to configure the interface specifically for SONET or SDH mode. That difference is in the setting of two bits (the ss-bits) in the pointer. SONET equipment ignores these bits, but SDH equipment uses them to distinguish a VC-4 payload from other types. When configured in SDH mode, Juniper Networks SONET/SDH PICs set the ss-bits to **s1s0 2** (binary 10). For more information, see the *JUNOS System Basics Configuration Guide*.



CAUTION: To extend the life of the laser, when a SONET/SDH PIC is not being actively used with any valid links, take the PIC offline until you are ready to establish a link to another device. To do this, issue the **request chassis pic offline fpc-slot slot-number pic-slot slot-number** operational mode command:

```
user@host> request chassis pic offline fpc-slot slot-number pic-slot slot-number
```

After you have connected the PIC to another device, bring the PIC back online by issuing the **request chassis pic online fpc-slot slot-number pic-slot slot-number** operational mode command.

```
user@host> request chassis pic online fpc-slot slot-number pic-slot slot-number
```

For information about taking a PIC offline or online, see the **request chassis pic offline** command and the **request chassis pic online** command in the *JUNOS System Basics and Services Command Reference*.

Configuring SONET/SDH Physical Interface Properties

To configure SONET/SDH physical interface properties, include the **sonet-options** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces so-fpc/pic/port]
framing (sdh | sonet);
sonet-options {
  aggregate asx;
  aps {
    advertise-interval milliseconds;
    annex-b
    authentication-key key;
    force;
    hold-time milliseconds;
    lockout;
```

```

neighbor address;
paired-group group-name;
protect-circuit group-name;
request;
revert-time seconds;
switching-mode (bidirectional | unidirectional);
working-circuit group-name;
}
bytes {
  e1-quiet value;
  f1 value;
  f2 value;
  s1 value;
  z3 value;
  z4 value;
}
fcs (16 | 32);
loopback (local | remote);
mpls {
  pop-all-labels {
    required-depth number;
  }
}
path-trace trace-string;
(payload-scrambler | no-payload-scrambler);
rfc-2615;
trigger {
  defect ignore;
  defect hold-time up milliseconds down milliseconds;
}
}
vtmapping (itu-t | klm);
(z0-increment | no-z0-increment);
speed (oc3 | oc12 | oc48);

```

Note that when you configure SONET/SDH OC48 interfaces for channelized (multiplexed) mode (by including the **no-concatenate** statement at the [edit chassis fpc slot-number pic pic-number] hierarchy level), the **bytes f1** statement has no effect. Currently, the **bytes e1-quiet** statement is ignored if you include it in the configuration. The **bytes f2**, **bytes z3**, **bytes z4**, and **path-trace** options work correctly on channel 0 and work in the transmit direction only on channels 1, 2, and 3. When using **no-concatenate**, you must specify a channel. For more information, see the *JUNOS System Basics Configuration Guide*.

For DS3 channels on a channelized OC12 interface, the **bytes f1**, **bytes f2**, **bytes z3**, and **bytes z4** options have no effect. The **bytes s1** option is supported only for channel 0; it is ignored if configured on channels 1 through 11. The **bytes s1** value configured on channel 0 applies to all channels on the interface.

You can also include some of the statements in the **sonet-options** statement to set SONET/SDH parameters on ATM interfaces.

You can configure the following SONET/SDH physical interface properties:

- Configuring SONET/SDH Framing on page 846
- Configuring SONET/SDH Interface Speed on page 847
- Configuring SONET/SDH Header Byte Values on page 849
- Configuring an Incrementing STM ID on page 850
- Configuring the SONET/SDH Frame Checksum on page 851
- Configuring Channelized IQ and IQE SONET/SDH Loop Timing on page 852
- Configuring SONET/SDH Loopback Capability on page 852
- Configuring the SONET/SDH Path Trace Identifier on page 853
- Configuring SONET/SDH HDLC Payload Scrambling on page 854
- Configuring SONET/SDH RFC 2615 Support on page 855
- Configuring SONET/SDH Defect Triggers to Be Ignored on page 855
- Configuring SONET/SDH Defect Hold Times on page 856
- Configuring Virtual Tributary Mapping on page 858
- Configuring APS and MSP on page 859
- Configuring SONET Options for 10-Gigabit Ethernet Interfaces on page 872

Configuring SONET/SDH Framing

The 4-port OC48 PIC with SFP installed, the next-generation SONET/SDH PICs with SFP, and the 4-port OC192 PIC on M Series, MX Series, and T Series routers, support SONET or SDH framing on a per-port basis. This functionality allows you to mix SONET and SDH modes on interfaces on a single PIC. You can use the **framing** statement to configure incoming SDH links from Europe and outgoing SONET links to the US on the same PIC. Traffic flowing through other ports of the same PIC will not be affected.

When you change SONET/SDH mode on a port, only the port's framing type is changed. The PIC does not go offline.

To configure framing on a per-port basis, include the **framing (sdh | sonet)** statement at the **[edit interfaces so-fpc/pic/port]** hierarchy level:

```
[edit interfaces]
so-fpc/pic/port {
    framing (sdh | sonet);
}
```



NOTE: Per-port framing configuration is applicable for SONET interfaces in concatenated mode (default mode) only. When you configure a PIC to operate in nonconcatenated mode, the individual channels inherit framing configuration from the **[edit chassis fpc number pic number framing (sonet | sdh)]** hierarchy level.



NOTE: Automatic Protection Switching (APS) is used by SONET add/drop multiplexers (ADMs) to protect against circuit failures. If APS is configured, and you do not change the SONET/SDH mode on both the working and protection port, APS support will not function properly. Both the working and protection ports must have the same mode configuration.

To view interface information, use the operational mode command **show interfaces so-fpc/pic/port**.

Configuring SONET/SDH Interface Speed

You can configure the speed of SONET/SDH interfaces on next-generation SONET/SDH Type 1 and Type 2 PICs with SFP. The speed you select is dependent upon whether the PIC is in concatenated or nonconcatenated mode. In concatenated mode, the bandwidth of the interface is in a single channel. In nonconcatenated mode, the PIC operates in channelized (multiplexed) mode.

Table 72 on page 847 shows the mode combinations for the next-generation SONET/SDH Type 1 PICs with SFP.

Table 72: Type 1 PIC Mode Combinations

PIC	Mode	Speed Configuration	Default Mode
2-port OC3	2xOC3 concatenated	<i>fpc/pic/0 speed oc3</i>	Concatenated
4-port OC3	1xOC12 concatenated	<i>fpc/pic/0 speed oc12</i>	
	1xOC12 nonconcatenated	<i>fpc/pic/0:0 speed oc3</i>	Nonconcatenated
	4xOC3 concatenated	<i>fpc/pic/port speed oc3</i>	Concatenated
1-port OC12	1xOC12 concatenated	<i>fpc/pic/0 speed oc12</i>	Concatenated
	1xOC12 nonconcatenated	<i>fpc/pic/0:0 speed oc3</i>	Nonconcatenated
	1xOC3 concatenated	<i>fpc/pic/0 speed oc3</i>	

Table 73 on page 848 shows the mode combinations for the next-generation SONET/SDH Type 2 PICs with SFP.

Table 73: Type 2 PIC Mode Combinations

PIC	Mode	Speed Configuration	Default Mode
1-port OC48	1xOC48 concatenated	<i>fpc/pic/0 speed oc48</i>	Concatenated
	1xOC48 nonconcatenated	<i>fpc/pic/0:0 speed oc12</i>	Nonconcatenated
	1xOC12 concatenated	<i>fpc/pic/0 speed oc12</i>	
	1xOC12 nonconcatenated	<i>fpc/pic/0 0 speed oc3</i>	
	1xOC3 concatenated	<i>fpc/pic/0 speed oc3</i>	
4-port OC12	1xOC48 concatenated	<i>fpc/pic/0 speed oc48</i>	
	1xOC48 nonconcatenated	<i>fpc/pic/0:0 speed</i>	Nonconcatenated
	1xOC12 nonconcatenated	<i>fpc/pic/0 speed oc3</i>	
	4xOC12 concatenated	<i>fpc/pic/port speed oc3 oc12</i>	Concatenated
4-port OC3	1xOC12 concatenated	<i>fpc/pic/0 speed oc12</i>	
	1xOC12 nonconcatenated	<i>fpc/pic/0:0 speed oc3</i>	Nonconcatenated
	4xOC3 concatenated	<i>fpc/pic/port speed oc3</i>	Concatenated

By default, SONET/SDH PICs operate in concatenated mode. To specify interface speed in concatenated mode, include the **speed** statement with options at the [edit interfaces *so-fpc/pic/port*] hierarchy level:

```
[edit interfaces so-fpc/pic/port
speed (oc3 | oc12 | oc48);
```

For example, each port of 4-port OC12 PIC can be configured to be in OC3 or OC12 speed independently when this PIC is in 4xOC12 concatenated mode.

To specify interface speed in nonconcatenated mode, include the **speed** statement at the [edit interfaces *so-fpc/pic/port.channel*] hierarchy level:

```
[edit interfaces so-fpc/pic/port.channel]
speed (oc3 | oc12);
```

To configure the PIC to operate in channelized (multiplexed) mode, include the **no-concatenate** statement at the [edit chassis *fpc slot-number pic pic-number*] hierarchy level.

For more information about using the **no-concatenate** statement, see the *JUNOS System Basics Configuration Guide*.

Configuring SONET/SDH Header Byte Values

To configure values in SONET/SDH header bytes, include the **bytes** statement at the [edit interfaces *interface-name* sonet-options] hierarchy level:

```
[edit interfaces so-fpc/pic/port sonet-options]
bytes {
  c2 value;
  e1-quiet value;
  f1 value;
  f2 value;
  s1 value;
  z3 value;
  z4 value;
}
```

You can configure the following SONET/SDH header bytes:

- **c2**—Path signal label SONET/SDH overhead byte. SONET/SDH frames use the C2 byte to indicate the contents of the payload inside the frame. SONET/SDH interfaces use the C2 byte to indicate whether the payload is scrambled. For the c2 byte, *value* can be from 0 through 255. The default value is 0xCF.
- **e1-quiet**—Default idle byte sent on the orderwire SONET/SDH overhead bytes. The router does not support the orderwire channel, and hence sends this byte continuously.
- **f1, f2, z3, z4**—SONET/SDH overhead bytes. For these bytes, *value* can be from 0 through 255. The default value is 0x00.
- **s1**—Synchronization message SONET/SDH overhead byte. This byte is normally controlled as a side effect of the system reference clock configuration and the state of the external clock coming from an interface if the system reference clocks have been configured to use an external reference. For the s1 byte, *value* can be from 0 through 255.

Table 74 on page 849 displays JUNOS Software framing bytes for several specific speeds.

Table 74: SONET/SDH Framing Bytes for Specific Speeds

Overhead Bytes	STM4	STM16	STM64	OC12	OC48	OC192
A1	F6	F6	F6	F6	F6	F6
A2	28	28	28	28	28	28
C1	—	—	—	1..12	1..48	1..192
H1/H2	6A0A	6A0A	6A0A	620A	620A	620A
Z0	01/CC	01/CC	01/CC	—	—	—

Table 74: SONET/SDH Framing Bytes for Specific Speeds (*continued*)

Overhead Bytes	STM4	STM16	STM64	OC12	OC48	OC192
Concatenated mode	93FF	93FF	93FF	93FF	93FF	93FF

When you configure SONET/SDH header bytes, note the following:

- The C2 byte is the path signal label. If the C2 byte value on an interface does not match the C2 byte value on the remote interface, the path label mismatch (PLM-P) or unequipped (UNEQ-P) alarm might occur.
- When you configure SONET/SDH OC48 interfaces for channelized (multiplexed) mode (by including the `no-concatenate` statement at the `[edit chassis fpc slot-number pic pic-number]` hierarchy level), the `bytes f1` statement has no effect.
- Currently, the `bytes e1-quiet` statement is ignored if you include it in the configuration.
- The `bytes f2`, `bytes z3`, `bytes z4`, and `path-trace` options work correctly on channel 0 and work in the transmit direction only on channels 1, 2, and 3.
- For DS3 channels on a channelized OC12 interface, the `bytes f1`, `bytes f2`, `bytes z3`, and `bytes z4` options have no effect.
- The `bytes s1` option is supported only for channel 0; it is ignored if configured on channels 1 through 11. The `bytes s1` value configured on channel 0 applies to all channels on the interface.
- Embedded operations channel (EOC) D1, D2, and D3 bytes are not supported.
- For channelized OC12 IQE and channelized OC48 IQE PICs with SFPs:
 - Only C2 (Path signal label) and S1 byte setting is supported.
 - Following header bytes are not supported. The router will syslog an INFO message if a command for an unsupported header byte is received.

F1—Section user channel byte

F2—Path user channel byte

Z3, Z4—SONET/SDH overhead bytes

E1—quiet default idle byte

Configuring an Incrementing STM ID

When configured in SDH framing mode, SONET/SDH interfaces on a Juniper Networks router might not interoperate with some older versions of ADMs or regenerators that require an incrementing STM ID.

Current SDH standards specify a set of $3 * n$ overhead bytes in an $STMn$ that includes the J0 section trace byte. The rest are essentially unused (spare Z0) and contain

hexadecimal values (0x01, 0xCC, 0xCC ... 0xCC). The older version of the standard specified that the same set of bytes should contain an incrementing sequence: 1, 2, 3, ..., 3*n. Their use was still unspecified although they might have been used to assist in frame alignment. You can configure an incrementing STM ID to enable your Juniper Networks router to interoperate with older equipment that relies on these bytes for frame alignment.

The STM identifier has a precise definition in the SDH specifications. In ITU-T Recommendation G.707, *Network node interface for the synchronous digital hierarchy (SDH)* (03/96), Section 9.2.2.2.

You can explicitly configure an incrementing STM ID rather than a static one in the SDH overhead by including the `z0-increment` statement at the `[edit interfaces interface-name sonet-options]` hierarchy level. You should include this statement only for SDH mode; do not use it for SONET mode.

```
[edit interfaces so-fpc/pic/port sonet-options]
z0-increment;
```

To explicitly disable incrementing of the STM ID, include the following statement:

```
[edit interfaces so-fpc/pic/port sonet-options]
no-z0-increment;
```

Configuring the SONET/SDH Frame Checksum

By default, SONET/SDH interfaces use a 16-bit frame checksum. You can configure a 32-bit checksum, which provides more reliable packet verification. However, some older equipment might not support 32-bit checksums.

To configure a 32-bit checksum, include the `fcs` statement at the `[edit interfaces interface-name sonet-options]` hierarchy level:

```
[edit interfaces so-fpc/pic/port sonet-options]
fcs 32;
```

To return to the default 16-bit frame checksum, delete the `fcs 32` statement from the configuration:

```
[edit]
user@host# delete interfaces so-fpc/pic/port sonet-options fcs 32
```

To explicitly configure a 16-bit checksum, include the `fcs` statement at the `[edit interfaces interface-name sonet-options]` hierarchy level:

```
[edit interfaces so-fpc/pic/port sonet-options]
fcs 16;
```

On a channelized OC12 interface, the `sonet-options fcs` statement is not supported. To configure the frame checksum sequence (FCS) on each DS3 channel, you must include the `t3-options fcs` statement in the configuration for each channel.

Configuring Channelized IQ and IQE SONET/SDH Loop Timing

By default, internal clocking (line timing) is used on channelized IQ and IQE interfaces. To configure SONET/SDH or DS3-level clocking, include the `loop-timing` statement:

```
loop-timing;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *ct3-fpc/pic/port* t3-options]
- [edit interfaces *stm1-fpc/pic/port* sonet-options]

To explicitly configure the default line timing, include the `no-loop-timing` statement in the configuration:

```
no-loop-timing;
```

The `loop-timing` and `no-loop-timing` statements apply only to E1 and T1 interfaces you configure on channelized IQ and IQE PICs. If you attempt to include these statements on any other interface type, they are ignored.

For all channelized IQ and IQE PICs, the `clocking` statement is supported on all channels. To configure clocking on individual interfaces, include the `clocking` statement at the [edit interfaces *type-fpc/pic/port:channel*] hierarchy level. If you do not include the `clocking` statement, the individual interfaces use internal clocking by default.

For more information, see “Configuring the Clock Source” on page 128 and “Clock Sources on Channelized Interfaces” on page 390.

Configuring SONET/SDH Loopback Capability

To configure loopback capability on a SONET/SDH interface, include the `loopback` statement at the [edit interfaces *interface-name* sonet-options] hierarchy level:

```
[edit interfaces so-fpc/pic/port sonet-options]
loopback (local | remote);
```

To exchange BERT patterns between a local router and a remote router, include the `loopback remote` statement in the interface configuration at the remote end of the link. From the local router, issue the `test interface` command.

For more information about configuring BERT, see “Interface Diagnostics” on page 134. For more information about using operational mode commands to test interfaces, see the *JUNOS System Basics and Services Command Reference*.

To turn off the loopback capability, remove the `loopback` statement from the configuration:

```
[edit]
user@host# delete interfaces so-fpc/pic/port sonet-options loopback
```

For channel 0 on channelized interfaces only, you can include the **loopback** statement at the [edit interfaces *interface-name* *interface-type-options*] hierarchy level. The loopback setting configured for channel 0 applies to all channels on the channelized interface. The **loopback** statement is ignored if you include it at this hierarchy level in the configuration of other channels. To configure loopbacks on individual channels, you must include the *channel-type-options* **loopback** statement in the configuration for each channel. This allows each channel to be put in loopback mode independently.

For example, for DS3 channels on a channelized OC12 interface, the **sonet-options loopback** statement is supported only for channel 0; it is ignored if included in the configuration for channels 1 through 11. The SONET/SDH loopback configured for channel 0 applies to all 12 channels equally. To configure loopbacks on the individual DS3 channels, you must include the **t3-options loopback** statement in the configuration for each channel. This allows each DS3 channel can be put in loopback mode independently.

You can determine whether there is an internal problem or an external problem by checking the error counters in the output of the **show interface *interface-name* extensive** command:

```
user@host> show interfaces so-fpc/pic/port extensive
```

Example: Configuring SONET/SDH Loopback Capability

To determine whether a problem is internal or external, loop packets on both the local and the remote router. To do this, include the **no-keepalives** and **encapsulation cisco-hdlc** statements at the [edit interfaces *interface-name*] hierarchy level, and the **loopback local** statement at the [edit interfaces *interface-name* sonet-options] hierarchy level. With this configuration, the link stays up, so you can loop ping packets to a remote router. The **loopback local** statement causes the interface to loop within the PIC just before the data reaches the transceiver.

```
[edit interfaces]
so-1/0/0 {
  no-keepalives;
  encapsulation cisco-hdlc;
  sonet-options {
    loopback local;
  }
  unit 0 {
    family inet {
      address 10.100.100.1/24;
    }
  }
}
```

Configuring the SONET/SDH Path Trace Identifier

The SONET/SDH *path trace identifier* is a text string that identifies the circuit. If the string contains spaces, enclose it in quotation marks.

By default, the JUNOS Software uses the router and interface names for the path trace identifier. Depending on the router and interface names, the default path trace

identifier might be longer than 16 bytes. The SDH standards define a maximum 16-byte path trace. For this reason, the default path trace identifier might be truncated in SDH mode. You can prevent the path trace identifier from being truncated in SDH mode by configuring a path trace identifier that is under 16-bytes long. In SONET mode, a path trace identifier can be up to 64-bytes long.

For DS3 channels on a channelized OC12 interface, you can configure a unique path trace for each of the 12 channels. Each path trace can be up to 16 bytes. For channels on a channelized OC12 intelligent queuing (IQ and IQE) interface, each path trace can be up to 64 bytes.

To configure a path trace identifier, include the **path-trace** statement at the [edit interfaces *interface-name* sonet-options] hierarchy level:

```
[edit interfaces interface-name sonet-options]
path-trace trace-string;
```

A common convention is to use the circuit identifier as the path trace identifier.

To display the local router's path trace identifier, issue the **show interfaces** command on the remote router.

Configuring SONET/SDH HDLC Payload Scrambling

SONET/SDH HDLC payload scrambling, which is enabled by default, provides better link stability. Both sides of a connection must either use or not use scrambling.



NOTE: HDLC payload scrambling conflicts with traffic shaping configured using leaky bucket properties. If you configure leaky bucket properties, you must disable payload scrambling, because the JUNOS Software rejects configurations that have both features enabled. For more information, see “Configuring Receive and Transmit Leaky Bucket Properties on SONET/SDH Interfaces” on page 876.

On a channelized OC12 interface, the **sonet-options payload-scrambler** statement is ignored. To configure scrambling on the DS3 channels on the interface, include the **t3-options payload-scrambler** statement in the configuration for each DS3 channel.

To disable HDLC payload scrambling, include the **no-payload-scrambler** statement at the [edit interfaces *interface-name* sonet-options] hierarchy level:

```
[edit interfaces so-fpc/pic/port sonet-options]
no-payload-scrambler;
```

To return to the default, that is, to re-enable payload scrambling, delete the **no-payload-scrambler** statement from the configuration:

```
[edit]
user@host# delete interfaces so-fpc/pic/port sonet-options no-payload-scrambler
```

To explicitly enable payload scrambling, include the **payload-scrambler** statement at the [edit interfaces *interface-name* sonet-options] hierarchy level:

```
[edit interfaces so-fpc/pic/port sonet-options]
payload-scrambler;
```

Configuring SONET/SDH RFC 2615 Support

RFC 2615, *PPP over SONET/SDH*, requires certain C2 header byte and FCS settings that vary from the default values configured in accordance with RFC 1619 (the previous version of RFC 2615). The newer values are optimized for stronger error detection, especially when combined with payload scrambling at higher bit rate links.

Table 75 on page 855 shows the older (RFC 1619) and newer (RFC 2615) values, together with the Juniper Networks default values.

Table 75: SONET/SDH Default Settings

Value	RFC 1619	Default	RFC 2615
SONET/SDH C2 header byte	0XCF	0XCF	0X16
Frame checksum (bit)	16	16	32
Payload scrambling	n/a	Enabled	Enabled

To enable support for the RFC 2615 features, include the `rfc-2615` statement at the [edit interfaces *interface-name* sonet-options] hierarchy level:

```
[edit interfaces so-fpc/pic/port sonet-options]
rfc-2615;
```

Configuring SONET/SDH Defect Triggers to Be Ignored

A trigger is a defect alarm that causes a physical interface to be marked down. By default, all defects are honored with no hold time. For SONET/SDH and ATM over SONET/SDH interfaces only, you can configure individual triggers to ignore a defect, honor a defect, and apply up and down hold timers to the defect.

Table 76 on page 855 lists the defects you can configure.

Table 76: SONET/SDH and ATM Active Alarms and Defects

Alarm	Description
Physical	
pll	Phase-locked loop out of lock
lol	Loss of light
Section	
lof	Loss of frame
los	Loss of signal

Table 76: SONET/SDH and ATM Active Alarms and Defects (*continued*)

Alarm	Description
Line	
ais-l	Alarm indication signal—line
rfi-l	Remote failure indication—line
ber-sd	Bit error rate defect-signal degrade
ber-sf	Bit error rate fault-signal fail
Path	
ais-p	Alarm indication signal—path
locd (ATM only)	Loss of cell delineation
lop-p	Loss of pointer—path
plm-p	Payload label mismatch
rfi-p	Remote failure indication—path
uneq-p	Path unequipped

To configure defects to be ignored, include the **trigger** statement at the [edit interfaces *interface-name* sonet-options] hierarchy level:

```
[edit interfaces interface-name sonet-options]
trigger {
    defect ignore;
}
```

If you configure a defect to be ignored, that defect does not contribute to the interface being marked down or up.

After you configure a defect to be ignored, the JUNOS Software reevaluates the state of the defect on the interface. If the defect is outstanding and has caused the interface to be marked down, the interface is marked up.

When you configure a trigger on a low-level defect—for example, an LOS—only the low-level defect is affected. Higher-level defects that might result from the lower-level defect are not affected by the low-level trigger configuration. Therefore, you must configure higher-level defects as well.

Configuring SONET/SDH Defect Hold Times

By default, an interface is marked down as soon as a defect is detected, and is marked up as soon as the defect is absent. You might want to apply hold times to defects for the following reasons:

- To prevent route flaps from happening before a defect has been outstanding for a longer period than would be expected for an Automatic Protection Switching (APS) cutover
- To reduce the number of interface transitions



NOTE: On M Series and T Series routers with Channelized SONET IQ PICs and Channelized SONET IQE PICs, the SONET defect alarm trigger **hold-time** statement is not supported.

When you apply a “down” hold time to a defect, the defect must be present for at least the hold-time period before the interface is marked down. When you apply an “up” hold time to a defect, the defect must remain absent for at least the hold-time period before the interface is marked up, assuming no other defect is outstanding.

When you configure hold timers and the interface goes from up to down, the interface transition is not advertised to the rest of the system until the interface has remained down for the hold-time period. Similarly, when an interface goes from down to up, the interface transition is not advertised until the interface has remained up for the hold-time period.

To configure hold timers, include the **hold-time** statement at the [edit interfaces *interface-name* sonet-options trigger defect] hierarchy level:

```
[edit interfaces interface-name sonet-options trigger defect]
hold-time up milliseconds down milliseconds;
```

The time can be a value from 1 through 65,534 milliseconds.

When you configure defect hold times, you should note the following:

- You can configure an up hold time, a down hold time, or both.
- Each interface on a SONET/SDH PIC controls certain aspects of the SONET/SDH overhead. For example, when you configure an OC48 PIC to be nonconcatenated, four interfaces are created. Each interface has its own path overhead. However, all four path interfaces share the same physical, section, and line overhead. This means the following:
 - Each interface’s path trigger configuration is honored.
 - The physical, section, and line trigger configuration for the primary interface (*so-fpc/pic/slot:0*) is applied to all four interfaces.

Therefore, if you configure the *so-fpc/pic/slot:0* interface to have a hold time for the LOS trigger, when an LOS event occurs, all four interfaces remain up until the trigger expires, and then all four interfaces are marked down.

- The hold timers on the SONET/SDH defects are applied in addition to any other hold timers you configure on the interface. For example, if an interface is up and you configure a SONET/SDH trigger down hold time of 100 milliseconds and an interface down hold time of 250 milliseconds, when the SONET/SDH defect occurs, the SONET/SDH trigger timer starts. After 100 milliseconds, assuming

the defect is still present, the SONET/SDH defect starts the 250 millisecond down timer. After this has expired and again assuming the defect is still outstanding, the interface will be marked down. For more information about interface hold timers, see “Damping Interface Transitions” on page 138.

- Some defects are reported through a periodic poll (once every second). For these defects, there could be up to one second lost before the defect is detected and the hold timer is started. The hold timer expires in precisely the amount of time configured. At that point, the existence of the defect is checked again and the interface is marked up or down accordingly. These defects are as follows:
 - lol
 - pll
 - ber-sf
 - ber-sd
- We recommend the following settings:
 - Configure SONET/SDH defect timers on no more than 64 interfaces per FPC.
 - Configure a combined up hold time and down hold time for a SONET/SDH defect to be at least 100 milliseconds.

Example: Configuring SONET/SDH Defects to Be Ignored

Prevent an LOS from bringing down an interface. An LOS can lead to the following defects:

- AIS-L
- LOF
- PLL
- RFI-L
- RFI-P

```
[edit interfaces sonet-options trigger]
ais-l ignore;
lof ignore;
los ignore;
pll ignore;
rfi-l ignore;
rfi-p ignore;
```

Configuring Virtual Tributary Mapping

You can configure virtual tributary mapping to use KLM mode or ITU-T mode. By default, virtual tributary mapping uses KLM mode.

For the Channelized STM1 IQ and IQE PICs, you can configure virtual tributary mapping by including the `vtmapping` statement at the `[edit interfaces cau4-fpc/pic/port sonet-options]` hierarchy level:

```
[edit interfaces cau4-fpc/pic/port sonet-options]
vtmapping (klm | itu-t);
```

For the STM1 PIC, you can configure virtual tributary mapping by including the `vtmapping` statement at the `[edit chassis fpc slot-number pic pic-number]` hierarchy level:

```
[edit chassis fpc slot-number pic pic-number]
vtmapping (klm | itu-t);
```

Table 42 on page 473 lists the KLM mappings used by the Channelized STM1-to-E1 PIC interfaces.

Configuring APS and MSP

Automatic Protection Switching (APS) is used by SONET add/drop multiplexers (ADMs) to protect against circuit failures. The JUNOS implementation of APS allows you to protect against circuit failures between an ADM and one or more routers, and between multiple interfaces in the same router. When a circuit or router fails, a backup immediately takes over.



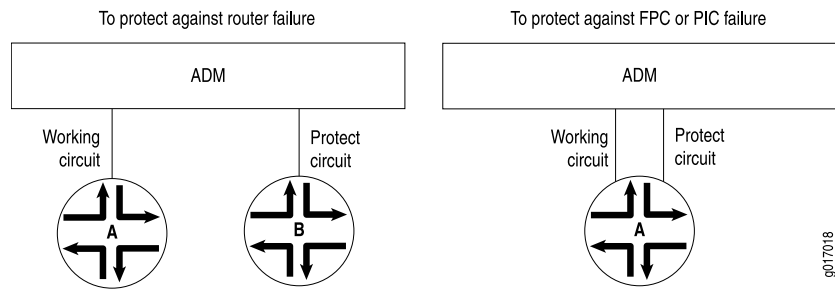
NOTE: For SDH interfaces, the JUNOS Software supports multiplex section protection (MSP). You configure MSP with the same CLI statements you use to configure APS.

The JUNOS Software supports APS 1 + 1 switching, either revertive or nonrevertive mode, and bidirectional mode only (although you can configure interoperability with line-terminating equipment [LTE] provisioned for unidirectional mode). The JUNOS Software does not transmit identical data on the working and protect circuits, as the APS specification requires for 1 + 1 switching, but this causes no operational impact.

For DS3 channels on a channelized OC12 interface, you can configure APS on channel 0 only. If you configure APS on channels 1 through 11, it is ignored.

With APS and MSP, you configure two circuits, a *working circuit* and a *protect circuit*. Normally, traffic is carried on the working circuit (that is, the working circuit is the active circuit), and the protect circuit is disabled. If the working circuit fails or degrades, or if the working router fails, the ADM and the protect router switch the traffic to the protect circuit, and the protect circuit becomes the active circuit.

To configure APS or MSP, you configure a working and a protect circuit, as shown in Figure 77 on page 860. To protect against a router failure, you connect two routers to the ADM, configuring one of them as the working router and the second as the protect router. To protect against a PIC or FPC failure, you connect one router to the ADM through both the working and protect circuits, configuring one of the PICs or FPCs as the working circuit and the second as the protect circuit.

Figure 77: APS/MSP Configuration Topologies

To configure APS or MSP, include the **aps** statement at the [edit interfaces *interface-name* sonet-options] hierarchy level:

```
[edit interfaces interface-name sonet-options]
aps {
  advertise-interval milliseconds;
  annex-b
  authentication-key key;
  force;
  hold-time milliseconds;
  lockout;
  neighbor address;
  paired-group group-name;
  protect-circuit group-name;
  request;
  revert-time seconds;
  switching-mode (bidirectional | unidirectional);
  working-circuit group-name;
}
```

This section includes the following topics:

- Configuring Basic APS Support on page 861
- Configuring Container Interfaces on page 863
- Configuring Switching Between the Working and Protect Circuits on page 866
- Configuring Revertive Mode on page 867
- Configuring Unidirectional Switching Mode Support on page 867
- Configuring APS Timers on page 868
- Configuring Link PIC Redundancy on page 869
- Example: Configuring Link PIC Redundancy on page 870
- Configuring APS Load Sharing Between Circuit Pairs on page 870
- Example: Configuring APS Load Sharing Between Circuit Pairs on page 872



NOTE: This implementation of APS is not supported on Layer 2 circuits. For Layer 2 circuits, configure APS by including the `protect-interface` statement. You can include this statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* protocols l2circuit neighbor *neighbor-id* interface *interface-name*]
- [edit protocols l2circuit neighbor *neighbor-id* interface *interface-name*]

For more information and a configuration example, see the *JUNOS VPNs Configuration Guide*.

When configuring the APS `annex-b` option, the APS options *must* be configured as follows:

- `switching-mode` *cannot* be uni-directional
- `revert-time` *cannot* be configured
- `manual-request` *cannot* be configured
- `exercise-request` *cannot* be configured
- `lockout-request` *cannot* be configured
- `wait-to-restore-time` is allowed *only* when Annex-B is configured
- `protect-circuit` *must* be configured
- `working-circuit` *must* be configured

Configuring Basic APS Support

To set up a basic APS configuration, configure one interface to be the working circuit and a second to be the protect circuit. If you are using APS to protect against router failure, configure one interface on each router. If you are using APS to protect against FPC failure, configure two interfaces on the router, one on each FPC.

For each working–protect circuit pair, configure the following:

- Group name—Creates the association between the two circuits. Configure the same group name for both the working and protect routers.
- Authentication key—You configure this on both interfaces. Configure the same key for both the working and protect routers.
- Address of the other interface on the other router—If you are configuring one router to be the working router and a second to be the protect router, you must configure the address of the remote interface. You configure this on one or both of the interfaces.

The address you specify for the neighbor must never be routed through the interface on which APS is configured, or instability will result. APS neighbor only applies to inter-router configurations. We strongly recommend that you directly connect the

working and protect routers and that you configure the interface address of this shared network as the neighbor address.

The working and protect configurations on the routers must match the circuit configurations on the ADM; that is, the working router must be connected to the ADM's working circuit and the protect router must be connected to the protect circuit.

To set up a basic APS configuration, include the following statements at the [edit interfaces *interface-name* sonet-options] hierarchy level:

On the Working Circuit	<pre>[edit interfaces so-fpc/pic/port sonet-options] aps { working-circuit group-name; authentication-key key; neighbor address; # Include if protect circuit is on a different router }</pre>
On the Protect Circuit	<pre>aps { protect-circuit group-name; authentication-key key; neighbor address; # Include if working circuit is on a different router }</pre>

Example: Configuring Basic APS Support

Configure Router A to be the working router and Router B to be the protect router.

On Router A (the Working Router)	<pre>[edit interfaces so-6/1/1 sonet-options] aps { working-circuit San-Jose; authentication-key " \$9\$B2612345" ; }</pre>
On Router B (the Protect Circuit)	<pre>[edit interfaces so-0/0/0 sonet-options] aps { protect-circuit San-Jose; authentication-key " \$9\$B2612345" ; neighbor 192.168.1.2;# Address of Router A on the link between A and B }</pre>
On a Single Platform, One Interface as the Working Circuit and Another Interface as the Protect Circuit	<pre>[edit interfaces so-2/1/1 sonet-options] aps { working-circuit bayward; authentication-key blarney; } [edit interfaces so-3/0/2 sonet-options] aps { protect-circuit bayward; authentication-key blarney; }</pre>

Configuring Container Interfaces

The JUNOS Software supports container interfaces for APS on SONET links. Physical interfaces and logical interfaces remain up on switchover, and their APS parameters are auto-copied from the container interface to the member links. See “Container Interfaces” on page 37 for more information.

Container interfaces support the following features:

- Cisco HDLC or PPP encapsulation methods.
- Unpaired groups.
- Bidirectional APS.
- Non-container and container-based APS on the same system.
- Use of any combination of (nonchannelized) SONET interfaces installed on the same router.

To configure a container interface, you must first create the number of container devices that you require. You can create up to a maximum of 128 container interfaces per router using the **device-count** statement at the **[edit chassis container-devices]** hierarchy level. You can create more container interfaces later if required, up to 128 (total). The resulting container interfaces are designated sequentially from **ci0** up to a maximum of **ci127**, depending on the **device-count number** specified. SONET interfaces can be assigned to any container interface **cin**.

To configure each container interface, you must assign two SONET interfaces (**so-fpc/pic/port**) using the **container-list cin** statement, and specify the **member-interface-speed speed** and **container-options** for each SONET interface.

Within each of the two SONET interfaces' container options, you must set one **container-type** as **primary** (corresponding to an APS working circuit) and the other as **standby** (corresponding to an APS protect circuit). For each SONET interface, you can also use the **allow-configuration-override** statement to allow the physical configuration of a member link to override the container configuration.

The following configuration steps are required:

1. Specify the total number of container interfaces (up to 128) to create using the **device-count number** statement at the **[edit chassis container-devices]** hierarchy level:

```
[edit chassis container-devices]
user@host# set device-count number
```

2. Configure the container interface parameters for a specified container **cin** as follows:
 - a. Specify the container interface using the numbered identifier **cin**:

```
[edit interfaces]
user@host# edit cin
```

- b. Specify the container interface encapsulation as `cisco-hdlc` or `ppp`:

```
[edit interfaces cin]
user@host# set encapsulation (cisco-hdlc | ppp)
```

- c. Specify the container options `container-type` as `aps`; a SONET interface is required for APS selection:

```
[edit interfaces cin]
user@host# set container-options container-type aps
```

- d. Specify the container interface member-interface type as `sonet`:

```
[edit interfaces cin]
user@host# set interfaces cin container-options member-interface-type sonet
```

- e. Specify the container member-interface-speed `speed` to match the specified installed SONET interface links; the available values are `OC3`, `OC12`, `OC48`, `OC192`, `OC768`, or `mixed`. The member-interface-speed `speed` statement setting applies to all SONET member interfaces of the specified container `cin`.

```
[edit interfaces cin]
user@host# set interfaces cin container-options member-interface-type sonet member-interface-speed speed
```

- f. Specify the container interface's unit number, family, IP address, and mask:

```
[edit interfaces cin]
user@host# set interfaces cin unit number family inet address ip-address/mask
```

- 3. Configure each of the required two SONET interfaces as follows:

- a. Specify the SONET interfaces and their container options; including the `container-list`, identified by its `cin`.
- b. Specify the `container-type` as *primary* (corresponding to an APS working-circuit) or *standby* (corresponding to an APS protect-circuit).

For example, setting `so-0/0/0` as the primary and `so-0/0/1` as the standby SONET interfaces for container interface `ci0`:

```
[edit]
user@host#edit interfaces so-0/0/0 # Enter config mode for interface so-0/0/0
[edit interfaces so-0/0/0]
user@host# set container-options container-list ci0 primary # Set so-0/0/0 as APS primary interface
[edit interfaces so-0/0/0]
user@host# top
[edit]
user@host#edit interfaces so-0/0/1 # Enter config mode for interface so-0/0/1
[edit interfaces so-0/0/1]
```



```
user@host# set container-options container-list ci0 standby # Set so-0/0/1
as APS standby interface
```

Optionally, you can set the `allow-configuration-override` statement to allow the physical configuration of a member link to override the container configuration:

```
[edit interfaces so-0/0/1]
user@host# set container-options container-list ci0 standby
allow-configuration-override
```

Example Container Interface Configuration

The following is a sample container interface configuration:

```
[edit chassis]
container-devices {
  device-count 1;
}
[edit interfaces]
so-1/0/2 {
  container-options {
    container-list ci0;
    primary;
  }
}
so-1/0/3 {
  container-options {
    container-list ci0;
    standby;
  }
}
ci0 {
  encapsulation cisco-hdlc;
  container-options {
    container-type aps {
      member-interface-type sonet {
        member-interface-speed mixed;
      }
    }
  }
  unit 0 {
    family inet {
      address 192.168.11.1/24;
    }
  }
}
```

You can run the `show aps` command to display the APS container interface configuration, as follows:

```
user@host> show aps
```

Interface	Group	Circuit	Intf state
ci0	CONTAINER_ci0	Container	enabled, up

so-1/2/2	MEMBER_OF_ci0	Working	enabled, up
so-1/2/3	MEMBER_OF_ci0	Protect	disabled, up

Configuring Switching Between the Working and Protect Circuits

When there are multiple reasons to switch between the working and protect circuits, a priority scheme is used to decide which circuit to use. The routers and the ADM might automatically switch traffic between the working and protect circuits because of circuit and router failures. You can also choose to switch traffic manually between the working and protect circuits.

When an ATM2 PIC is configured for APS, and the protect circuit comes online for the first time, there are no open VCs and the PIC discards the input traffic received on the protect circuit. The **show interface extensive** or **show monitor interface traffic** commands display the statistics as zero since the PIC drops the packets at the VC.

When the APS switches from the working circuit to the protect circuit, VCs are created on the protect circuit to accept traffic. However, the VCs on the working circuit remain open to support any future APS switches even though the interface is down or disabled. The input traffic received on the working circuit (current backup) is accepted by the PIC but discarded in the PFE. The **show interface extensive** or **show monitor interface traffic** commands displays live statistics for the traffic since it is accepted by the PIC.

When APS switches from the protect circuit to the working circuit again, the VCs on the protect circuit remain open to support a future APS switch even though the interface is down or disabled. The input traffic received on the current backup protect circuit is accepted by the PIC but discarded in the PFE. The **show interface extensive** or the **show monitor interface traffic** command displays live statistics for this traffic since it is accepted by the PIC.

There are three priority levels of manual configuration, listed here in order from lowest to highest priority:

- Request (also known as manual switch)—Overridden by signal failures, signal degradations, or any higher-priority reasons.
- Force (also known as forced switch)—Overrides manual switches, signal failures, and signal degradation.
- Lockout (also known as lockout of protection)—Do not switch between the working and protect circuits.



NOTE: Do not use the **disable** statement at the `[edit interfaces interface-name aps]` hierarchy level to switch between interface working and protect circuits; it can cause loss of traffic on the disabled interface. Use only the **request** statement or the **force** statement at the `[edit interfaces interface-name aps]` hierarchy level to modify interface status.

A router failure is considered to be equivalent to a signal failure on a circuit.

To perform a manual switch, include the **request** statement at the [edit interfaces *interface-name* sonet-options aps] hierarchy level. This statement is honored only if there are no higher-priority reasons to switch.

```
[edit interfaces so-fpc/pic/port sonet-options aps]
request (protect | working);
```

When the working circuit is operating in nonrevertive mode, use the **request working** statement to switch the circuit manually to being the working circuit or to override the revert timer.

To perform a forced switch, include the **force** statement at the [edit interfaces *interface-name* sonet-options aps] hierarchy level. This statement is honored only if there are no higher-priority reasons to switch. This configuration can be overridden by a signal failure on the protect circuit, thus causing a switch to the working circuit.

```
[edit interfaces so-fpc/pic/port sonet-options aps]
force (protect | working);
```

To configure a lockout of protection, forcing the use of the working circuit and locking out the protect circuit regardless of anything else, include the **lockout** statement at the [edit interfaces *interface-name* sonet-options aps] hierarchy level:

```
[edit interfaces so-fpc/pic/port sonet-options aps]
lockout;
```

Configuring Revertive Mode

By default, APS is nonrevertive, which means that if the protect circuit becomes active, traffic is not switched back to the working circuit unless the protect circuit fails or you manually configure a switch to the working circuit. In revertive mode, traffic is automatically switched back to the working circuit.

You should configure the ADM and routers consistently with regard to revertive or nonrevertive mode.

To configure revertive mode, include the **revert-time** statement, specifying the amount of time to wait after the working circuit has again become functional before making the working circuit active again:

```
[edit interfaces so-fpc/pic/port sonet-options aps]
revert-time seconds;
```

If you are using nonrevertive APS, you can use the **request working** statement to switch the circuit manually to being the working circuit or to override the revert timer (configured with the **revert-time** statement).

Configuring Unidirectional Switching Mode Support

You can configure interoperability with SONET/SDH Line Terminating Equipment (LTE) that is provisioned for unidirectional linear APS in 1 + 1 architecture on the following interfaces:

- Unchannelized OC3, OC12, and OC48 SONET/SDH interfaces on T Series routers

- SONET/SDH interfaces on the M40e router
- ATM over SONET interfaces

By default, APS supports only SONET/SDH LTE that is provisioned for bidirectional mode.

In bidirectional switching mode, the working interface switches to the protect interface for both receipt and transmission of data, regardless of whether the signal failure is in the transmit or receive direction.

In true unidirectional mode, the working interface switches to the protect interface only for the direction in which signal failure occurs; for example, if there is a signal failure in the transmit direction, the working interface switches over to the protect interface for transmission but not receipt of data. When the protect interface operates in unidirectional mode, the working and protect interfaces must cooperate to operate the transmit and receive interfaces in a bidirectional fashion.

The JUNOS Software does not support true unidirectional mode. Instead the software supports interoperation with SONET/SDH LTE provisioned for unidirectional switching. This means that the SONET/SDH LTE on the router receives and transmits on one interface, even when you configure unidirectional support. The JUNOS implementation of unidirectional mode support allows the router to do the following:

- Accept a unidirectional mode as valid
- Trigger the peer (ADM) selector to switch receive from working interface to protect interface or the reverse
- Not send reverse requests to the far end (ADM)

To configure unidirectional mode support, include the **switching-mode unidirectional** statement, at the [edit interfaces *interface-name* sonet-options aps] hierarchy level:

```
[edit interfaces interface-name sonet-options aps]
switching-mode unidirectional;
```



NOTE: On interfaces with unidirectional APS support configured, revertive mode and load sharing between circuits are not supported.

To restore the default behavior, include the **switching-mode bidirectional** statement, at the [edit interfaces *interface-name* sonet-options aps] hierarchy level:

```
[edit interfaces interface-name sonet-options aps]
switching-mode bidirectional;
```

Configuring APS Timers

The protect and working routers periodically send packets to their neighbors to advertise that they are operational. By default, these advertisement packets are sent every 1000 milliseconds. A router considers its neighbor to be operational for a period, called the hold time, that is, by default, three times the advertisement interval.

If the protect router does not receive an advertisement packet from the working router within the hold time configured on the protect router, the protect router assumes that the working router has failed and becomes active.

APS is symmetric; either side of a circuit can time out the other side (for example, when detecting a crash of the other). Under normal circumstances, the failure of the protect router does not cause any changes because the traffic is already moving on the working router. However, if you had configured **request protect** and the protect router failed, the working router would enable its interface.

To modify the advertisement interval, include the **advertise-interval** at the [edit interfaces *interface-name* sonet-options aps] hierarchy level:

```
[edit interfaces so-fpc/pic/port sonet-options aps]
advertise-interval milliseconds;
```

To modify the hold time, include the **hold-time** at the [edit interfaces *interface-name* sonet-options aps] hierarchy level:

```
[edit interfaces so-fpc/pic/port sonet-options aps]
hold-time milliseconds;
```

The advertisement intervals and hold times on the protect and working routers can be different.

Configuring Link PIC Redundancy

Link state replication, also called interface preservation, is an addition to the SONET Automatic Protection Switching (APS) functionality that helps promote redundancy of link PICs used in LSQ configurations, providing MLPPP link redundancy at the port level.

Link state replication provides the ability to add two sets of links, one from the active SONET PIC and the other from the standby SONET PIC, to the same bundle. If the active SONET PIC fails, links from the standby PIC are used without link renegotiation. All the negotiated state is replicated from the active links to the standby links to prevent link renegotiation. For more information about LSQ configurations, see the *JUNOS Services Interfaces Configuration Guide*.

To configure link state replication, include the **preserve-interface** statement at the [edit interfaces *interface-name* sonet-options aps] hierarchy level on the interfaces on both PICs:

```
preserve-interface;
```

APS functionality must be available on the SONET PICs and the interface configurations must be identical on both ends of the link. Any configuration mismatch causes the commit operation to fail.

This feature is supported with SONET APS and the following link PICs:

- Channelized OC3 IQ and IQE PICs
- Channelized OC12 IQ and IQE PICs

■ Channelized STM1 IQ and IQE PICs

Link state replication supports MLPPP and PPP over Frame Relay (`frame-relay-ppp`) encapsulation, and fully supports GRES.

Example: Configuring Link PIC Redundancy

Configure link state replication configuration between the ports `coc3-1/0/0` and `coc3-2/0/0`.

```

interfaces {
  coc3-1/0/0 {
    sonet-options {
      aps {
        preserve-interface;
        working-circuit aps-group-1;
      }
    }
  }
  coc3-2/0/0 {
    sonet-options {
      aps {
        preserve-interface;
        working-circuit aps-group-1;
      }
    }
  }
}

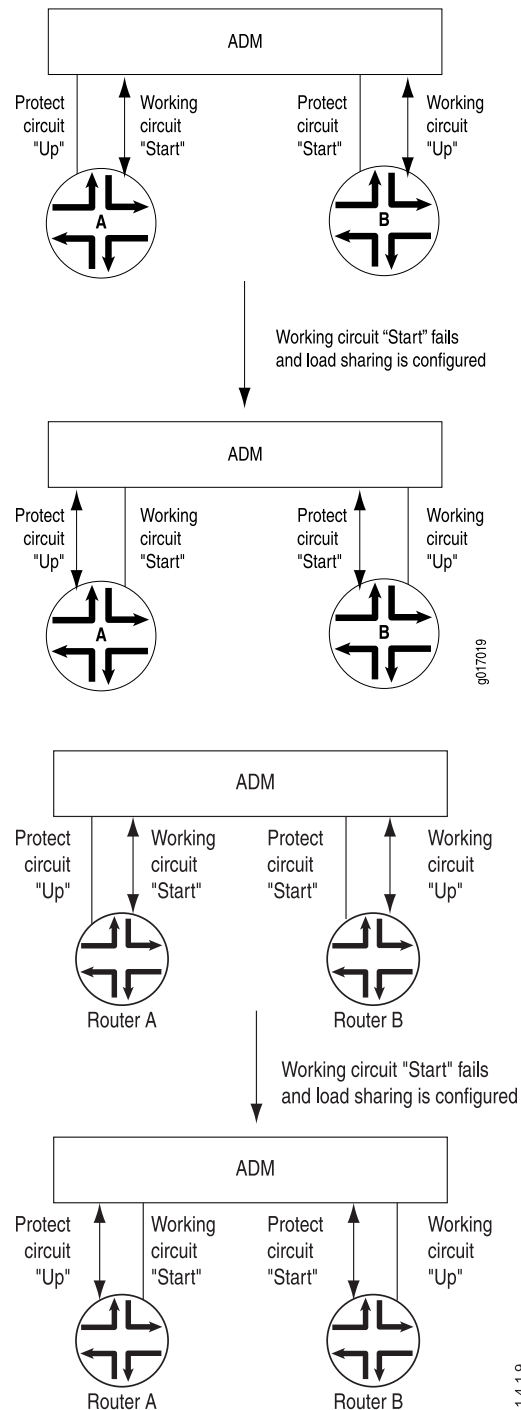
```

Configuring APS Load Sharing Between Circuit Pairs

When two routers are connected to a single ADM, you can have them back up each other on two different pairs of circuits. This arrangement provides load balancing between the routers if one of the working circuits fails.

Figure 78 on page 871 illustrates load sharing between circuits on two routers. Router A has a working circuit “Start” and a protect circuit “Up,” and Router B has a working circuit “Up” and a protect circuit “Start.” Under normal circumstances, Router A carries the “Start” circuit traffic and Router B carries the “Up” circuit traffic. If the working circuit “Start” were to fail, Router B would end up carrying all the traffic for both the “Start” and “Up” circuits.

To balance the load between the circuits, you pair the two circuits. In this case, you pair the “Start” and “Up” circuits. Then, if the working circuit “Start” fails, the two routers automatically switch the “Up” traffic from the working to the protect circuit so that each router is still carrying only one circuit’s worth of traffic. That is, the working circuit on Router A would be “Up” and the working circuit on Router B would be “Start.”

Figure 78: APS Load Sharing Between Circuit Pairs

To configure load sharing between two working-protect circuit pairs, include the **paired-group** statement when configuring one of the circuits on one of the routers. In this statement, the **group-name** is the name of the group you assigned to one of the circuits with the **working-circuit** and **protect-circuit** statements. The JUNOS Software

automatically configures the remainder of the load-sharing setup based on the group name.

```
[edit interfaces so-fpc/pic/port sonet-options aps]
paired-group group-name;
```

Example: Configuring APS Load Sharing Between Circuit Pairs

Configure APS load sharing to match the configuration shown in Figure 78 on page 871:

```
On Router A [edit interfaces so-7/0/0 sonet-options aps]
user@host# set working-circuit start
[edit interfaces so-7/0/0 sonet-options aps]
user@host# set authentication-key linsey
[edit interfaces so-7/0/0 sonet-options aps]
user@host# set paired-group "Router A-Router B"
...
[edit interfaces so-0/0/0 sonet-options aps]
user@host# set protect-circuit up
[edit interfaces so-0/0/0 sonet-options aps]
user@host# set authentication-key woolsey
[edit interfaces so-0/0/0 sonet-options aps]
user@host# set paired-group "Router A-Router B"

On Router B [edit interfaces so-1/0/0 sonet-options aps]
user@host# set working-circuit up
[edit interfaces so-1/0/0 sonet-options aps]
user@host# set authentication-key woolsey
[edit interfaces so-1/0/0 sonet-options aps]
user@host# set paired-group "Router A-Router B"
...
[edit interfaces so-6/0/0 sonet-options aps]
user@host# set protect-circuit start
[edit interfaces so-6/0/0 sonet-options aps]
user@host# set authentication-key linsey
[edit interfaces so-6/0/0 sonet-options aps]
user@host# set paired-group "Router A-Router B"
```

Configuring SONET Options for 10-Gigabit Ethernet Interfaces

The 10-Gigabit Ethernet IQ2 and IQ2-E PIC is supported on the M120, M320, and T Series routers. The PIC provides one external interface running at 10 Gbps. The interface operates in either LAN PHY or WAN PHY mode. When the external interface is running in WAN PHY mode, it uses the WIS sublayer to transport 10-Gigabit Ethernet frames in an OC192c SONET payload, and can interoperate with SONET section or line level repeaters. This creates an advantage when the interface is used for long-distance, point-to-point 10-Gigabit Ethernet links.

When the external interface is running in WAN PHY mode, you can configure specific physical SONET options. To configure SONET options, include the `loopback`, `mpls`, `path-trace`, and `trigger` statements at the `[edit interfaces interface-name sonet-options]` hierarchy level:


```

[edit interfaces]
xe-0/0/0 {
  sonet-options {
    loopback (local | remote);
    mpls {
      pop-all-labels {
        required-depth number;
      }
    }
    path-trace trace-string;
    trigger {
      defect ignore {
        defect hold-time up milliseconds down milliseconds;
      }
    }
  }
}

```

For information about using the `loopback` statement, see “Configuring SONET/SDH Loopback Capability” on page 852. For information about using the `mpls` statement, see “Removing MPLS Labels from Incoming Packets” on page 874. For information about using the `path-trace` statement, see “Configuring the SONET/SDH Path Trace Identifier” on page 853. For information about using the `trigger` statement, see “Configuring SONET/SDH Defect Triggers to Be Ignored” on page 855.

Configuring the Media MTU on SONET/SDH Interfaces

The default media MTU size used on a physical interface depends on the encapsulation being used on that interface. For a listing of MTU sizes for each encapsulation type, see “Configuring the Media MTU” on page 98. For information about configuring the encapsulation on an interface, see “Configuring Interface Encapsulation on SONET/SDH Interfaces” on page 878.

To modify the default media MTU size for a physical interface, include the `mtu` statement at the `[edit interfaces interface-name]` hierarchy level:

```

[edit interfaces interface-name]
mtu bytes;

```

If you change the size of the media MTU, you must ensure that the size is equal to or greater than the sum of the protocol MTU and the encapsulation overhead. You configure the protocol MTU by including the `mtu` statement at the following hierarchy levels:

- `[edit interfaces interface-name unit logical-unit-number family family]`
- `[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family]`

For more information, see “Setting the Protocol MTU” on page 191.

Enabling Passive Monitoring on SONET/SDH Interfaces

The Monitoring Services I and Monitoring Services II PICs are designed to enable IP services. If you have a Monitoring Services PIC and a SONET/SDH PIC installed in an M1 60, M40e, or T Series router, you can monitor IPv4 traffic from another router.

On SONET/SDH interfaces, you enable packet flow monitoring by including the `passive-monitor-mode` statement:

```
passive-monitor-mode;
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *so-fpc/pic/port* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *so-fpc/pic/port* unit *logical-unit-number*]

If you include the `passive-monitor-mode` statement in the configuration, the SONET/SDH interface does not send keepalives or alarms, and does not participate actively on the network.

On monitoring services interfaces, you enable packet flow monitoring by including the `family` statement at the [edit interfaces *mo-fpc/pic/port* unit *logical-unit-number*] hierarchy level, specifying the `inet` option:

```
[edit interfaces mo-fpc/pic/port unit logical-unit-number]  
family inet;
```

For conformity with cflowd record structure, you must include the `receive-options-packets` and `receive-ttl-exceeded` statements at the [edit interfaces *mo-fpc/pic/port* unit *logical-unit-number* family *inet*] hierarchy level:

```
[edit interfaces mo-fpc/pic/port unit logical-unit-number family inet]  
receive-options-packets;  
receive-ttl-exceeded;
```

For the monitoring services interface, you can configure multiservice physical interface properties. For more information, see “Configuring Multiservice Physical Interface Properties” on page 138 and the *JUNOS Services Interfaces Configuration Guide*.

Removing MPLS Labels from Incoming Packets

The JUNOS Software can forward only IPv4 packets to a Monitoring Services PIC. IPv4 packets with MPLS labels cannot be forwarded to a Monitoring Services PIC. By default, if packets with MPLS labels are forwarded to the Monitoring Services PIC, they are discarded. To monitor packets with MPLS labels, you must remove the MPLS labels as the packets arrive on the interface.

You can remove up to two MPLS labels from an incoming packet by including the `pop-all-labels` statement at the [edit interfaces *interface-name* sonet-options *mpls*] hierarchy level:

```
[edit interfaces interface-name sonet-options mpls]
pop-all-labels {
    required-depth number;
}
```

By default, the **pop-all-labels** statement takes effect for incoming packets with one or two labels. You can specify the number of MPLS labels an incoming packet must have for the **pop-all-labels** statement to take effect by including the **required-depth** statement at the `[edit interfaces interface-name atm-options mpls pop-all-labels]` hierarchy level:

```
[edit interfaces interface-name atm-options mpls pop-all-labels]
required-depth number;
```

The required depth can be 1, 2, or [1 2]. If you include the **required-depth 1** statement, the **pop-all-labels** statement takes effect for incoming packets with one label only. If you include the **required-depth 2** statement, the **pop-all-labels** statement takes effect for incoming packets with two labels only. If you include the **required-depth [1 2]** statement, the **pop-all-labels** statement takes effect for incoming packets with one or two labels. A required depth of [1 2] is equivalent to the default behavior of the **pop-all-labels** statement.

When you remove MPLS labels from incoming packets, note the following:

- The **pop-all-labels** statement has no effect on IP packets with three or more MPLS labels.
- When you enable MPLS label removal, you must configure all ports on a PIC with the same label popping mode and required depth.
- You use the **pop-all-labels** statement to enable passive monitoring applications, not active monitoring.
- You cannot apply MPLS filters or accounting to the MPLS labels because the labels are removed as soon as the packet arrives on the interface.

Configuring the Clock Source on SONET/SDH Interfaces

For interfaces such as SONET/SDH that can use different clock sources, you can configure the source of the transmit clock on each interface. The source can be internal or external. The default source is internal, which means that each interface uses the router's internal Stratum 3 clock.

For DS3 channels on a channelized OC12 interface, the **clocking** statement is supported only for channel 0; it is ignored if included in the configuration of channels 1 through 11. The clock source configured for channel 0 applies to all channels on the channelized OC12 interface. The individual DS3 channels use a gapped 45-MHz clock as the transmit clock. For more information, see "Clock Sources on Channelized Interfaces" on page 390.



NOTE: On channelized STM1 interfaces, you should configure the clock source at one side of the connection to be internal (the default JUNOS Software configuration) and configure the other side of the connection to be external.

To configure loop timing on an interface, include the **clocking external** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]  
clocking external;
```

To explicitly configure line timing on an interface, include the **clocking internal** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]  
clocking internal;
```

Configuring Receive and Transmit Leaky Bucket Properties on SONET/SDH Interfaces

Congestion control is particularly difficult in high-speed networks with high volumes of traffic. When congestion occurs in such a network, it is usually too late to react. You can avoid congestion by regulating the flow of packets into your network. Smoother flows prevent bursts of packets from arriving at (or being transmitted from) the same interface and causing congestion.

For all interface types except ATM, Fast Ethernet, Gigabit Ethernet, and channelized IQ and IQE, you can configure leaky bucket properties, which allow you to limit the amount of traffic received on and transmitted by a particular interface. You effectively specify what percentage of the interface's total capacity can be used to receive or transmit packets. You might want to set leaky bucket properties to limit the traffic flow from a link that is known to transmit high volumes of traffic.



NOTE: Instead of configuring leaky bucket properties, you can limit traffic flow by configuring policers. Policers work on all interfaces. For more information, see the *JUNOS Policy Framework Configuration Guide*.

The leaky bucket is used at the host-network interface to allow packets into the network at a constant rate. Packets might be generated in a bursty manner, but after they pass through the leaky bucket, they enter the network evenly spaced. In some cases, you might want to allow short bursts of packets to enter the network without smoothing them out. By controlling the number of packets that can accumulate in the bucket, the **threshold** property controls burstiness. The maximum number of packets entering the network in *t* time units is **threshold + rate * t**.

By default, leaky buckets are disabled and the interface can receive and transmit packets at the maximum line rate.

For each DS3 channel on a channelized OC12 interface, you can configure unique receive and transmit buckets.



NOTE: HDLC payload scrambling conflicts with traffic shaping configured using leaky bucket properties. If you configure leaky bucket properties, you must disable payload scrambling, because the JUNOS Software rejects configurations that have both features enabled. For more information, see “Configuring SONET/SDH HDLC Payload Scrambling” on page 854.

To configure leaky bucket properties, include one or both of the **receive-bucket** and **transmit-bucket** statements at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
receive-bucket {
  overflow (discard | tag);
  rate percentage;
  threshold bytes;
}
transmit-bucket {
  overflow discard;
  rate percentage;
  threshold bytes;
}
```

In the **rate** statement, specify the percentage of the interface line rate that is available to receive or transmit packets. The percentage can be a value from 0 (none of the interface line rate is available) to 100 (the maximum interface line rate is available). For example, when you set the line rate to 33, the interface receives or transmits at one third of the maximum line rate.

In the **threshold** statement, specify the bucket threshold, which controls the burstiness of the leaky bucket mechanism. The larger the value, the more bursty the traffic, which means that over a very short amount of time the interface can receive or transmit close to line rate, but the average over a longer time is at the configured bucket rate. The threshold can be a value from 0 through 65,535 bytes. For ease of entry, you can enter *number* either as a complete decimal number or as a decimal number followed by the abbreviation k (1,000). For example, the entry **threshold 2k** corresponds to a threshold of 2,000 bytes.

In the **overflow** option, specify how to handle packets that exceed the threshold:

- **tag**—(receive-bucket only) Tag, count, and process received packets that exceed the threshold.
- **discard**—Discard received packets that exceed the threshold. No counting is done.

Damping Interface Transitions on SONET/SDH Interfaces

By default, when an interface changes from being up to being down, or from down to up, this transition is advertised immediately to the hardware and the JUNOS Software. In some situations—for example, when an interface is connected to an add-drop multiplexer (ADM) or wavelength-division multiplexer (WDM), or to protect against SONET/SDH framer holes—you might want to damp interface transitions.

This means not advertising the interface's transition until a certain period of time has transpired.

When you have damped interface transitions and the interface goes from up to down, the interface is not advertised to the rest of the system as being down until it has remained down for the hold-time period. Similarly when an interface goes from down to up, it is not advertised as being up until it has remained up for the hold-time period.

To damp interface transitions, include the **hold-time** statement at the **[edit interfaces interface-name]** hierarchy level:

```
hold-time up milliseconds down milliseconds;
```

The time can be a value from 0 through 65,534 milliseconds. The default value is 0, which means that interface transitions are not damped. The JUNOS Software advertises the transition within 100 milliseconds of the time value you specify.

Configuring Interface Encapsulation on SONET/SDH Interfaces

Point-to-Point Protocol (PPP) encapsulation is the default encapsulation type for physical interfaces. You need not configure encapsulation for any physical interfaces that support PPP encapsulation. If you do not configure encapsulation, PPP is used by default. For physical interfaces that do not support PPP encapsulation, you must configure an encapsulation to use for packets transmitted on the interface. You can optionally configure an encapsulation on a logical interface, which is the encapsulation used within certain packet types.

Configuring the Encapsulation on a Physical SONET/SDH Interface

For SONET/SDH interfaces, the physical interface encapsulation can be one of the following:

- Point-to-Point Protocol (PPP)—PPP encapsulation is defined in RFC 1661, *The Point-to-Point Protocol (PPP) for the Transmission of Multiprotocol Datagrams over Point-to-Point Links*. PPP is the default encapsulation type for physical interfaces. Two related versions are supported:
 - Circuit cross-connect (CCC) version (**ppp-ccc**)—The logical interfaces do not require an encapsulation statement. When you use this encapsulation type, you can configure the **ccc** family only.
 - Translational cross-connect (TCC) version (**ppp-tcc**)—Similar to CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- Cisco HDLC—E1, E3, SONET/SDH, T1, and T3 interfaces can use Cisco HDLC encapsulation. Two related versions are supported:
 - CCC version (**cisco-hdlc-ccc**)—The logical interfaces do not require an encapsulation statement. When you use this encapsulation type, you can configure the **ccc** family only.

- TCC version (**cisco-hdlc-tcc**)—Similar to CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- Frame Relay—Defined in RFC 1490, *Multiprotocol Interconnect over Frame Relay*. E1, E3, SONET/SDH, T1, and T3 interfaces can use Frame Relay encapsulation. Two related versions are supported:
 - CCC version (**frame-relay-ccc**)—The same as standard Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to CCC. This numbering restriction does not apply to IQ and IQE interfaces. The logical interface must also have **frame-relay-ccc** encapsulation.
 - TCC version (**frame-relay-tcc**)—Similar to Frame Relay CCC and has the same configuration restrictions, but used for circuits with different media on either side of the connection.
- Frame Relay Ether Type (**frame-relay-ether-type**)—Physical interfaces can use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. IETF Frame Relay encapsulation identifies the payload format using NLPID and SNAP formats. Cisco-compatible Frame Relay encapsulation uses the Ethernet type to identify the type of payload. Two related versions are supported:
 - TCC version (**frame-relay-ether-type-tcc**)—Cisco-compatible Frame Relay for DLCIs 0 through 511. DLCIs 512 through 1022 are dedicated to TCC. This numbering restriction does not apply to IQ and IQE interfaces. This encapsulation is used for circuits with different media on either side of the connection.
 - Extended TCC version (**extended-frame-relay-ether-type-tcc**)—This encapsulation allows you to dedicate Cisco-compatible Frame Relay TCC for DLCIs 1 through 1022. This encapsulation is used for circuits with different media on either side of the connection. All ether type TCC encapsulation is supported on the same PICs as non-ether type Frame Relay TCC encapsulation.



NOTE: When the encapsulation type is set to Cisco-compatible Frame Relay encapsulation, ensure that the LMI type is set to ANSI or Q933-A.

To configure the encapsulation on a physical interface, include the **encapsulation** statement at the [edit interfaces *interface-name*] hierarchy level:

```
[edit interfaces interface-name]
encapsulation (cisco-hdlc | cisco-hdlc-ccc | cisco-hdlc-tcc | frame-relay | frame-relay-ccc |
frame-relay-tcc | frame-relay-tcc | ppp | ppp-ccc | ppp-tcc);
```

When you configure a point-to-point encapsulation (such as PPP or Cisco HDLC) on a physical interface, the physical interface can have only one logical interface (that is, only one unit statement) associated with it. When you configure a multipoint encapsulation (such as Frame Relay), the physical interface can have multiple logical units, and the units can be either point to point or multipoint. Use PPP if you are running Cisco IOS Release 12.0 or later. If you need to run Cisco HDLC, the JUNOS Software automatically configures an ISO family MTU of 4469 in the router. This is due to an extra byte of padding used by Cisco.

For more information about physical interface encapsulation, see “Configuring the Encapsulation on a Physical Interface” on page 106.

Example: Configuring the Encapsulation on a Physical SONET/SDH Interface

Configure PPP encapsulation on a SONET/SDH interface. The second two family statements allow IS-IS and MPLS to run on the interface.

```
[edit interfaces]
so-7/0/0 {
  encapsulation ppp;
  unit 0 {
    point-to-point;
    family inet {
      address 192.168.1.113/32 {
        destination 192.168.1.114;
      }
    }
    family iso;
    family mpls;
  }
}
```

Configuring the Encapsulation on a Logical SONET/SDH Interface

Generally, you configure an interface’s encapsulation at the [edit interfaces *interface-name*] hierarchy level. However, for Frame Relay encapsulation, you can also configure the encapsulation type that is used inside the Frame Relay packet itself. To do this, include the *encapsulation* statement, specifying the *frame-relay-ccc*, *frame-relay-tcc*, *frame-relay-ether-type*, or *frame-relay-ether-type-tcc* option:

```
encapsulation (frame-relay-ccc | frame-relay-tcc | frame-relay-ether-type |
  frame-relay-ether-type-tcc);
```

You can include this statement at the following hierarchy levels:

- [edit interfaces *interface-name* unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

The ATM encapsulations are defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*.

With the *atm-nlpid*, *atm-cisco-nlpid*, and *atm-vc-mux* encapsulations, you can configure the *inet* family only. With the circuit cross-connect (CCC) encapsulations, you cannot configure a family on the logical interface. A logical interface cannot have *frame-relay-ccc* encapsulation unless the physical device also has *frame-relay-ccc* encapsulation. A logical interface cannot have *frame-relay-tcc* encapsulation unless the physical device also has *frame-relay-tcc* encapsulation. In addition, you must assign this logical interface a DLCI from 512 through 1022. This numbering restriction does not apply to IQ and IQE interfaces. You must configure the logical interface as point-to-point.

For more information about logical interface encapsulation, see “Configuring the Encapsulation on a Logical Interface” on page 160.

Example: Configuring SONET/SDH Interfaces

SONET/SDH interfaces can use either PPP or Cisco HDLC encapsulation. Use PPP if you are running Cisco IOS Release 12.0 or later. If you need to run Cisco HDLC, the JUNOS Software automatically configures an ISO family MTU of 4469 in the router. This is due to an extra byte of padding used by Cisco. The following configuration, which uses PPP encapsulation, is sufficient to get a SONET/SDH OC3 or OC12 interface up and running:

```
[edit interfaces]
so-fpc/pic/port {
  encapsulation ppp;
  unit 0 {
    family inet {
      address local-address {
        destination remote-address;
      }
    }
  }
}
```

Configuring Aggregated SONET/SDH Interfaces

The JUNOS Software enables link aggregation of SONET/SDH interfaces; this is similar to Ethernet link aggregation, but is not defined in a public standard. The JUNOS Software balances traffic across the member links within an aggregated SONET/SDH bundle based on the Layer 3 information carried in the packet. This implementation uses the same load balancing algorithm used for per-packet load balancing. For information about per-packet load balancing, see the *JUNOS Routing Protocols Configuration Guide*.

You configure an aggregated SONET/SDH virtual link by specifying the link number as a physical device and then associating a set of physical interfaces that have the same speed. Channelized OC IQ and IQE PICs do not support SONET aggregation.

By default, no aggregated SONET/SDH interfaces are created. You must define the number of aggregated SONET/SDH interfaces by including the **device-count** statement at the [edit chassis aggregated-devices sonet] hierarchy level:

```
[edit chassis aggregated-devices sonet]
device-count number;
```

The maximum number of aggregated interfaces is 16. The aggregated SONET/SDH interfaces are numbered from **as0** through **as15**. For more information, see the *JUNOS Services Interfaces Configuration Guide*.



NOTE: SONET/SDH aggregation is proprietary to the JUNOS Software and might not work with other software.

To configure aggregated SONET/SDH interfaces, assign a number for the aggregated SONET/SDH interface `asx` at the `[edit interfaces]` hierarchy level:

```
[edit interfaces]
asx {
  ...
}
```

The following example shows an aggregated SONET/SDH configuration:

```
[edit interfaces]
as0 {
  aggregated-sonet-options {
    minimum-links 1;
    link-speed oc3;
  }
  unit 0 {
    family inet {
      address 10.2.11.1/30;
    }
  }
}
```

You also need to specify the constituent physical interfaces by including the **aggregate** statement at the `[edit interfaces interface-name sonet-options]` hierarchy level; for more information, see “Configuring SONET/SDH Link Aggregation” on page 882. You can optionally specify other physical properties that apply specifically to the aggregated SONET/SDH interfaces; for details, see “Configuring SONET/SDH Physical Interface Properties” on page 844. For a sample configuration, see “Example: Configuring Aggregated SONET/SDH Interfaces” on page 885.

To remove the configuration statements related to `asx` and set the aggregated SONET/SDH interface to down state, delete the interface from the configuration:

```
[edit]
user@host# delete interfaces asx
```

However, the aggregated SONET/SDH interface is not deleted until you delete the `chassis aggregated-devices sonet device-count` configuration statement.

You can configure the following aggregated SONET/SDH properties:

- Configuring SONET/SDH Link Aggregation on page 882
- Configuring Aggregated SONET/SDH Link Speed on page 883
- Configuring Aggregated SONET/SDH Minimum Links on page 883
- Configuring Filters or Sampling on Aggregated SONET/SDH Links on page 884
- Example: Configuring Aggregated SONET/SDH Interfaces on page 885

Configuring SONET/SDH Link Aggregation

On SONET/SDH interfaces, you can associate a physical interface with an aggregated SONET/SDH interface. To associate the interface with an aggregated SONET/SDH

link, include the **aggregate** statement at the [edit interfaces *interface-name* sonet-options] hierarchy level:

```
[edit interfaces interface-name sonet-options]
aggregate asx;
```

x is the interface instance number and can be from 0 through 15, for a total of 16 aggregated interfaces. You should not mix SONET and SDH mode on the same aggregated interface. You must also include a statement configuring *asx* at the [edit interfaces] hierarchy level. For a sample configuration, see “Example: Configuring Aggregated SONET/SDH Interfaces” on page 885.

Configuring Aggregated SONET/SDH Link Speed

On aggregated SONET/SDH interfaces, you can set the required link speed for all interfaces included in the bundle, or specify that the bundle contains interfaces with mixed interface speeds.



NOTE: For nonconcatenated interfaces on aggregated SONET/SDH interfaces, you can configure the link speed of the aggregate to match the speed of the nonconcatenated interface. For example, an OC12 PIC can have nonconcatenated interfaces with a link speed of OC3.

To set the required link speed or specify mixed interface speeds, include the **link-speed** statement at the [edit interfaces *interface-name* aggregated-sonet-options] hierarchy level:

```
[edit interfaces interface-name aggregated-sonet-options]
link-speed (speed | mixed);
```

The link speed can be one of the following values:

- **oc3**—Links are OC3c or STM1c.
- **oc12**—Links are OC12c or STM4c.
- **oc48**—Links are OC48c or STM16c.
- **oc192**—Links are OC192c or STM64c.
- **oc768**—Links are OC768c or STM256c.

Configuring Aggregated SONET/SDH Minimum Links

On aggregated SONET/SDH interfaces, you can configure the minimum number of links that must be up for the bundle as a whole to be labeled up. By default, only one link must be up for the bundle to be labeled up.

To configure the minimum number of links, include the **minimum-links** statement at the [edit interfaces *interface-name* aggregated-sonet-options] hierarchy level:

```
[edit interfaces interface-name aggregated-sonet-options]
minimum-links number;
```

On a T Series, TX Matrix router with SONET interfaces, the valid range for minimum-links *number* is from 1 through 16. When the maximum value (16) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

On all other router routers, the range of valid values for minimum-links *number* is 1 through 8 and the maximum number of links supported in an aggregate is eight. When the maximum value (8) is specified, all configured links of a bundle must be up for the bundle to be labeled up.

Configuring Filters or Sampling on Aggregated SONET/SDH Links

To set up firewall filters or sampling on aggregated SONET/SDH interfaces, you must configure the *asx* interface with these properties. The filters function in the same manner as on other interfaces.

To configure a filter, include the *filter* statement:

```
filter {
  input input-filter-name;
  output output-filter-name;
}
```

You can include this statement at the following hierarchy levels:

- [edit interfaces as x unit *logical-unit-number*]
- [edit logical-systems *logical-system-name* interfaces asx unit *logical-unit-number*]

You must also configure separate statements that define the properties of the filter. For more information, see the *JUNOS Policy Framework Configuration Guide* and “Examples: Configuring Filters or Sampling on Aggregated SONET/SDH Links” on page 884.

Examples: Configuring Filters or Sampling on Aggregated SONET/SDH Links

Configure filtering on aggregated SONET/SDH interfaces:

```
[edit interfaces]
asx {
  unit 0 {
    family inet {
      address 10.2.11.1/30;
      filter {
        input input-filter-name;
        output output-filter-name;
      }
    }
  }
}
```

Defining the Filter

```
[edit firewall]
filter input-filter-name {
```

```

        term match-any-input {
            then {
                accept;
            }
        }
    }
}
filter output-filter-name {
    term match-any-output {
        then {
            accept;
        }
    }
}
}

```

**Configuring Sampling on
an Aggregated
SONET/SDH Interface**

```

[edit interfaces]
asx {
    unit 0 {
        family inet {
            address 10.2.11.1/30;
            filter {
                input input-sampler-name;
            }
        }
    }
}

```

**Defining the Sampling
Filter and the
Forwarding Action**

```

[edit firewall]
filter input-sampler-name {
    term match-any-input {
        then {
            sample;
            accept;
        }
    }
}
[edit forwarding-options]
sampling {
    input {
        family inet {
            rate 10000;
            run-length 1;
        }
    }
}
}

```

Example: Configuring Aggregated SONET/SDH Interfaces

The following configuration is sufficient to get an aggregated SONET/SDH interface up and running:

```

[edit interfaces]
as0 {
    aggregated-sonet-options {

```

```
        minimum-links 1;
        link-speed oc3;
    }
    unit 0 {
        family inet {
            address 10.2.11.1/30;
        }
    }
}
[edit chassis]
aggregated-devices {
    sonet {
        device-count 15;
    }
}
[edit interfaces]
so-1/3/0 {
    sonet-options {
        aggregate as0;
    }
}
```

Part 13

Interface Configuration Statements

- Summary of Interface Configuration Statements on page 889

Chapter 60

Summary of Interface Configuration Statements

The following descriptions explain each of the interface configuration statements. The statements are organized alphabetically.

802.3ad

Syntax	802.3ad { aex (primary backup); lacp { port-priority; } }
Hierarchy Level	[edit interfaces <i>interface-name</i> fastether-options], [edit interfaces <i>interface-name</i> ggether-options]
Release Information	Statement introduced before JUNOS Release 7.4. primary and backup options added in JUNOS Release 8.3.
Description	Specify aggregated Ethernet logical interface number.
Options	aex—Aggregated Ethernet logical interface number. Range: 0 through 15 primary—For link protection configurations, specify the primary link for egress traffic. backup—For link protection configurations, specify the backup link for egress traffic.
Usage Guidelines	See “Configuring Ethernet Link Aggregation” on page 625 and “Configuring Aggregated Ethernet Link Protection” on page 626.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	link-protection

accept

Syntax	accept inet;
Hierarchy Level	[edit interfaces <i>interface-name</i> auto-configure stacked-vlan-ranges dynamic-profile <i>profile-name</i>], [edit interfaces <i>interface-name</i> auto-configure vlan-ranges dynamic-profile <i>profile-name</i>]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	Specify the type of VLAN Ethernet packet accepted by an interface that is associated with a VLAN dynamic profile or stacked VLAN dynamic profile.
Options	inet—IPv4 Ethernet and ARP packet type.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<ul style="list-style-type: none">■ Configuring the VLAN Ethernet Packet Type for Single-Tag VLAN Dynamic Profiles■ Configuring the VLAN Ethernet Packet Type for Stacked VLAN Dynamic Profiles

accept-source-mac

Syntax

```
accept-source-mac {
  mac-address mac-address {
    policer {
      input cos-policer-name;
      output cos-policer-name;
    }
  }
}
```

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description For Gigabit Ethernet intelligent queuing (IQ) interfaces only, accept traffic from and to the specified remote media access control (MAC) address.

The **accept-source-mac** statement is equivalent to the **source-address-filter** statement, which is valid for aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only.

On untagged Gigabit Ethernet interfaces you should not configure the **source-address-filter** statement and the **accept-source-mac** statement simultaneously. On tagged Gigabit Ethernet interfaces you should not configure the **source-address-filter** statement and the **accept-source-mac** statement with an identical MAC address specified in both filters.

The statements are explained separately.

Usage Guidelines See “Configuring MAC Address Filtering” on page 761.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics source-filtering

access-concentrator

Syntax	<code>access-concentrator <i>name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> pppoe-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> pppoe-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For J Series Services Routers with Point-to-Point Protocol over Ethernet (PPPoE) interfaces, configure the name of the access concentrator (AC).
Options	<i>name</i> —Name of the AC.
Usage Guidelines	See “Identifying the Access Concentrator” on page 791.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

access-profile

Syntax	access-profile <i>name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> ppp-options chap], [edit interfaces <i>interface-name</i> ppp-options pap], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options chap], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options pap], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options chap], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options pap]
Release Information	Statement introduced before JUNOS Release 7.4. Support for PAP added in JUNOS Release 8.3.
Description	<p>For CHAP authentication, the mapping between peer names (or “clients”) and the secrets associated with their respective links. For PAP authentication, the peer's username and password.</p> <p>For Asynchronous Transfer Mode 2 (ATM2) IQ interfaces only, you can configure a Challenge Handshake Authentication Protocol (CHAP) access profile on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:</p> <ul style="list-style-type: none"> ■ atm-ppp-llc—PPP over AAL5 logical link control (LLC) encapsulation. ■ atm-ppp-vc-mux—PPP over AAL5 multiplex encapsulation.
Options	<i>name</i> —Name of the access profile.
Usage Guidelines	See “Configuring the PPP Challenge Handshake Authentication Protocol” on page 112 and “Configuring the PPP Password Authentication Protocol” on page 114.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	default-chap-secret, <i>JUNOS System Basics Configuration Guide</i>

accounting

Syntax	<pre> accounting { destination-class-usage; source-class-usage { direction; } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Enable IP packet counters on an interface.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Enabling Source Class and Destination Class Usage” on page 214.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

accounting-profile

Syntax	accounting-profile <i>name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable collection of accounting data for the specified physical or logical interface.
Options	<i>name</i> —Name of the accounting profile.
Usage Guidelines	See “Applying an Accounting Profile to the Physical Interface” on page 130 and “Applying an Accounting Profile to the Logical Interface” on page 158.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

acfc

Syntax	acfc;
Hierarchy Level	[edit interfaces <i>interface-name</i> ppp-options compression], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options compression], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options compression]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For interfaces with PPP encapsulation, configure compression of the Data Link Layer address and control fields. The acfc option is not supported with frame-relay-ppp encapsulation. On M320, M120, and T Series routers, address and control field compression (ACFC) is not supported for any ISO family protocols. Do not include the acfc statement at the [edit interfaces <i>interface-name</i> ppp-options compression] hierarchy level when you include the family iso statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.
Usage Guidelines	See “Configuring PPP Address and Control Field Compression” on page 120.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ack-delay-time

Syntax	ack-delay-time <i>time</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series Services Routers only. On Ethernet interfaces configured for DLSw, configure the maximum time allowed for incoming Information-frames (I-frames) to remain unacknowledged.
Options	<i>time</i> —Number of milliseconds. Range: 1 through 60,000 milliseconds Default: 100 milliseconds
Usage Guidelines	See “Configuring LLC2 Options” on page 180.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

ack-max

Syntax	<code>ack-max count;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series Services Routers only. On Ethernet interfaces configured for DLSw, configure the maximum number of Information-frames (I-frames) received before acknowledgment is sent.
Options	<i>count</i> —Number of I-frames. Range: 1 through 127 I-frames Default: 3 I-frames
Usage Guidelines	See “Configuring LLC2 Options” on page 180.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

acknowledge-retries

Syntax	<code>acknowledge-retries number;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services and voice services interfaces only, configure the number of retransmission attempts to be made for consecutive hello or remove link messages following the expiration of the acknowledgment timer.
Options	<i>number</i> —Number of retransmission attempts to be made following the expiration of the acknowledgment timer. Range: 1 through 5 Default: 2
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	action-red-differential-delay, hello-timer

acknowledge-timer

Syntax	acknowledge-timer <i>milliseconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services and voice services interfaces only, configure the maximum time, in milliseconds, to wait for an add link acknowledgment, hello acknowledgment, or remove link acknowledgment message.
Options	<p>milliseconds—Time, in milliseconds, to wait for an add link acknowledgment, hello acknowledgment, or remove link acknowledgment message.</p> <p>Range: 1 through 10 milliseconds</p> <p>Default: 4 milliseconds</p>
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	address, hello-timer

action

See the following sections:

- action (OAM) on page 898
- action (Policer) on page 898

action (OAM)

Syntax

```
action {
    syslog (OAM Action);
    link-down;
    send-critical-event;
}
```

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile]

Release Information Statement introduced in JUNOS Release 8.5.

Description Define the action or actions to be taken when the OAM fault event occurs.

Usage Guidelines See “Specifying the Actions to Be Taken for Link-Fault Management Events” on page 749.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

action (Policer)

Syntax

```
action {
    loss-priority high then discard;
}
```

Hierarchy Level [edit firewall three-color-policer *policer-name*]

Release Information Statement introduced in JUNOS Release 8.2.

Description This statement discards high loss priority traffic as part of a configuration using tricolor marking on a logical interface.

Usage Guidelines See the *JUNOS Class of Service Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics logical-interface-policer

action-profile

See the following sections:

- action-profile (Applying to CFM) on page 899
- action-profile (Defining for CFM) on page 899
- action-profile (Defining for LFM) on page 900

action-profile (Applying to CFM)

Syntax	<code>action-profile <i>profile-name</i>;</code>
Hierarchy Level	<code>[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> maintenance-association <i>ma-name</i> mep <i>mep-id</i> remote-mep <i>mep-id</i>]</code>
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Identify the action profile to use.
Options	<i>profile-name</i> —Name of the action profile to use.
Usage Guidelines	See “Configuring a Remote Maintenance End Point Action Profile” on page 688.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

action-profile (Defining for CFM)

Syntax	<pre>action-profile <i>profile-name</i> { default-action { interface-down; } }</pre>
Hierarchy Level	<code>[edit protocols oam ethernet connectivity-fault-management]</code>
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Configure a name and default action for an action profile.
Options	<i>profile-name</i> —Name of the action profile. The remaining statements are explained separately.
Usage Guidelines	See “Configuring a Connectivity-Fault Management Action Profile” on page 688.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

action-profile (Defining for LFM)

Syntax action-profile *profile-name* {
 action {
 syslog;
 link-down;
 send-critical-event;
 }
 event {
 link-adjacency-loss;
 link-event-rate {
 frame-error *count*;
 frame-period *count*;
 frame-period-summary *count*;
 symbol-period *count*;
 }
 }
 protocol-down;
 }

Hierarchy Level [edit protocols oam ethernet link-fault-management]

Release Information Statement introduced in JUNOS Release 8.5.

Description Configure a name, one or more actions, and the events that trigger the action for an action profile.

Options *profile-name*—Name of the action profile.

The remaining statements are explained separately.

Usage Guidelines See “Configuring an OAM Action Profile” on page 748.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

action-red-differential-delay

Syntax	action-red-differential-delay (disable-tx remove-link);
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services and voice services interfaces only, configure the action to be taken when the differential delay exceeds the red limit.
Options	disable-tx—Disable transmission on the bundle link. remove-link—Remove bundle link from service. Default: disable-tx
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	remote, yellow-differential-delay

activation-delay

Syntax	activation-delay <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>dln</i> unit <i>logical-unit-number</i> dialer-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	(J Series Services Routers) For ISDN interfaces, configure the ISDN dialer activation delay. Used only for dialer backup and dialer watch cases.
Options	<i>seconds</i> —Interval before the backup interface is activated after the primary interface has gone down. Range: 1 through 4,294,967,295 seconds
Usage Guidelines	See “Configuring the Dialer Interface” on page 835.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

activation-priority

Syntax	<code>activation-priority <i>priority</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> dynamic-call-admission-control], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> dynamic-call-admission-control]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	(J4350 and J6350 Services Routers supporting voice over IP with the TGM550 media gateway module) For Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation, configure the dynamic call admission control (dynamic CAC) activation priority value.
Options	<p><i>priority</i>—The activation priority in which the interface is used for providing call bandwidth. The interface with the highest activation priority value is used as the primary link for providing call bandwidth. If the primary link becomes unavailable, the TGM550 switches over to the next active interface with the highest activation priority value, and so on.</p> <p>Range: 0 through 255</p> <p>Default: 50</p>
Usage Guidelines	See “Configuring Dynamic Call Admission Control” on page 166.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Services Interfaces Configuration Guide, J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

address

Syntax `address address {`
 `arp ip-address (mac | multicast-mac) mac-address <publish>;`
 `broadcast address;`
 `destination address;`
 `destination-profile name;`
 `eui-64;`
 `master-only;`
 `multipoint-destination address dlcid dlcid-identifier;`
 `multipoint-destination address {`
 `epd-threshold cells;`
 `inverse-arp;`
 `oam-liveness {`
 `up-count cells;`
 `down-count cells;`
 `}`
 `oam-period (disable | seconds);`
 `shaping {`
 `(cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained`
 `rate burst length);`
 `queue-length number;`
 `}`
 `vci vpi-identifier.vci-identifier;`
 `}`
 `primary;`
 `preferred;`
 `(vrrp-group | vrrp-inet6-group) group-number {`
 `(accept-data | no-accept-data);`
 `advertise-interval seconds;`
 `authentication-type authentication;`
 `authentication-key key;`
 `fast-interval milliseconds;`
 `(preempt | no-preempt) {`
 `hold-time seconds;`
 `}`
 `priority-number number;`
 `track {`
 `priority-cost seconds;`
 `priority-hold-time interface-name {`
 `interface priority;`
 `bandwidth-threshold bits-per-second {`
 `priority;`
 `}`
 `}`
 `route ip-address/mask routing-instance instance-name priority-cost cost;`
 `}`
 `virtual-address [addresses];`
 `}`
`}`

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family*],

[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*
family *family*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the interface address.

Options *address*—Address of the interface.

The remaining statements are explained separately.

Usage Guidelines See “Configuring the Protocol Family” on page 172.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics negotiate-address, unnumbered-address, *JUNOS System Basics Configuration Guide*

advertise-interval

See the following sections:

- advertise-interval (APS) on page 905
- advertise-interval (DLSw) on page 906



NOTE: For information about the `advertise-interval` statement at the [edit interfaces *interface-name* unit *unit-number* family inet address *address* (vrrp-group | vrrp-inet6-group) *group-number*] or [edit logical-systems *logical-system-name* interfaces *interface-name* unit *unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-number*] hierarchy level, see the *JUNOS High Availability Configuration Guide*.

advertise-interval (APS)

Syntax	<code>advertise-interval milliseconds;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Modify the Automatic Protection Switching (APS) interval at which the protect and working routers send packets to their neighbors to advertise that they are operational. A router considers its neighbor to be operational for a period, called the hold time, that is, by default, three times the advertisement interval.
Options	<i>milliseconds</i> —Interval between advertisement packets. Range: 1 through 65,534 milliseconds Default: 1000 milliseconds
Usage Guidelines	See “Configuring APS Timers” on page 868.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	hold-time

advertise-interval (DLSw)

Syntax	advertise-interval <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i>]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	For J Series Services Routers only. On Ethernet interfaces configured for DLSw Ethernet redundancy, configure the advertisement interval of DLSw neighbors on the network. All routers in the redundancy group must use the same advertisement interval.
Options	<i>seconds</i> —Interval between advertisement packets. Range: 1 through 255 seconds Default: 1 second
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 181.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Services Interfaces Configuration Guide</i>

age

Syntax	age (30m 10m 1m 30s 10s);
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management linktrace]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Time to wait (in minutes or seconds) for a response. If no response is received, the request and response entry is deleted from the linktrace database.
Default	10 minutes
Usage Guidelines	See “Configuring the Linktrace Path Age Timer” on page 689.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

aggregate

See the following sections:

- aggregate (Gigabit Ethernet CoS Policer) on page 907
- aggregate (Hierarchical Policer) on page 908
- aggregate (SONET/SDH) on page 908

aggregate (Gigabit Ethernet CoS Policer)

Syntax	aggregate { bandwidth-limit (Policer for Gigabit Ethernet Interfaces) <i>bps</i> ; burst-size-limit (Policer for Gigabit Ethernet Interfaces) <i>bytes</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> gigether-options ethernet-switch-profile ethernet-policer-profile policer <i>cos-policer-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define a policer to apply to nonpremium traffic. The statements are explained separately.
Usage Guidelines	See “Configuring Gigabit Ethernet Policers” on page 757.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	premium, ieee802.1p

aggregate (Hierarchical Policier)

Syntax aggregate {
 if-exceeding {
 bandwidth-limit *bandwidth*;
 burst-size-limit *burst*;
 }
 then {
 discard;
 }
 }

Hierarchy Level [edit firewall hierarchical-policer]

Release Information Statement introduced in JUNOS Release 9.5.

Description On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, configure an aggregate hierarchical policer.

Options Options are described separately.

Usage Guidelines See “Applying Policers” on page 194 and the *JUNOS Class of Service Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

aggregate (SONET/SDH)

Syntax aggregate asx;

Hierarchy Level [edit interfaces *interface-name* sonet-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify aggregated SONET/SDH logical interface number.

Options asx—Aggregated SONET/SDH logical interface number.
 Range: 0 through 15

Usage Guidelines See “Configuring Aggregated SONET/SDH Interfaces” on page 881.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

aggregate-ports

Syntax	aggregate-ports;
Hierarchy Level	[edit chassis fpc <i>slot-number</i> pic <i>pic-number</i>]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	For T Series routers only, specify OC768-over-OC192 mode on the 4-port OC192C PIC. Four OC192 links are aggregated into one OC768 link with one logical interface.
Usage Guidelines	See “Specifying OC768-over-OC192 Mode” on page 95.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

aggregated-ether-options

```

Syntax  aggregated-ether-options {
    ethernet-switch-profile {
        ethernet-policer-profile {
            input-priority-map {
                ieee802.1p premium [ values ];
            }
            output-priority-map {
                classifier {
                    premium {
                        forwarding-class class-name {
                            loss-priority (high | low);
                        }
                    }
                }
            }
        }
        policer cos-policer-name {
            aggregate {
                bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
                burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
            }
            premium {
                bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
                burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
            }
        }
    }
    (mac-learn-enable | no-mac-learn-enable);
}
(flow-control | no-flow-control);
lacp {
    (active | passive);
    link-protection {
        disable;
    }
    (revertive | non-revertive);
    periodic interval;
    system-priority priority;
}
link-protection;
link-speed speed;
(loopback | no-loopback);
minimum-links number;
source-address-filter {
    mac-address;
    (source-filtering | no-source-filtering);
}
}

```

Hierarchy Level [edit interfaces aex]

Release Information Statement introduced before JUNOS Release 7.4.

Description	Configure aggregated Ethernet-specific interface properties. The statements are explained separately.
Usage Guidelines	See “Configuring Ethernet Interfaces” on page 585.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

aggregated-sonet-options

Syntax	aggregated-sonet-options { link-speed <i>speed</i> ; minimum-links <i>number</i> ; }
Hierarchy Level	[edit interfaces <i>asx</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure aggregated SONET/SDH-specific interface properties. The statements are explained separately.
Usage Guidelines	See “Configuring Aggregated SONET/SDH Interfaces” on page 881.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

allow-any-vci

Syntax	allow-any-vci;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit 0], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit 0]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Dedicate entire ATM device to ATM cell relay circuit.
Usage Guidelines	See “Configuring an ATM1 Cell-Relay Circuit” on page 332.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

allow-fragmentation

Syntax	allow-fragmentation;
Hierarchy Level	[edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Enable fragmentation of generic routing encapsulation (GRE) encapsulated packets regardless of maximum transmission unit (MTU) value.
Default	By default, the GRE encapsulated packets are dropped if the packet size exceeds the MTU setting of the egress interface.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	reassemble-packets

allow-remote-loopback

Syntax	allow-remote-loopback;
Hierarchy Level	[edit protocols oam link-fault-management interface <i>interface-name</i> negotiation-options]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Enable the remote loopback on IQ2 and IQ2-E Gigabit Ethernet interfaces, and all Ethernet interfaces on the MX Series routers.
Usage Guidelines	See “Enabling Remote Loopback Support on the Local Interface” on page 751.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

annex

Syntax	annex (annex-a annex-b);
Hierarchy Level	[edit interfaces <i>interface-name</i> shdsl-options],; [edit interfaces <i>interface-name</i> sonet-options aps],; [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> shdsl-options]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series Services Routers only, configure the type of SHDSL annex. For M320 and M120 routers only, for Multiplex Section Protection (MSP) switching on SDH interfaces, set annex-b. You must also configure the working protection circuit under the [edit interfaces <i>so-fpc/pic/port</i> sonet-options aps] hierarchy level.
Default	annex-b
Options	annex-a—Use for North American SHDSL network implementations. annex-b—Use for European SHDSL network implementations.
Usage Guidelines	See “Configuring ATM-over-SHDSL Interfaces” on page 361.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

apply-action-profile

Syntax	apply-action-profile <i>profile-name</i> ;
Hierarchy Level	[edit protocols oam ethernet link-fault-management interface]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Apply the specified action profile to the interface for link-fault management.
Usage Guidelines	See “Applying an Action Profile” on page 751.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

aps

Syntax `aps {
 advertise-interval milliseconds;
 annex-b
 authentication-key key;
 force;
 hold-time milliseconds;
 lockout;
 neighbor address;
 paired-group group-name;
 preserve-interface;
 protect-circuit group-name;
 request;
 revert-time seconds;
 switching-mode (bidirectional | unidirectional);
 working-circuit group-name;
 }`

Hierarchy Level [edit interfaces *interface-name* sonet-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure Automatic Protection Switching (APS) on the router.

For DS3 channels on a channelized OC12 interface, configure APS on channel 0 only. If you configure APS on channels 1 through 11, it is ignored.

The statements are explained separately.

Usage Guidelines See “Configuring APS and MSP” on page 859.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

arp

Syntax	<code>arp ip-address (mac multicast-mac) mac-address <publish>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, configure Address Resolution Protocol (ARP) table entries, mapping IP addresses to MAC addresses.
Options	<p><i>ip-address</i>—IP address to map to the MAC address. The IP address specified must be part of the subnet defined in the enclosing address statement.</p> <p><i>mac mac-address</i>—MAC address to map to the IP address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p><i>multicast-mac</i>—Multicast MAC address to map to the IP address. Specify the multicast MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p><i>publish</i>—(Optional) Have the router reply to ARP requests for the specified IP address. If you omit this option, the router uses the entry to reach the destination but does not reply to ARP requests.</p>
Usage Guidelines	See “Configuring Static ARP Table Entries” on page 669.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

asynchronous-notification

Syntax	(asynchronous-notification no-asynchronous-notification);
Hierarchy Level	[edit interfaces <i>ge-fpc/pic/port</i> <i>gigether-options</i>]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	For all 10-Gigabit Ethernet interfaces, M120, M320, and T Series routers, configure support for notification of link down alarm generation and transfer. <ul style="list-style-type: none"> ■ asynchronous-notification—Support notification of link down alarm generation and transfer. ■ no-asynchronous-notification—Prohibit notification of link down alarm generation and transfer.
Default	Support for notification of link down alarm generation and transfer is not enabled.
Usage Guidelines	See “Configuring 10-Gigabit Ethernet Notification of Link Down Alarm” on page 783.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

atm-encapsulation

Syntax	atm-encapsulation (direct plcp);
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> <i>e3-options</i>], [edit interfaces <i>at-fpc/pic/port</i> <i>t3-options</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure encapsulation for E3 and T3 traffic over ATM interfaces.
Default	Physical Layer Convergence Protocol (PLCP) encapsulation is the default for T3 traffic and for E3 traffic using G.751 framing.
Options	direct —Use direct encapsulation. G.832 framing on E3 interfaces requires direct encapsulation. plcp —Use PLCP encapsulation.
Usage Guidelines	See “Configuring E3 and T3 Parameters on ATM Interfaces” on page 337.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	encapsulation

atm-options

Syntax

```

atm-options {
  cell-bundle-size cells;
  ilmi;
  linear-red-profiles profile-name {
    high-plp-max-threshold percent;
    low-plp-max-threshold percent;
    queue-depth cells high-plp-threshold percent low-plp-threshold percent;
  }
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  pic-type (atm1 | atm2);
  plp-to-clp;
  promiscuous-mode {
    vpi vpi-identifier;
  }
  scheduler-maps map-name {
    forwarding-class class-name {
      epd-threshold cells plp1 cells;
      linear-red-profile profile-name;
      priority (high | low);
      transmit-weight (cells number | percent number);
    }
    vc-cos-mode (alternate | strict);
  }
  vpi vpi-identifier {
    maximum-vcs maximum-vcs;
    oam-liveness {
      up-count cells;
      down-count cells;
    }
    oam-period (disable | seconds);
    shaping {
      (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
        rate burst length);
      queue-length number;
    }
  }
}

```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure ATM-specific physical interface properties.

The statements are explained separately.

Usage Guidelines See “Configuring ATM Interfaces” on page 281.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics multipoint-destination, shaping, vci

atm-scheduler-map

Syntax atm-scheduler-map (*map-name* | default);

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Associate a scheduler map with a virtual circuit on a logical interface.

Options *map-name*—Name of scheduler map that you define at the [edit interfaces *interface-name* atm-options scheduler-maps] hierarchy level.

default—The default scheduler mapping.

Usage Guidelines See “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics scheduler-maps

authentication-key

Syntax authentication-key *key*;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the Automatic Protection Switching (APS) authentication key (password).

Options *key*—Authentication password. It can be 1 through 8 characters long. Configure the same key for both the working and protect routers.

Usage Guidelines See “Configuring Basic APS Support” on page 861.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics For information about the authentication-key statement at the [edit interfaces *interface-name* unit *unit-number* family inet address *address* (vrrp-group | vrrp-inet6-group) *group-number*] or [edit logical-systems *logical-system-name* interfaces *interface-name* unit *unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-number*] hierarchy level, see the *JUNOS High Availability Configuration Guide*.

authentication-profile-name

Syntax	authentication-profile-name <i>access-profile-name</i> ;
Hierarchy Level	[edit protocols dot1x authenticator]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specify the RADIUS authentication profile to use for user authentication when establishing an IEEE 802.1x Port-Based Network Access Control (dot1x) connection.
Usage Guidelines	See “Configuring IEEE 802.1x Port-Based Network Access Control” on page 741.
Required Privilege Level	interface—To view this statement in the configuration. interface control—To add this statement to the configuration.
Related Topics	dot1x, authenticator

authenticator

Syntax authenticator {
 authentication-profile-name *access-profile-name*;
 interface *interface-id* {
 maximum-requests *integer*;
 quiet-period *seconds*;
 reauthentication (disable | interval *seconds*);
 retries *integer*;
 server-timeout *seconds*;
 supplicant (*single*);
 supplicant-timeout *seconds*;
 transmit-period *seconds*;
 }
 }

Hierarchy Level [edit protocols dot1x]

Release Information Statement introduced in JUNOS Release 9.3.

Description Specify an authentication profile for user or client authentication and configure the Ethernet interface for 802.1x protocol operation.

Options authentication-profile-name *access-profile-name*—Specifies the RADIUS authentication profile for user or client authentication.

interface *interface-ids*—Configures the Ethernet interface for 802.1x protocol operation. See interface (IEEE 802.1x) for descriptions of interface statement subordinate options.

Usage Guidelines See “Configuring IEEE 802.1x Port-Based Network Access Control” on page 741.

Required Privilege Level protocols—To view this statement in the configuration.
 protocols-control—To add this statement to the configuration.

Related Topics authentication-profile-name, dot1x, interface (IEEE 802.1x)

auto-configure

Syntax	<pre> auto-configure { vlan-ranges { dynamic-profile (VLAN) <i>profile-name</i> { accept (inet); ranges (Dynamic VLAN) (any <i>low-tag</i>) - (any <i>high-tag</i>); } } stacked-vlan-ranges { dynamic-profile (Stacked VLAN)<i>profile-name</i> { accept (inet); ranges (Dynamic Stacked VLAN) (any <i>low-tag - high-tag</i>) , (any <i>low-tag - high-tag</i>); } } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	<p>Enables the configuration of dynamic, auto-sensed VLANs.</p> <p>The remaining statements are explained separately.</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	

auto-discovery

Syntax	auto-discovery;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> maintenance-association <i>ma-name</i> mep <i>mep-id</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Enable the MEP to accept continuity check messages from all remote MEPs.
Usage Guidelines	See “Enabling Maintenance End Point Automatic Discovery” on page 686.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

auto-negotiation

See the following sections:

- auto-negotiation (Gigabit Ethernet) on page 922
- auto-negotiation (J Series uPIM) on page 923

auto-negotiation (Gigabit Ethernet)

Syntax	(auto-negotiation no-auto-negotiation) remote-fault <local-interface-online local-interface-offline>;
Hierarchy Level	[edit interfaces <i>interface-name</i> gigeether-options]
Release Information	Statement introduced in JUNOS Release 7.6.
Description	For Gigabit Ethernet interfaces only, explicitly enable autonegotiation and remote fault. In this mode, you can manually configure remote fault options. Include the no-auto-negotiation statement to disable autonegotiation. When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled.
Default	Autonegotiation is automatically enabled. No explicit action is taken after the autonegotiation is complete or if the negotiation fails.
Options	<p>no-auto-negotiation—Disables explicit autonegotiation.</p> <p>remote-fault local-interface-online local-interface-offline—(Optional) Manually configure remote fault on an interface. Specify remote fault as online or offline. Default: online</p>
Usage Guidelines	See “Configuring Gigabit Ethernet Autonegotiation” on page 767.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

auto-negotiation (J Series uPIM)

Syntax	(auto-negotiation no-auto-negotiation);
Hierarchy Level	[edit interfaces <i>ge-pim</i> /0/0 switch-options switch-port <i>port-number</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	For universal Physical Interface Modules (uPIMs) on J Series Services Routers only, explicitly enable autonegotiation. If the link speed and duplex are also configured, the interfaces use the values configured as the desired values in the negotiation. Include the no-auto-negotiation statement to disable autonegotiation. If autonegotiation is disabled, the link speed and link mode must be configured.
Default	Autonegotiation is enabled by default.
Options	auto-negotiation —Enables autonegotiation. no-auto-negotiation —Disables autonegotiation.
Usage Guidelines	See “Configuring J Series Services Router Switching Interfaces” on page 589.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

auto-reconnect

Syntax	auto-reconnect <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> pppoe-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> pppoe-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For J Series Services Routers with PPP over Ethernet interfaces, configure the amount of time to wait before reconnecting after a session has terminated.
Options	seconds —Time to wait before reconnecting after a session has terminated. Range: 0 through 4,294,967,295 seconds Default: 0 (immediately)
Usage Guidelines	See “Configuring the PPPoE Automatic Reconnect Wait Timer” on page 792.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

backup-destination

Syntax	backup-destination <i>address</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For tunnel interfaces, specify the remote address of the backup tunnel.
Options	<i>address</i> —Address of the remote side of the connection.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	destination

backup-interface

Syntax	backup-interface <i>es-fpc/pic/port</i> ;
Hierarchy Level	[edit interfaces <i>es-fpc/pic/port</i> es-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure a backup ES Physical Interface Card (PIC). If the primary ES PIC fails, the backup becomes active, inherits all the tunnels and security associations (SAs), and acts as the new next hop for IP Security (IPsec) traffic.
Options	<i>es-fpc/pic/port</i> —Name of ES interface to serve as the backup.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

backup-options

Syntax	backup-options { interface <i>interface-name</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure an interface to be used as a backup interface if the primary interface goes down. This is used to support ISDN dial backup operation. The remaining statement is explained separately.
Usage Guidelines	See “Configuring an ISDN Dialer Interface as a Backup Interface” on page 826.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

bandwidth

Syntax	bandwidth <i>rate</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure an informational-only bandwidth value for an interface. This statement is valid for all logical interface types, except multilink and aggregated interfaces.
Options	<i>rate</i> —Peak rate, in bits per second (bps) or cells per second (cps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation c; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps. Range: Not limited.
Usage Guidelines	See “Configuring the Interface Bandwidth” on page 159.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

bandwidth-limit

See the following sections:

- [bandwidth-limit \(Hierarchical Policer\) on page 926](#)
- [bandwidth-limit \(Policer for Gigabit Ethernet Interfaces\) on page 927](#)

bandwidth-limit (Hierarchical Policer)

Syntax `bandwidth-limit bandwidth;`

Hierarchy Level `[edit firewall hierarchical-policer aggregate if-exceeding]`
`[edit firewall hierarchical-policer premium if-exceeding]`

Release Information Statement introduced in JUNOS Release 9.5.

Description On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, to define a policer to apply to nonpremium traffic in a hierarchical policer, use the `bandwidth-limit` statement at the `[edit firewall hierarchical-policer aggregate if-exceeding]` or `[edit firewall hierarchical-policer premium if-exceeding]` hierarchy level.

Options

Usage Guidelines See “Applying Policers” on page 194 and the *JUNOS Class of Service Configuration Guide*.

Required Privilege Level `interface`—To view this statement in the configuration.
`interface-control`—To add this statement to the configuration.

bandwidth-limit (Policer for Gigabit Ethernet Interfaces)

Syntax	bandwidth-limit <i>bps</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> together-options ethernet-switch-profile ethernet-policer-profile policer <i>cos-policer-name</i> aggregate], [edit interfaces <i>interface-name</i> together-options ethernet-switch-profile ethernet-policer-profile policer <i>cos-policer-name</i> premium]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define a policer to apply to nonpremium traffic.
Options	<i>bps</i> —Bandwidth limit, in bits per second. Specify either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). Range: 32 Kbps through 32 gigabits per second (Gbps). For IQ2 and IQ2-E interfaces 65,536 bps through 1 Gbps. For 10-Gigabit IQ2 and IQ2-E interfaces 65,536 bps through 10 Gbps.
Usage Guidelines	See “Configuring Gigabit Ethernet Policers” on page 757.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	burst-size-limit (Policer for Gigabit Ethernet Interfaces)

bchannel-allocation

Syntax	bchannel-allocation (ascending descending);
Hierarchy Level	[edit interfaces <i>interface-name</i> isdn-options]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	(J Series Services Routers equipped with a Dual-Port Channelized T1/E1 PIM) For Integrated Services Digital Network Primary Rate Interfaces (ISDN PRI), allocate PRI dialout B-channels in ascending or descending order.
Options	(ascending descending)—Allocate the B-channels in ascending (from low to high) or descending (from high to low) order. Default: Descending order
Usage Guidelines	See “Allocating B-Channels for Dialout” on page 513.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

bearer-bandwidth-limit

Syntax	<code>bearer-bandwidth-limit <i>kilobits-per-second</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> dynamic-call-admission-control], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> dynamic-call-admission-control]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	(J4350 and J6350 Services Routers supporting voice over IP with the TGM550 media gateway module) For Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation, configure the bearer bandwidth limit (BBL). BBL is used for dynamic call admission control (dynamic CAC) to provide enhanced control over WAN bandwidth.
Options	<i>kilobits-per-second</i> —The bearer bandwidth limit to be reported to a TGM550 media gateway module, in kilobits per second (kbps). Range: 0 through 9999 kbps Default: 1 (dynamic CAC is not enabled on the interface)
Usage Guidelines	See “Configuring Dynamic Call Admission Control” on page 166.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Services Interfaces Configuration Guide, J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

bert-algorithm

Syntax	<code>bert-algorithm <i>algorithm</i>;</code>
Hierarchy Level	[edit interfaces <i>ce1-fpc/pic/port</i>], [edit interfaces <i>ct1-fpc/pic/port</i>], [edit interfaces <i>interface-name</i> ds0-options], [edit interfaces <i>interface-name</i> e1-options], [edit interfaces <i>interface-name</i> e3-options], [edit interfaces <i>interface-name</i> t1-options], [edit interfaces <i>interface-name</i> t3-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the pattern to send in the bit stream during a bit error rate test (BERT). Applies to T1, E3, T3, and multichannel DS3 interfaces, the channelized interfaces (DS3, OC12, STM1), and channelized IQ and IQE interfaces (E1, E3 and DS3).



NOTE: When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the `bert-algorithm` statement must be included at the [edit interfaces *ce1-fpc/pic/port*] or [edit interfaces *ct1-fpc/pic/port*] hierarchy level as appropriate.

Options	<p><i>algorithm</i>—Pattern to send in the bit stream. There are two categories of test patterns: pseudorandom and repetitive. Both patterns conform to CCITT/ITU O.151, O.152, O.153, and O.161 standards. The algorithm can be one of the following patterns:</p> <ul style="list-style-type: none"> ■ <i>all-ones-repeating</i>—Pattern is all ones. ■ <i>all-zeros-repeating</i>—Pattern is all zeros. ■ <i>alternating-double-ones-zeros</i>—Pattern is alternating pairs of ones and zeros. ■ <i>alternating-ones-zeros</i>—Pattern is alternating ones and zeros. ■ <i>pseudo-2e3</i>—Pattern is $2^3 - 1$. ■ <i>pseudo-2e4</i>—Pattern is $2^4 - 1$. ■ <i>pseudo-2e5</i>—Pattern is $2^5 - 1$. ■ <i>pseudo-2e6</i>—Pattern is $2^6 - 1$. ■ <i>pseudo-2e7</i>—Pattern is $2^7 - 1$. ■ <i>pseudo-2e9-o153</i>—Pattern is $2^9 - 1$, as defined in the O153 standard. ■ <i>pseudo-2e10</i>—Pattern is $2^{10} - 1$. ■ <i>pseudo-2e11-o152</i>—Pattern is $2^{11} - 1$, as defined in the O152 standard. ■ <i>pseudo-2e15-o151</i>—Pattern is $2^{15} - 1$, as defined in the O151 standard. ■ <i>pseudo-2e17</i>—Pattern is $2^{17} - 1$. ■ <i>pseudo-2e18</i>—Pattern is $2^{18} - 1$.
----------------	--

- `pseudo-2e20-o151`—Pattern is $2^{20} - 1$, as defined in the O151 standard.
- `pseudo-2e20-o153`—Pattern is $2^{20} - 1$, as defined in the O153 standard.
- `pseudo-2e21`—Pattern is $2^{21} - 1$.
- `pseudo-2e22`—Pattern is $2^{22} - 1$.
- `pseudo-2e23-o151`—Pattern is $2^{23} - 1$, as defined in the O151 standard.
- `pseudo-2e25`—Pattern is $2^{25} - 1$.
- `pseudo-2e28`—Pattern is $2^{28} - 1$.
- `pseudo-2e29`—Pattern is $2^{29} - 1$.
- `pseudo-2e31`—Pattern is $2^{31} - 1$.
- `pseudo-2e32`—Pattern is $2^{32} - 1$.
- `repeating-1-in-4`—One bit in four is set to 1; the others are set to 0.
- `repeating-1-in-8`—One bit in eight is set to 1; the others are set to 0.
- `repeating-3-in-24`—Three bits in twenty four are set to 1; the others are set to 0.

Default: `pseudo-2e3`

Usage Guidelines See “Interface Diagnostics” on page 134, “Configuring E3 BERT Properties” on page 552, “Configuring T1 BERT Properties” on page 560, “Configuring T3 BERT Properties” on page 570, and “Examples: Configuring T3 Interfaces” on page 579.

Required Privilege Level `interface`—To view this statement in the configuration.
`interface-control`—To add this statement to the configuration.

Related Topics `bert-error-rate`, `bert-period`

bert-error-rate

Syntax `bert-error-rate rate;`

Hierarchy Level [edit interfaces *ce1-fpc/pic/port*],
[edit interfaces *ct1-fpc/pic/port*],
[edit interfaces *interface-name* ds0-options],
[edit interfaces *interface-name* e1-options],
[edit interfaces *interface-name* e3-options],
[edit interfaces *interface-name* t1-options],
[edit interfaces *interface-name* t3-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the bit error rate to use in a BERT procedure. Applies to E1, E3, T1, or T3 interfaces, and to the channelized interfaces (DS3, OC3, OC12, and STM1).



NOTE: When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the `bert-error-rate` statement must be included at the [edit interfaces *ce1-fpc/pic/port*] or [edit interfaces *ct1-fpc/pic/port*] hierarchy level as appropriate.

Options *rate*—Bit error rate.
Range: 0 through 7, which corresponds to 10^{-1} (1 error per bit) to 10^{-7} (1 error per 10 million bits)
Default: 0

Usage Guidelines See “Interface Diagnostics” on page 134, “Configuring E3 BERT Properties” on page 552, “Configuring T1 BERT Properties” on page 560, “Configuring T3 BERT Properties” on page 570, and “Examples: Configuring T3 Interfaces” on page 579.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics `bert-algorithm`, `bert-period`

bert-period

Syntax bert-period *seconds*;

Hierarchy Level [edit interfaces *ce1-fpc/pic/port*],
 [edit interfaces *ct1-fpc/pic/port*],
 [edit interfaces *interface-name* ds0-options],
 [edit interfaces *interface-name* e1-options],
 [edit interfaces *interface-name* e3-options],
 [edit interfaces *interface-name* t1-options],
 [edit interfaces *interface-name* t3-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the duration of a BERT test. Applies to E1, E3, T1, and T3 interfaces, and to E1, E3, T1, and T3 partitions on the channelized interfaces (CE1, CT1, DS3, OC3, OC12, OC48, STM1, STM4, and STM16).

E1 and T1 IQ, IQE, and standard interfaces support an extended BERT period range, up to 86,400 seconds (24 hours).



NOTE: When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the `bert-period` statement must be included at the [edit interfaces *ce1-fpc/pic/port*] or [edit interfaces *ct1-fpc/pic/port*] hierarchy level as appropriate.

Options

Range:

seconds—Test duration. Range and default values vary by interface type.

- PIC-dependent—Normal BERT period: either 1 through 239 seconds or 1 through 240 seconds
- PIC-dependent—Extended BERT period: from 1 through 86,400 seconds

Default:

- Normal BERT period: 10 seconds
- Extended BERT period (on supported E1 interfaces): 10 seconds
- Extended BERT period (on supported T1 interfaces): 240 seconds

Usage Guidelines See “Interface Diagnostics” on page 134, “Configuring E1 BERT Properties” on page 544, “Configuring E3 BERT Properties” on page 552, “Configuring T1 BERT Properties” on page 560, and “Configuring T3 BERT Properties” on page 570.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics bert-algorithm, bert-error-rate

bridge-domain

Syntax	bridge-domain <i>name</i> ;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain routing-instances <i>name</i>]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	Specify the OAM Ethernet CFM maintenance domain bridge domain.
Options	<i>name</i> —Specify the name of the bridge domain.
Usage Guidelines	See “Configuring Maintenance Intermediate Points” on page 683.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	maintenance-domain

broadcast

Syntax	broadcast <i>address</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the broadcast address on the network or subnet. On a subnet you cannot specify a host address of 0, nor can you specify a broadcast address.
Default	The default broadcast address has a host portion of all ones.
Options	<i>address</i> —Broadcast address. The address must have a host portion of either all ones or all zeros. You cannot specify the addresses 0.0.0.0 or 255.255.255.255.
Usage Guidelines	See “Configuring the Interface Address” on page 174.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

buildout


See the following sections:

- buildout (E3 or T3 over ATM Interfaces) on page 934
- buildout (T1 Interfaces) on page 935

buildout (E3 or T3 over ATM Interfaces)

Syntax	buildout <i>feet</i> ;
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> e3-options], [edit interfaces <i>at-fpc/pic/port</i> t3-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For E3 and T3 traffic over ATM interfaces, set the buildout value.
Options	<i>feet</i> —The buildout value in feet. Range: 0 through 450 feet (137 meters) Default: 10 feet (3 meters)
Usage Guidelines	See “Configuring E3 and T3 Parameters on ATM Interfaces” on page 337.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

buildout (T1 Interfaces)

Syntax	buildout <i>value</i> ;
Hierarchy Level	[edit interfaces <i>ct1-fpc/pic/port</i>] [edit interfaces <i>interface-name</i> t1-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For T1 interfaces, set the buildout value.
<hr/>	
	NOTE: When configuring CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the buildout statement must be included at the hierarchy level.
<hr/>	
Default	The default buildout value is 0 through 132 feet.
Options	<p>You can set the buildout value to one of the following:</p> <ul style="list-style-type: none"> ■ 0-132—0 through 132 feet (0 through 40 meters) ■ 133-265—133 through 265 feet (40 through 81 meters) ■ 266-398—266 through 398 feet (81 through 121 meters) ■ 399-531—399 through 531 feet (121 through 162 meters) ■ 532-655—532 through 655 feet (162 through 200 meters) ■ long-0db—For J Series routers only, long buildout with 0 decibel (dB) transmit attenuation ■ long-7.5db—For J Series routers only, long buildout with 7.5 dB transmit attenuation ■ long-15db—For J Series routers only, long buildout with 15 dB transmit attenuation ■ long-22.5db—For J Series routers only, long buildout with 22.5 dB transmit attenuation
Usage Guidelines	See “Configuring the T1 Buildout” on page 561.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

bundle

Syntax	<code>bundle (ml-fpc/pic/port ls-fpc/pic/port);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Associate the multilink interface with the logical interface it is joining. You can include this statement for the <code>mlfr-end-to-end</code> and <code>mlfr-uni-nni</code> protocol families only.
Options	<code>ml-fpc/pic/port</code> —Name of the multilink interface you are linking. <code>ls-fpc/pic/port</code> —Name of the link services interface you are linking.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.

burst-size-limit

See the following sections:

- burst-size-limit (Hierarchical Policer) on page 937
- burst-size-limit (Policer for Gigabit Ethernet Interfaces) on page 937


burst-size-limit (Hierarchical Policer)

Syntax	bandwidth-limit <i>bandwidth</i> ;
Hierarchy Level	[edit firewall hierarchical-policer aggregate if-exceeding], [edit firewall hierarchical-policer premium if-exceeding]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, to define a policer to apply to nonpremium traffic in a hierarchical policer, use the burst-size-limit statement at the [edit firewall hierarchical-policer aggregate if-exceeding] or [edit firewall hierarchical-policer premium if-exceeding] hierarchy level.
Options	
Usage Guidelines	See “Applying Policers” on page 194 and the <i>JUNOS Class of Service Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

burst-size-limit (Policer for Gigabit Ethernet Interfaces)

Syntax	burst-size-limit <i>bytes</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> gigether-options ethernet-switch-profile ethernet-policer-profile policer <i>cos-policer-name</i> aggregate], [edit interfaces <i>interface-name</i> gigether-options ethernet-switch-profile ethernet-policer-profile policer <i>cos-policer-name</i> premium]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define a policer to apply to nonpremium traffic.
Options	<i>bytes</i> —Burst length. Range: 1500 through 100,000,000 bytes
Usage Guidelines	See “Configuring Gigabit Ethernet Policers” on page 757.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	bandwidth-limit (Policer for Gigabit Ethernet Interfaces)

byte-encoding

Syntax	byte-encoding (nx56 nx64);
Hierarchy Level	[edit interfaces <i>t1-fpc/pic/port</i>], [edit interfaces <i>interface-name</i> ds0-options], [edit interfaces <i>interface-name</i> t1-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the byte encoding on a DS0 or T1 interface to use 7 bits per byte or 8 bits per byte.
<hr/>	
	NOTE: When configuring T1 interfaces on the 10-port Channelized E1/T1 IQE PIC, the byte-encoding statement must be included at the [edit interfaces <i>t1-fpc/pic/port</i>] hierarchy level.
<hr/>	
Default	The default byte encoding is 8 bits per byte (nx64).
Options	nx56—Use 7 bits per byte. nx64—Use 8 bits per byte.
Usage Guidelines	See “Configuring T1 Byte Encoding” on page 561.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

bytes

Syntax	<pre>bytes { c2 value; e1-quiet value; f1 value; f2 value; s1 value; z3 value; z4 value; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set values in some SONET/SDH header bytes.
Options	<p>c2 value—Path signal label SONET/SDH overhead byte. SONET/SDH frames use the C2 byte to indicate the contents of the payload inside the frame. SONET/SDH interfaces use the C2 byte to indicate whether the payload is scrambled. Range: 0 through 255 Default: 0xCF</p> <p>e1-quiet value—Default idle byte sent on the orderwire SONET/SDH overhead bytes. The router does not support the orderwire channel, and hence sends this byte continuously. Range: 0 through 255 Default: 0x7F</p> <p>f1 value, f2 value, z3 value, z4 value—SONET/SDH overhead bytes. Range: 0 through 255 Default: 0x00</p> <p>s1 value—Synchronization message SONET overhead byte. This byte is normally controlled as a side effect of the system reference clock configuration and the state of the external clock coming from an interface if the system reference clocks have been configured to use an external reference. Range: 0 through 255 Default: 0xCC</p>
Usage Guidelines	See “Configuring SONET/SDH Header Byte Values” on page 849.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	no-concatenate in the <i>JUNOS System Basics Configuration Guide</i>

callback

Syntax `callback;`

Hierarchy Level [edit interfaces *dl n* unit *logical-unit-number* dialer-options incoming-map],
[edit logical-systems *logical-system-name* interfaces *dl n* unit *logical-unit-number*
dialer-options incoming-map]

Release Information Statement introduced in JUNOS Release 7.5.

Description On J Series Services Routers with interfaces configured for ISDN, configure the dialer to terminate the incoming call and call back the originator after the callback wait period. The default wait time is 5 seconds. To configure the wait time, include the `callback-wait-period` statement at the [edit interfaces *dl n* unit *logical-unit-number* dialer-options] hierarchy level.



NOTE: The `incoming-map` statement is mandatory for the router to accept any incoming ISDN calls.

If the `callback` statement is configured, you cannot use the `caller caller-id` statement at the [edit interfaces *dl n* unit *logical-unit-number* dialer-options] hierarchy level.

Usage Guidelines See “Configuring Dial-In and Callback” on page 832.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics *J-series Services Router Basic LAN and WAN Access Configuration Guide*

callback-wait-period

Syntax	<code>callback-wait-period <i>time</i>;</code>
Hierarchy Level	[edit interfaces <i>dl</i> <i>n</i> unit <i>logical-unit-number</i> dialer-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>dl</i> <i>n</i> unit <i>logical-unit-number</i> dialer-options]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	<p>On J Series Services Routers with interfaces configured for ISDN with callback, specify the amount of time the dialer waits before calling back the caller. The default wait time is 5 seconds. The wait time is necessary because, when a call is rejected, the switch waits for up to 4 seconds on point-to-multipoint connections to ensure no other device accepts the call before sending the DISCONNECT message to the originator of the call. However, the default time of 5 seconds may not be sufficient for different switches or may not be needed on point-to-point connections.</p> <p>To configure callback mode, include the <code>callback</code> statement at the [edit interfaces <i>dl</i><i>n</i> unit <i>logical-unit-number</i> dialer-options] hierarchy level.</p> <p>If the <code>callback</code> statement is configured, you cannot use the <code>caller <i>caller-id</i></code> statement at the [edit interfaces <i>dl</i><i>n</i> unit <i>logical-unit-number</i> dialer-options] hierarchy level.</p>
Options	<i>time</i> —Time the dialer waits before calling back the caller.
Usage Guidelines	See “Configuring Dial-In and Callback” on page 832.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

caller

Syntax	caller (<i>caller-id</i> accept-all);
Hierarchy Level	[edit interfaces <i>dlr</i> unit <i>logical-unit-number</i> dialer-options incoming-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>dlr</i> unit <i>logical-unit-number</i> dialer-options incoming-map]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	On J Series Services Routers with interfaces configured for ISDN, specify the dialer to accept a specified caller number or accept all incoming calls.
Options	<p>caller-id—Incoming caller number. You can configure multiple caller IDs on a dialer. The caller ID of the incoming call is matched against all caller IDs configured on all dialers. The dialer matching the caller ID is looked at for further processing. Only a precise match is a valid match. For example, the configured caller ID 1-222-333-4444 or 222-333-4444 will match the incoming caller ID 1-222-333-4444.</p> <p>If the incoming caller ID has fewer digits than the number configured, it is not a valid match. Duplicate caller IDs are not allowed on different dialers; however, for example, the numbers 1-408-532-1091, 408-532-1091, and 532-1091 can still be configured on different dialers.</p> <p>Only one B-channel can map to one dialer. If one dialer is already mapped, any other call mapping to the same dialer is rejected (except in the case of a multilink dialer). If no dialer caller is configured on a dialer, that dialer will not accept any calls.</p> <p>accept-all—Any incoming call in an associated interface is accepted.</p>
Usage Guidelines	See “Configuring ISDN Interfaces” on page 819.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

calling-number

Syntax	calling-number <i>number</i> ;
Hierarchy Level	[edit interfaces <i>br-pim</i> /0/ <i>port</i> isdn-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	On J Series Services Routers with ISDN interfaces, configure the calling number to include in outgoing calls.
Options	<i>number</i> —Calling number.
Usage Guidelines	See “Configuring ISDN Physical Interface Properties” on page 821.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

cbit-parity

Syntax	(cbit-parity no-cbit-parity);
Hierarchy Level	[edit interfaces <i>interface-name</i> t3-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For T3 interfaces only, enable or disable C-bit parity mode, which controls the type of framing that is present on the transmitted T3 signal. When C-bit parity mode is enabled, the C-bit positions are used for the far-end block error (FEBE), far-end alarm and control (FEAC), terminal data link, path parity, and mode indicator bits, as defined in ANSI T1.107a-1989. For ATM and ATM2 IQ2 and IQ2-E interfaces, M23 framing is used when the no-cbit-parity statement is included. For all other interfaces, M13 framing is used when the no-cbit-parity statement is included.
Default	C-bit parity mode is enabled.
Usage Guidelines	See “Configuring E3 and T3 Parameters on ATM Interfaces” on page 337 and “Disabling T3 C-Bit Parity Mode” on page 571.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

cbr

Syntax	<code>cbr rate;</code>
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> atm-options vpi <i>vpi-identifier</i> shaping], [edit interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping], [edit interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i> shaping], [edit logical-systems <i>logical-system-name</i> interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping], [edit logical-systems <i>logical-system-name</i> interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i> shaping]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM encapsulation only, define a constant bit rate bandwidth utilization in the traffic-shaping profile.
Default	Unspecified bit rate (UBR); that is, bandwidth utilization is unlimited.
Options	<p>rate—Peak rate, in bits per second (bps) or cells per second (cps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation c; values expressed in cells per second are converted to bits per second by means of the formula</p> $1 \text{ cps} = 384 \text{ bps.}$ <p>For ATM1 OC3 interfaces, the maximum available rate is 100 percent of <i>line-rate</i>, or 135,600,000 bps. For ATM1 OC12 interfaces, the maximum available rate is 50 percent of <i>line-rate</i>, or 271,263,396 bps. For ATM2 IQ interfaces, the maximum available rate is 542,526,792 bps.</p>
Usage Guidelines	See “Defining the ATM Traffic-Shaping Profile” on page 319.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	rtvbr, shaping, vbr

cell-bundle-size

Syntax	cell-bundle-size <i>cells</i> ;
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> atm-options], [edit interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces using ATM Layer 2 circuit cell-relay transport mode only, you can configure the maximum number of ATM cells per frame.
Options	<i>cells</i> —Maximum number of cells. Default: 1 cell Range: 1 through 176 cells
Usage Guidelines	See “Configuring the Layer 2 Circuit Cell-Relay Cell Maximum” on page 313.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

chap

Syntax	<pre>chap { access-profile <i>name</i>; default-chap-secret <i>name</i>; local-name <i>name</i>; passive; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> ppp-options], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Allows each side of a link to challenge its peer, using a “secret” known only to the authenticator and that peer. The secret is not sent over the link.</p> <p>By default, PPP CHAP is disabled. If CHAP is not explicitly enabled, the interface makes no CHAP challenges and denies all incoming CHAP challenges.</p> <p>For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:</p> <ul style="list-style-type: none"> ■ atm-ppp-llc—PPP over AAL5 LLC encapsulation. ■ atm-ppp-vc-mux—PPP over AAL5 multiplex encapsulation. <p>The statements are explained separately.</p>
Usage Guidelines	See “Configuring the PPP Challenge Handshake Authentication Protocol” on page 112.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS System Basics Configuration Guide</i>

chap-secret

Syntax	<code>chap-secret <i>chap-secret</i>;</code>
Hierarchy Level	[edit access profile <i>profile-name</i> client <i>client-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For interfaces with PPP encapsulation on which the PPP Challenge Handshake Authentication Protocol (CHAP) is configured, configure the shared secret, as defined in RFC 1994.
Options	<i>chap-secret</i> —The secret key associated with a peer.
Usage Guidelines	See “Configuring PPP CHAP Authentication” on page 163.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Topics	pap-password and the <i>JUNOS System Basics Configuration Guide</i>

cisco-interoperability

Syntax	<code>cisco-interoperability send-lip-remove-link-for-link-reject;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	FRF.16 interoperability settings.
Options	<i>send-lip-remove-link-for-link-reject</i> —Send Link Integrity Protocol remove link when an add-link rejection message is received.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

classifier

Syntax	<pre> classifier { per-unit-scheduler { forwarding-class <i>class-name</i> { loss-priority (high low); } } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the classifier for the output priority map to be applied to outgoing frames on this interface.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Specifying an Output Priority Map” on page 759.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	input-priority-map

clear-dont-fragment-bit

Syntax	clear-dont-fragment-bit;
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</p>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Clear the don't-fragment (DF) bit on all IP version 4 (IPv4) packets entering a generic routing encapsulation (GRE) tunnel. If the encapsulated packet's size exceeds the tunnel's maximum transmission unit (MTU), the packet is fragmented before encapsulation.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

client

Syntax	client;
Hierarchy Level	[edit interfaces pp0 unit <i>logical-unit-number</i> pppoe-options], [edit logical-systems <i>logical-system-name</i> interfaces pp0 unit <i>logical-unit-number</i> pppoe-options]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	On J Series Services Routers, configure the router to operate in the PPPoE client mode.
Usage Guidelines	See “Configuring the PPPoE Client Mode” on page 793.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

clocking

Syntax	clocking (external [interface <i>interface-name</i>] internal);
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4. interface option added in JUNOS Release 8.2.
Description	For interfaces that can use various clock sources, configure the source of the transmit clock on each interface.
Options	<p>external—The clock source is provided by the data communication equipment (DCE).</p> <p>interface <i>interface-name</i>—For interfaces operating on T1/E1 PIMs for J Series Services Routers only, configure clocking for the drop-and insert feature. When configuring this feature, both ports must use the same clock source: either the router's internal clock or an external clock on one of the interfaces. If an external clock source is required, one interface must specify clocking external and the other must specify the same clock.</p> <p>internal—Use the internal stratum 3 clock as the reference clock.</p> <p>Default: internal</p>
Usage Guidelines	See “Configuring the Clock Source” on page 128 or “Configuring the Clock Source on SONET/SDH Interfaces” on page 875 and “Clock Sources on Channelized Interfaces” on page 390, and “Configuring a Channelized T1/E1 Interface to Drop and Insert Time Slots” on page 510.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	loop-timing


clocking-mode

Syntax	clocking-mode (dce internal loop);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For EIA-530 and V.35 interfaces, configure the clock mode. You cannot configure clocking-mode dce on a DTE router using an X.21 serial line protocol (detected automatically when an X.21 cable is plugged into the serial interface).
Options	<p>dce—DCE timing (DTE mode only, not valid for X.21).</p> <p>internal—Internal baud timing.</p> <p>loop—Loop timing.</p> <p>Default: loop</p>
Usage Guidelines	See “Configuring the Serial Clocking Mode” on page 269.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration</p>
Related Topics	Configuring the DTE Clock Rate on page 270

clock-rate

Syntax	clock-rate <i>rate</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For EIA-530 and V.35 interfaces, configure the interface speed, in megahertz (MHz).
Options	<p><i>rate</i>—You can specify one of the following rates:</p> <ul style="list-style-type: none">■ 2.048 MHz■ 2.341 MHz■ 2.731 MHz■ 3.277 MHz■ 4.096 MHz■ 5.461 MHz■ 8.192 MHz■ 16.384 MHz <p>Default: 16.384mhz</p>
Usage Guidelines	See “Configuring the Serial Clocking Mode” on page 269.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

compatibility-mode

Syntax	compatibility-mode (adtran digital-link kentrox larscom verilink) <subrate <i>value</i> >;
Hierarchy Level	[edit interfaces <i>interface-name</i> e3-options], [edit interfaces <i>interface-name</i> t3-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the E3 or T3 interface to be compatible with the channel service unit (CSU) at the remote end of the line.
Default	If you omit this option, the full E3 or T3 rate is used.
Options	<p>adtran—For T3 IQ interfaces only, configure compatibility with Adtran CSUs.</p> <p>digital-link—Configure compatibility with Digital Link CSUs. If you include this option on an E3 interface, you must also disable payload scrambling.</p> <p>kentrox—Configure compatibility with Kentrox CSUs. Kentrox subrate is valid for E3 IQ and T3 IQ interfaces only.</p> <p>larscom—For T3 and T3 IQ interfaces only, configure compatibility with Larscom CSUs.</p> <p>verilink—For T3 IQ and T3 IQE interfaces only, configure compatibility with Verilink CSUs</p>
	NOTE: Verilink configuration is not functional if an IQ interface is paired with an IQE interface.
	<p>subrate <i>value</i>—Subrate of the E3 or T3 line.</p> <p>Range: For Kentrox CSUs on E3 IQ interfaces and T3 IQ interfaces the subrate value must match the value configured on the CSU. Each increment of the subrate value corresponds to a rate increment of about 0.5 Mbps.</p>
Usage Guidelines	See “Configuring the E3 CSU Compatibility Mode” on page 553 and “Configuring the T3 CSU Compatibility Mode” on page 572.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	payload-scrambler

compression

See the following sections:

- compression (PPP Properties) on page 954
- compression (Voice Services) on page 955

compression (PPP Properties)

Syntax compression {
 acfc;
 pfc;
 }

Hierarchy Level [edit interfaces *interface-name* ppp-options],
 [edit interfaces *interface-name* unit *logical-unit-number* ppp-options],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*
 ppp-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description For interfaces with PPP encapsulation, set Link Control Protocol (LCP) compression options.

The statements are explained separately.

Usage Guidelines See “Configuring PPP Address and Control Field Compression” on page 120 and
 “Configuring the PPP Protocol Field Compression” on page 121.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

compression (Voice Services)

Syntax	<pre> compression { rtp { f-max-period <i>number</i>; queues [<i>queue-numbers</i>]; port { minimum <i>port-number</i>; maximum <i>port-number</i>; } } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Configure the compression properties for voice services traffic.</p> <p>The remaining statements are described separately.</p>
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

compression-device

Syntax	compression-device <i>interface-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the compression interface for voice services traffic.
Options	<i>interface-name</i> —Logical interface used for compression.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

connections

Syntax connections {
 interface-switch *connection-name* {
 interface *interface-name.unit-number*;
 interface *interface-name.unit-number*;
 }
 }

Hierarchy Level [edit protocols]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the connection between two circuits in a circuit cross-connect (CCC) connection.

The statements are explained separately.

Usage Guidelines See “Defining the Connection for Switching Cross-Connects” on page 228.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Topics *JUNOS MPLS Applications Configuration Guide*

connectivity-fault-management

Syntax

```
connectivity-fault-management {
  action-profile profile-name {
    default-action {
      interface-down;
    }
  }
  performance-monitoring {
    hardware-assisted-timestamping;
  }
  linktrace {
    age (30m | 10m | 1m | 30s | 10s);
    path-database-size path-database-size;
  }
  maintenance-domain domain-name {
    interface interface-name;
    level number;
    name-format (character-string | none | dns | mac+2oct);
    maintenance-association ma-name {
      short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
      continuity-check {
        hold-interval (OAM) minutes;
        interval (10m | 10s | 1m | 1s| 100ms);
        loss-threshold number;
      }
      mep mep-id {
        auto-discovery;
        direction (up | down);
        interface interface-name;
        priority number;
        remote-mep mep-id {
          action-profile profile-name;
        }
      }
    }
  }
}
```

Hierarchy Level [edit protocols oam ethernet]

Release Information Statement introduced in JUNOS Release 8.4.

Description For Ethernet interfaces on M320, MX Series, and T Series routers, specify connectivity fault management for IEEE 802.1ag Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately.

Usage Guidelines See “Configuring IEEE 802.1ag OAM Connectivity-Fault Management” on page 679.

Required Privilege Level interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

container-devices

Syntax	container-devices { device-count <i>number</i> ; }
Hierarchy Level	[edit chassis]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Specify the container devices configuration. The <i>number</i> option specifies the number of sequentially numbered container interfaces, from ci0 to ci127 maximum.
Options	number—Number of container devices. Range: 1 through 128
Usage Guidelines	See “Configuring Container Interfaces” on page 863.
Required Privilege Level	chassis—To view this statement in the configuration. chassis-control—To add this statement to the configuration.

container-list

Syntax	container-list [<i>container-interface-names</i>];
Hierarchy Level	[edit interfaces container-options]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Specify a list of container interfaces; for example: ci0, ci1, and up to ci127.
Options	<i>container-interface-names</i> —Name of each container interface.
Usage Guidelines	See “Configuring Container Interfaces” on page 863.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	container-options

container-options

Syntax	<pre> container-options { container-list [<i>container-interface-names</i>]; container-type aps; member-interface-type sonet { member-interface-speed [<i>speed</i>]; } } </pre>
Hierarchy Level	[edit interfaces]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Specify the container interface options.
Options	<p>interface-name—Name of the SONET or the container interface.</p> <p>aps—Specify the member link interface type of the container as APS.</p> <p>sonet—Protocol type of the container interface.</p> <p>speed—Set interface speed to OC3, OC12, OC48, OC192, OC768, or mixed.</p>
Usage Guidelines	See “Configuring Container Interfaces” on page 863.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

container-type

Syntax	<pre> container-type aps; </pre>
Hierarchy Level	[edit interfaces container-options]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Specify the container-options interface type.
Options	aps —Configure the interface type to be Automatic Protection Switching (APS).
Usage Guidelines	See “Configuring Container Interfaces” on page 863
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

continuity-check

Syntax	<pre>continuity-check { hold-interval (OAM) <i>minutes</i>; interval (10m 10s 1m 1s 100ms 10ms); loss-threshold <i>number</i>; interface-status-tlv; port-status-tlv; }</pre>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> maintenance-association <i>ma-name</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Specify continuity check protocol options.
Options	<p>hold-interval <i>minutes</i>—Specify the continuity check hold-interval, in minutes.</p> <p>interval (<i>10m</i> <i>10s</i> <i>1m</i> <i>1s</i> <i>100ms</i> <i>10ms</i>)—Specify the continuity check interval.</p> <p>loss-threshold <i>minutes</i>—Specify the loss-threshold, in minutes.</p> <p>interface-status-tlv—Enable interface-status-tlv transmission.</p> <p>port-status-tlv—Enable port-status-tlv transmission.</p>
Usage Guidelines	See “Configuring the Continuity Check” on page 685.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

control-channel

Syntax	control-channel <i>channel-name</i> { vlan <i>vlan-id</i> ; }
Hierarchy Level	[edit protocols protection-group ethernet-ring <i>name</i> (east-interface west-interface)]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	Configure the Ethernet RPS control channel logical interface to carry the RAPS PDU. The related physical interface is the physical ring port.
Options	vlan <i>vlan-id</i> —If the control channel logical interface is a trunk port, then a dedicated vlan <i>vlan-id</i> defines the dedicated VLAN channel to carry the RAPS traffic. Only configure the <i>vlan-id</i> when the control channel logical interface is the trunk port.
Usage Guidelines	See “Configuring Ethernet Ring Protection Switching” on page 799.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

control-polarity

Syntax	control-polarity (negative positive);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For X.21 interfaces only, configure the control signal polarity.
Options	positive—Positive signal polarity. negative—Negative signal polarity. Default: positive
Usage Guidelines	See “Configuring Serial Signal Polarities” on page 274.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

control-signal

Syntax	control-signal (assert de-assert normal);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options dce-options], [edit interfaces <i>interface-name</i> serial-options dte-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For X.21 interfaces only, configure the to-DCE signal.
Options	<p>assert—The to-DCE signal must be asserted.</p> <p>de-assert—The to-DCE signal must be deasserted.</p> <p>normal—Normal request-to-send (RTS) signal handling, as defined by ITU-T Recommendation X.21.</p> <p>Default: normal</p>
Usage Guidelines	See “Configuring the Serial Signal Handling” on page 271.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

copy-tos-to-outer-ip-header

Syntax	copy-tos-to-outer-ip-header;
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	For GRE tunnel interfaces only, enable the inner IP header’s TOS bits to be copied to the outer IP packet header.
Default	If you omit this statement, the TOS bits in the outer IP header are set to 0.
Usage Guidelines	See the <i>JUNOS Class of Service Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

core-dump

Syntax	(core-dump no-core-dump);
Hierarchy Level	[edit interfaces <i>mo-fpc/pic/port</i> multiservice-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For monitoring services interfaces only, a useful tool for isolating the cause of a problem. Core dumping is enabled by default. The directory <code>/var/tmp</code> contains core files. The JUNOS Software saves the current core file (0) and the four previous core files, which are numbered 1 through 4 (from newest to oldest): <ul style="list-style-type: none"> ■ core-dump—Enable the core dumping operation. ■ no-core-dump—Disable the core dumping operation.
Usage Guidelines	See “Configuring Multiservice Physical Interface Properties” on page 138 or the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

crc-major-alarm-threshold

Syntax	crc-major-alarm-threshold (1e-3 5e-4 1e-4 5e-5 1e-5);
Hierarchy Level	[edit interfaces <i>interface-name</i> t1-options]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Major alarm error thresholds for T1 CRC errors. When the threshold is exceeded for one second, a defect condition is declared. If the defect condition continues for the monitoring period, an alarm condition is declared.
Default	10-second monitoring period for all settings except 1e-5. The 1e-5 value uses a 50-second monitoring period.
Options	rate—Error rate expressed as the number of errors per number of bits. The value 1e-3 is one error in 10^{-3} bits and 5e-4 is five errors in 10^{-4} bits. Default: 5e-5
Usage Guidelines	See “Configuring T1 CRC Error Major Alarm Thresholds” on page 562.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

crc-minor-alarm-threshold

Syntax	crc-minor-alarm-threshold (1e-3 5e-4 1e-4 5e-5 1e-5 5e-6 1e-6);
Hierarchy Level	[edit interfaces <i>interface-name</i> t1-options]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Minor alarm error thresholds for T1 CRC errors. When the threshold is exceeded for one second, a defect condition is declared. If the defect condition continues for the monitoring period, an alarm condition is declared.
Default	10-second monitoring period for values 1e-3, 5e-4, 1e-4, and 5e-5. The 1e-5 value uses a 50-second monitoring period. The 5e-6 value uses a 100-second monitoring period. The 1e-6 value uses a 500-second monitoring period.
Options	rate—Error rate expressed as the number of errors per number of bits. The value 1e-3 is one error in 10^{-3} bits and 5e-4 is five errors in 10^{-4} bits. Default: 5e-6
Usage Guidelines	See “Configuring T1 CRC Error Minor Alarm Thresholds” on page 562.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

cts

Syntax	cts (ignore normal require);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options dce-options], [edit interfaces <i>interface-name</i> serial-options dte-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For EIA-530 and V.35 interfaces only, configure the from-DCE signal, clear-to-send (CTS).
Options	ignore—The from-DCE signal is ignored. normal—Normal CTS signal handling as defined by the TIA/EIA Standard 530. require—The from-DCE signal must be asserted. Default: normal
Usage Guidelines	See “Configuring the Serial Signal Handling” on page 271.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

cts-polarity

Syntax	cts-polarity (negative positive);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure CTS signal polarity.
Options	positive—Positive signal polarity. negative—Negative signal polarity. Default: positive
Usage Guidelines	See “Configuring Serial Signal Polarities” on page 274.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

current

Syntax	current <i>margin</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> shdsl-options snr-margin], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> shdsl-options snr-margin]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series Services Routers only, configure the current target signal-to-noise ratio (SNR) margin to be used when training the SHDSL line. The current margin is the difference between desired SNR and the actual SNR. When configured, the line trains at higher than the current margin plus SNR threshold.
Options	<i>margin</i> —Desired current SNR margin. Specify either disabled or a value from 0 dB through 10 dB. Default: 0 dB
Usage Guidelines	See “Configuring ATM-over-SHDSL Interfaces” on page 361.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

data-input

Syntax	data-input (system interface <i>interface-name</i>);
Hierarchy Level	[edit interfaces ds-pim/0/port:channel]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	<p>For interfaces operating on T1/E1 PIMs for J Series Services Routers only, configure whether an interface should send and receive data from the Routing Engine or from a given interface name. On channelized T1/E1 interfaces partitioned into channels, you can insert time slots from one port directly into the other port on the same PIM, to replace time slots coming through the Routing Engine.</p> <p>To avoid slips, both ports must use the same clock source: either the router's internal clock or an external clock on one of the interfaces. If an external clock source is required, one interface must specify clocking external and the other must specify the same clock by including the clocking external interface <i>interface-name</i> statement at the [edit interfaces <i>interface-name</i>] hierarchy level.</p>
Options	<p>system—Interface sends and receives data from the Routing Engine.</p> <p>interface <i>interface-name</i>—Interface sends and receives data from a specific interface.</p> <p>Default: Data is sent and received from the Routing Engine (system).</p>
Usage Guidelines	See “Configuring a Channelized T1/E1 Interface to Drop and Insert Time Slots” on page 510 and the <i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	clocking

dcd

Syntax	dcd (ignore normal require);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options dce-options], [edit interfaces <i>interface-name</i> serial-options dte-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For EIA-530 and V.35 interfaces only, configure the from-DCE signal, data-carrier-detect (DCD).
Options	ignore—The from-DCE signal is ignored. normal—Normal DCD signal handling as defined by the TIA/EIA Standard 530. require—The from-DCE signal must be asserted. Default: normal
Usage Guidelines	See “Configuring the Serial Signal Handling” on page 271.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dcd-polarity

Syntax	dcd-polarity (negative positive);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure DCD signal polarity.
Options	positive—Positive signal polarity. negative—Negative signal polarity. Default: positive
Usage Guidelines	See “Configuring Serial Signal Polarities” on page 274.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dce

Syntax	dce;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> serial-options clocking-mode]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Frame Relay only, respond to status enquiry message keepalives. When you configure the router to be a DCE, keepalives are disabled by default.
Default	The router operates in DTE mode.
Usage Guidelines	See “Configuring the Router as a DCE with Frame Relay Encapsulation” on page 380.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dce-options

Syntax	dce-options { control-signal (assert de-assert normal); cts (ignore normal require); dcd (ignore normal require); dsr (ignore normal require); dtr <i>signal-handling-option</i> ; ignore-all; indication (ignore normal require); rts (assert de-assert normal); tm (ignore normal require); }
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced in JUNOS Release 8.3. Statement previously known as control-leads .
Description	For J Series Services Routers, configure the serial interface signal characteristics. The statements are explained separately.
Usage Guidelines	See “Configuring the Serial Signal Handling” on page 271.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

deactivation-delay

Syntax	deactivation-delay <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>dl</i> unit <i>logical-unit-number</i> dialer-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	On J Series Services Routers with ISDN interfaces, configure the ISDN deactivation delay. Used only for dialer backup and dialer watch cases.
Options	<i>seconds</i> —Interval before the backup interface is deactivated after the primary interface has comes up. Range: 1 through 4,294,967,295 seconds Default: 0 (zero)
Usage Guidelines	See “Configuring ISDN Logical Interface Properties” on page 823.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

default-action

Syntax	default-action { interface-down; }
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management action-profile <i>profile-name</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Define the action to be taken when connectivity to the remote MEP is lost.
Default	If no action is configured, no action is taken.
Options	interface-down—When a remote MEP connectivity failure is detected, bring the interface down.
Usage Guidelines	See “Configuring a CFM Action Profile Action” on page 688.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

default-chap-secret

Syntax	default-chap-secret <i>name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> ppp-options chap], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options chap], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options chap]
Release Information	Statement introduced in JUNOS Release 8.0.
Description	<p>Define the default CHAP secret to be used when no matching CHAP access profile exists.</p> <p>For ATM2 IQ interfaces only, you can configure a default CHAP secret on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:</p> <ul style="list-style-type: none"> ■ atm-ppp-llc—PPP over AAL5 LLC encapsulation. ■ atm-ppp-vc-mux—PPP over AAL5 multiplex encapsulation.
Default	If you do not include the default-chap-secret statement in the configuration, and an interface receives a CHAP challenge or response from a peer that is not in the applied access profile, the link is immediately dropped.
Usage Guidelines	See “Configuring a Default CHAP Secret” on page 113.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	access-profile

default-pap-password

Syntax	default-pap-password <i>password</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options pap], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options pap]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	For PAP authentication, the default PAP password.
Usage Guidelines	See “Configuring a Default PAP Password” on page 165.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	access-profile

demux0

Syntax

```

demux0 {
    unit logical-unit-number {
        demux-options {
            underlying-interface interface-name
        }
        family family {
            demux-destination {
                destination-prefix;
            }
            demux-source {
                source-prefix;
            }
            unnumbered-address interface-name <preferred-source-address address>;
        }
    }
}

```

Hierarchy Level [edit interfaces],
[edit logical-systems *logical-system-name* interfaces]

Release Information Statement introduced in JUNOS Release 9.0.

Description Configure the logical demultiplexing (demux) interface.

The statements are explained separately.

Usage Guidelines See “Specifying the Demux Underlying Interface” on page 253 and “Configuring IP Demux Prefixes” on page 253.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

demux-destination

See the following sections:

- demux-destination (Underlying Interface) on page 972
- demux-destination (Demux Interface) on page 973

demux-destination (Underlying Interface)

Syntax demux-destination *family*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit
logical-unit-number],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced in JUNOS Release 9.0.
 Support for aggregated Ethernet added in JUNOS Release 9.4.

Description Configure the logical demultiplexing (demux) destination family type on the IP demux underlying interface.



NOTE: The IP demux interface feature currently supports only Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.

Usage Guidelines See “Configuring an IP Demux Underlying Interface” on page 252.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

demux-destination (Demux Interface)

Syntax	demux-destination { <i>destination-prefix</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced in JUNOS Release 9.0. Support for aggregated Ethernet added in JUNOS Release 9.4.
Description	Configure one or more logical demultiplexing (demux) destination prefixes. The prefixes are matched against the destination address of packets that the underlying interface receives. When a match occurs, the packet is processed as if it was received on the demux interface.
Usage Guidelines	See “Configuring IP Demux Prefixes” on page 253.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

demux-options

Syntax	demux-options { underlying-interface <i>interface-name</i> }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Configure logical demultiplexing (demux) interface options. The statement is explained separately.
Usage Guidelines	See “Specifying the Demux Underlying Interface” on page 253.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

demux-source

See the following sections:

- demux-source (Underlying Interface) on page 974
- demux-source (Demux Interface) on page 975

demux-source (Underlying Interface)

Syntax demux-source *family*;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit
logical-unit-number],
 [edit logical-systems *logical-system-name* routing-instances *routing-instance-name*
 interfaces *interface-name* unit *logical-unit-number*],

Release Information Statement introduced in JUNOS Release 9.0.
 Support for aggregated Ethernet added in JUNOS Release 9.4.

Description Configure the logical demultiplexing (demux) source family type on the IP demux underlying interface.



NOTE: The IP demux interface feature currently supports only Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet underlying interfaces.

Usage Guidelines See “Configuring an IP Demux Underlying Interface” on page 252.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

demux-source (Demux Interface)

Syntax	demux-source { source-prefix; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> routing-instances <i>routing-instance-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced in JUNOS Release 9.0. Support for aggregated Ethernet added in JUNOS Release 9.4.
Description	Configure one or more logical demultiplexing (demux) source prefixes. The prefixes are matched against the source address of packets that the underlying interface receives. When a match occurs, the packet is processed as if it was received on the demux interface.
Usage Guidelines	See “Configuring IP Demux Prefixes” on page 253.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

description

Syntax	<code>description text;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Provide a textual description of the interface or the logical unit. Any descriptive text you include is displayed in the output of the show interfaces commands, and is also exposed in the ifAlias Management Information Base (MIB) object. It has no effect on the operation of the interface or the router.</p> <p>The textual description can also be included in the extended DHCP relay option 82 Agent Circuit ID suboption.</p>
Options	<i>text</i> —Text to describe the interface. If the text includes spaces, enclose the entire text in quotation marks.
Usage Guidelines	See “Adding an Interface Description to the Configuration” on page 96 and “Adding a Logical Unit Description to the Configuration” on page 156. For information about including the textual description in the extended DHCP relay option 82 Agent Circuit ID suboption, see Enabling and Disabling Insertion of Option 82 Information in the <i>JUNOS Subscriber Access Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destination

See the following sections:

- destination (DLSw) on page 977
- destination (IPCP) on page 978
- destination (Routing Instance) on page 978
- destination (Tunnels) on page 979



NOTE: For information about the `destination` statement at the [edit interfaces *interface-name* unit *unit-number* family inet address *address* (vrrp-group | vrrp-inet6-group) *group-number*] or [edit logical-systems *logical-system-name* interfaces *interface-name* unit *unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-number*] hierarchy level, see the *JUNOS High Availability Configuration Guide*.

destination (DLSw)

Syntax	<code>destination mac-address priority-cost priority;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track dlsw], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track dlsw]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	For J Series Services Routers only. On Ethernet interfaces configured for DLSw Ethernet redundancy, enable tracking options for a destination MAC address.
Options	<p><i>mac-address</i>—Local MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p><i>priority-cost priority</i>—Cost value that is subtracted from the priority value when remote peer connectivity is lost. Specify a value from 1 through 254.</p>
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 181.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

destination (IPCP)

Syntax	destination address destination-profile <i>profile-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For unnumbered interfaces with PPP encapsulation, specify the IP address of the remote interface.
Options	<i>address</i> —IP address of the remote interface. The <i>destination-profile</i> statement is explained separately.
Usage Guidelines	See “Configuring IPCP Options” on page 177.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	address, negotiate-address, <i>JUNOS System Basics Configuration Guide</i>

destination (Routing Instance)

Syntax	destination <i>routing-instance-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel routing-instance], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel routing-instance]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the destination routing instance that points to the routing table containing the tunnel destination address.
Default	The default Internet routing table <i>inet.0</i> .
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

destination (Tunnels)

Syntax	destination address;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For encryption, PPP-encapsulated, and tunnel interfaces, specify the remote address of the connection.
Options	<i>address</i> —Address of the remote side of the connection.
Usage Guidelines	See “Configuring the Interface Address” on page 174 or the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	point-to-point

destination-class-usage

Syntax	destination-class-usage;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable packet counters on an interface that count packets that arrive from specific customers and are destined for specific prefixes on the provider core router.
Usage Guidelines	See “Enabling Source Class and Destination Class Usage” on page 214.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	accounting, source-class-usage

destination-profile

Syntax	<code>destination-profile <i>name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i> destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet unnumbered-address <i>interface-name</i> destination <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For interfaces with PPP encapsulation, assign PPP properties to the remote end. You define the profile at the [edit access group-profile <i>name</i> ppp] hierarchy level.
Options	<i>name</i> —Profile name defined at the [edit access group-profile <i>name</i> ppp] hierarchy level.
Usage Guidelines	See “Configuring IPCP Options” on page 177.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	destination, <i>JUNOS System Basics Configuration Guide</i> .

dialer

Syntax	<code>dialer <i>filter-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Apply a dialer filter to an interface. To create the dialer filter, include the <code>dialer-filter</code> statement at the [edit firewall filter <i>family</i> <i>family</i>] hierarchy level.
Options	<i>filter-name</i> —Dialer filter name.
Usage Guidelines	See “Applying the Dial-on-Demand Dialer Filter to the Dialer Interfaces” on page 828.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

dialer-options

Syntax

```
dialer-options {
  activation-delay seconds;
  callback;
  callback-wait-period time;
  deactivation-delay seconds;
  dial-string [ dial-string-numbers ];
  idle-timeout seconds;
  incoming-map {
    caller caller-number | accept-all;
    initial-route-check seconds;
    load-interval seconds;
    load-threshold percent;
    pool pool-name;
    redial-delay time;
    watch-list {
      [ routes ];
    }
  }
}
```

Hierarchy Level [edit interfaces umd0],
[edit interfaces dln unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces dln unit *logical-unit-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify the dialer options for configuring logical interfaces for group and user sessions.

The statements are explained separately.

Usage Guidelines See “Configuring ISDN Logical Interface Properties” on page 823 and “Specifying a USB Modem Interface on J Series Routers” on page 93.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics *JUNOS Services Interfaces Configuration Guide*

dialin

Syntax	dialin (console routable);
Hierarchy Level	[edit interfaces umd0 modem-options]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	For J Series Services Routers, configure a USB modem port to act as a dial-in console or WAN backup port.
Options	<p>console—Configure the USB modem port to operate as a dial-in console for management.</p> <p>routable—Configure the USB modem port to operate as a dial-in WAN backup interface.</p> <p>Default: console</p>
Usage Guidelines	See “Specifying a USB Modem Interface on J Series Routers” on page 93.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

dial-options

Syntax	<pre>dial-options { l2tp-interface-id <i>name</i>; (shared dedicated); }</pre>
Hierarchy Level	<p>[edit interfaces <i>sp-fpc/pic/port</i> unit <i>logical-unit-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]</p>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the Layer 2 Tunneling Protocol (L2TP) options for configuring logical interfaces for group and user sessions.
Options	<p><i>l2tp-interface-id name</i>—Interface identifier that you specified at the [edit access profile <i>name</i>] hierarchy level.</p> <p>(shared dedicated)—Specify whether a logical interface can host one (dedicated) or multiple (shared) sessions at one time.</p>
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Services Interfaces Configuration Guide</i>

dial-string

Syntax	dial-string [<i>dial-string-numbers</i>];
Hierarchy Level	[edit interfaces <i>br-pim</i> /0/ <i>port</i> unit <i>logical-unit-number</i> dialer-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>br-pim</i> /0/ <i>port</i> unit <i>logical-unit-number</i> dialer-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	On J Series Services Routers with ISDN interfaces, specify one or more ISDN dial strings used to reach a destination subnetwork.
Options	<i>dial-string-numbers</i> —One or more strings of numbers to call.
Usage Guidelines	See “Configuring the Dialer Interface” on page 835.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

direction

Syntax	direction (up down);
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> maintenance-association <i>ma-name</i> mep <i>mep-id</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Configure the direction of the MEP.
Options	up—An UP MEP CCM is transmitted out of every logical interface which is part of the same bridging or vpls instance except for the interface configured on this MEP. down—Down MEP CCMs are transmitted only out the interface configured on this MEP.
Usage Guidelines	See “Configuring the Maintenance End Point Direction” on page 686 and “Configuring IEEE 802.1ag OAM Connectivity-Fault Management” on page 679.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

disable

See the following sections:

- [disable \(Interface\)](#) on page 984
- [disable \(Link Protection\)](#) on page 984

disable (Interface)

Syntax	disable;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Disable a physical or a logical interface, effectively unconfiguring it.
Usage Guidelines	See “Disabling a Physical Interface” on page 140 and “Disabling a Logical Interface” on page 167.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

disable (Link Protection)

Syntax	disable;
Hierarchy Level	[edit interfaces aeX aggregated-ether-options lacp link-protection]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Disable LACP link protection on the interface.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

disable-mlppp-inner-ppp-pfc

Syntax	disable-mlppp-inner-ppp-pfc;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	For MLPPP interfaces only, disable compression of the inner PPP header in the MLPPP payload. By default, compression is enabled.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dlci

Syntax	dlci <i>dlci-identifier</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Frame Relay and Multilink Frame Relay (MLFR) user-to-network interface (UNI) network-to-network interface (NNI) encapsulation only, and for link services, voice services and point-to-point interfaces only, configure the data-link connection identifier (DLCI) for a permanent virtual circuit (PVC) or an switched virtual circuit (SVC). To configure a DLCI for a point-to-multipoint interface, use the multipoint-destination statement to specify the DLCI.
Options	<i>dlci-identifier</i> —Data-link connection identifier. Range: 16 through 1022. For Frame Relay DLCI ranges for channelized interfaces, see “Data-Link Connection Identifiers on Channelized Interfaces” on page 388.
Usage Guidelines	See “Configuring Frame Relay DLCIs” on page 380, “Configuring a Point-to-Point Frame Relay Connection” on page 380, and the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	encapsulation, multipoint-destination, multicast-dlci

dls

Syntax	dls { destination <i>mac-address</i> priority-cost <i>priority</i> ; peer <i>ip-address</i> priority-cost <i>priority</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	For J Series Services Routers only. On Ethernet interfaces configured for DLSw, enable tracking options for a remote peer or destination MAC address. The statements are explained separately.
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 181.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

do-not-fragment

Syntax	do-not-fragment;
Hierarchy Level	[edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i> tunnel]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Disable fragmentation of GRE encapsulated packets.
Default	By default fragmentation is disabled.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	reassemble-packets

dot1x

Syntax dot1x {
 authenticator {
 authentication-profile-name *access-profile-name*;
 interface (IEEE 802.1x) *interface-id* {
 maximum-requests *integer*;
 quiet-period *seconds*;
 reauthentication (disable | interval *seconds*);
 retries *integer*;
 server-timeout *seconds*;
 supplicant (*single*);
 supplicant-timeout *seconds*;
 transmit-period *seconds*;
 }
 }
}

Hierarchy Level [edit protocols]

Release Information Statement introduced in JUNOS Release 9.3.

Description For the MX Series only, specifies settings for using 802.1x Port-Based Network Access Control.

The remaining statements are explained separately.

Usage Guidelines See “Configuring IEEE 802.1x Port-Based Network Access Control” on page 741.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics authenticator, authentication-profile-name, interface (IEEE 802.1x)

down-count

Syntax	down-count <i>cells</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i> oam-liveness], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> oam-liveness], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i> oam-liveness], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> oam-liveness], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i> oam-liveness]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM encapsulation only, configure Operation, Administration, and Maintenance (OAM) F5 loopback cell count thresholds. This feature is not supported on ATM-over-SHDSL interfaces. For ATM2 IQ PICs only, configure OAM F4 loopback cell count thresholds at the [edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>] hierarchy level.
Options	<i>cells</i> —Minimum number of consecutive OAM F4 or F5 loopback cells lost before a VC is declared down. Range: 1 through 255 Default: 5 cells
Usage Guidelines	See “Configuring the ATM OAM F5 Loopback Cell Threshold” on page 329.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

drop-timeout

Syntax	drop-timeout <i>milliseconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services, multilink, and voice services interfaces only, configure the drop timeout period, in milliseconds.
Options	<i>milliseconds</i> —Drop timeout period. Range: 0 through 2000 milliseconds Default: 0 ms (disabled)
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ds0-options

Syntax	ds0-options { bert-algorithm <i>algorithm</i> ; bert-error-rate <i>rate</i> ; bert-period <i>seconds</i> ; byte-encoding (nx56 nx64); fcs (16 32); idle-cycle-flag (flags ones); invert-data; loopback payload; start-end-flag (filler shared); }
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure DS0-specific physical interface properties. The statements are explained separately.
Usage Guidelines	See “Configuring Channelized DS3-to-DS0 Interfaces” on page 482.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dsl-options

Syntax	dsl-options { loopback local; operating-mode <i>mode</i> ; }
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For J Series Services Routers only, modify the properties of the digital subscriber line for an ATM interface. The statements are explained separately.
Usage Guidelines	See “Configuring ATM-over-ADSL Interfaces” on page 355.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

dsr

Syntax	dsr (ignore normal require);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options dce-options], [edit interfaces <i>interface-name</i> serial-options dte-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For EIA-530 and V.35 interfaces only, configure the from-DCE signal, data-set-ready (DSR).
Options	ignore—The from-DCE signal is ignored. normal—Normal DSR signal handling as defined by the TIA/EIA Standard 530. require—The from-DCE signal must be asserted. Default: normal
Usage Guidelines	See “Configuring the Serial Signal Handling” on page 271.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dsr-polarity

Syntax	dsr-polarity (negative positive);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure DSR signal polarity.
Options	positive—Positive signal polarity. negative—Negative signal polarity. Default: positive
Usage Guidelines	See “Configuring Serial Signal Polarities” on page 274.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dte-options

Syntax	dte-options { control-signal (assert de-assert normal); cts (ignore normal require); dcd (ignore normal require); dsr (ignore normal require); dtr <i>signal-handling-option</i> ; ignore-all; indication (ignore normal require); rts (assert de-assert normal); tm (ignore normal require); }
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced in JUNOS Release 8.3. Statement previously known as control-leads .
Description	For M Series and T Series routers, configure the serial interface signal characteristics. The statements are explained separately.
Usage Guidelines	See “Configuring the Serial Signal Handling” on page 271.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dtr

Syntax	<code>dtr signal-handling-option;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options dce-options], [edit interfaces <i>interface-name</i> serial-options dte-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For EIA-530 and V.35 interfaces only, configure the to-DCE signal, data-transmit-ready (DTR).
Options	<p><i>signal-handling-option</i>—Signal handling for the DTR signal. The signal handling can be one of the following:</p> <p><i>assert</i>—The to-DCE signal must be asserted.</p> <p><i>auto-synchronize</i>—Normal DTR signal with automatic synchronization. This statement has two substatements:</p> <p><i>duration milliseconds</i>—Pulse duration of resynchronization. Range: 1 through 1000 milliseconds Default: 1000 milliseconds</p> <p><i>interval seconds</i>—Offset interval for resynchronization. Range: 1 through 31 seconds Default: 15 seconds</p> <p><i>de-assert</i>—The to-DCE signal must be deasserted.</p> <p><i>normal</i>—Normal DTR signal handling as defined by the TIA/EIA Standard 530. Default: normal</p>
Usage Guidelines	See “Configuring the Serial Signal Handling” on page 271.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

dtr-circuit

Syntax	dtr-circuit (balanced unbalanced);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For EIA-530 and V.35 interfaces only, configure a DTR circuit.
Options	balanced—Balanced DTR signal. unbalanced—Unbalanced DTR signal. Default: balanced
Usage Guidelines	See “Configuring the Serial DTR Circuit” on page 274.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dtr-polarity

Syntax	dtr-polarity (negative positive);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure DTR signal polarity.
Options	positive—Positive signal polarity. negative—Negative signal polarity. Default: positive
Usage Guidelines	See “Configuring Serial Signal Polarities” on page 274.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dump-on-flow-control

Syntax	dump-on-flow-control;
Hierarchy Level	[edit interfaces <i>interface-name</i> multiservice-options]
Description	This option supports high availability functionality and can be used with various service interfaces, including <i>rsp</i> , <i>rms</i> , <i>lsq</i> , and <i>rlsq</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

dynamic-call-admission-control

Syntax	dynamic-call-admission-control { activation-priority <i>priority</i> ; bearer-bandwidth-limit <i>kilobits-per-second</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	<p>(J4350 and J6350 Services Routers supporting voice over IP with the TGM550 media gateway module) For Fast Ethernet and Gigabit Ethernet interfaces, ISDN BRI interfaces, and serial interfaces with PPP or Frame Relay encapsulation, configure dynamic call admission control (CAC). Dynamic CAC provides enhanced control over WAN bandwidth. When dynamic CAC is configured on an interface responsible for providing call bandwidth, the TGM550 informs the Media Gateway Controller (MGC) of the bandwidth limit available for voice packets on the interface and requests the MGC to block new calls when the bandwidth is exhausted.</p> <p>Dynamic CAC must be configured on each Services Router interface responsible for providing call bandwidth.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Configuring Dynamic Call Admission Control” on page 166.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Services Interfaces Configuration Guide, J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

dynamic-profile

See the following sections:

- [dynamic-profile \(Stacked VLAN\) on page 995](#)
- [dynamic-profile \(VLAN\) on page 996](#)

dynamic-profile (Stacked VLAN)

Syntax `dynamic-profile profile-name {
 accept inet;
 ranges (Dynamic Stacked VLAN) (any | low-tag - high-tag) , (any | low-tag - high-tag);
}`

Hierarchy Level [edit interfaces *interface-name* auto-configure stacked-vlan-ranges]

Release Information Statement introduced in JUNOS Release 9.5.

Description Configure a dynamic profile for use when configuring dynamic stacked VLANs.

Options *profile-name*—Name of the dynamic profile that you want to use when configuring dynamic stacked VLANs.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics

- [Dynamic Profiles Overview](#)
- [Configuring a Basic Dynamic Profile](#)
- [Associating a Stacked VLAN Dynamic Profile to an Interface](#)

dynamic-profile (VLAN)

Syntax dynamic-profile *profile-name* {
 accept inet;
 ranges (Dynamic VLAN) (any | *low-tag*)-[any | *high-tag*];
 }

Hierarchy Level [edit interfaces *interface-name* auto-configure vlan-ranges]

Release Information Statement introduced in JUNOS Release 9.5.

Description Configure a dynamic profile for use when configuring dynamic VLANs.

Options *profile-name*—Name of the dynamic profile that you want to use when configuring dynamic VLANs.

The remaining statements are explained separately.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics ■ Dynamic Profiles Overview
 ■ Configuring a Basic Dynamic Profile
 ■ Associating a Single-Tag VLAN Dynamic Profile to an Interface

e1-options

Syntax e1-options {
 bert-algorithm *algorithm*;
 bert-error-rate *rate*;
 bert-period *seconds*;
 fcs (16 | 32);
 framing (g704 | g704-no-crc4 | unframed);
 idle-cycle-flag (flags | ones);
 invert-data;
 loopback (local | remote);
 start-end-flag (filler | shared);
 timeslots *time-slot-range*;
 }

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure E1-specific physical interface properties.

The statements are explained separately.

Usage Guidelines See “Configuring Channelized E1 Interfaces” on page 501, “Configuring Channelized STM1 Interfaces” on page 465, “Configuring E1 Interfaces” on page 543, and “Configuring T1 Interfaces” on page 559.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

e3-options

Syntax e3-options {
 atm-encapsulation (direct | plcp);
 bert-algorithm *algorithm*;
 bert-error-rate *rate*;
 bert-period *seconds*;
 buildout *feet*;
 compatibility-mode (digital-link | kentrox | larscom) <subrate *value*>;
 fcs (16 | 32);
 framing (g.751 | g.832);
 idle-cycle-flag *value*;
 invert-data;
 loopback (local | remote);
 (payload-scrambler | no-payload-scrambler);
 start-end-flag *value*;
 (unframed | no-unframed);
 }

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure E3-specific physical interface properties.

For ATM1 interfaces, you can configure a subset of E3 options statements.

The statements are explained separately.

Usage Guidelines See “Configuring E3 Interfaces” on page 551 and “Configuring T3 Interfaces” on page 569.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics atm-options

east-interface

Syntax east-interface {
 control-channel *channel-name* {
 vlan *number*;
 }
 }

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in JUNOS Release 9.4.

Description For Ethernet ring protection, each ring should have two interface ports: an **east-interface** and a **west-interface**.



NOTE: Always configure the **east-interface** first, before configuring the **west-interface**.

The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.

Options ring-protection-link-end—If this port is one side of the RPL, this flag should be set.

Usage Guidelines See “Configuring Ethernet Ring Protection Switching” on page 799.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics west-interface

encapsulation

See the following sections:

- encapsulation (Container Interface) on page 1000
- encapsulation (Logical Interface) on page 1001
- encapsulation (Physical Interface) on page 1004

encapsulation (Container Interface)

Syntax	encapsulation (cisco-hdlc ppp);
Hierarchy Level	[edit interfaces cin]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Container link-layer encapsulation type.
Options	cisco-hdlc—Use Cisco-compatible High-Level Data Link Control (HDLC) framing. ppp—Use serial PPP encapsulation.
Usage Guidelines	See “Configuring Container Interfaces” on page 863.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

encapsulation (Logical Interface)

Syntax	encapsulation (atm-ccc-cell-relay atm-ccc-vc-mux atm-cisco-nlpid atm-tcc-vc-mux atm-mlppp-llc atm-nlpid atm-ppp-llc atm-ppp-vc-mux atm-snap atm-tcc-snap atm-vc-mux ether-over-atm-llc ether-vpls-over-atm-llc ether-vpls-over-fr ether-vpls-over-ppp ethernet frame-relay-ccc frame-relay-ppp frame-relay-tcc frame-relay-ether-type frame-relay-ether-type-tcc multilink-frame-relay-end-to-end multilink-ppp ppp-over-ether ppp-over-ether-over-atm-llc vlan-bridge vlan-ccc vlan-vci-ccc vlan-tcc vlan-vpls);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Logical link-layer encapsulation type.
Options	<p>atm-ccc-cell-relay—Use ATM cell-relay encapsulation.</p> <p>atm-ccc-vc-mux—Use ATM virtual circuit (VC) multiplex encapsulation on CCC circuits. When you use this encapsulation type, you can configure the ccc family only.</p> <p>atm-cisco-nlpid—Use Cisco ATM network layer protocol ID (NLPID) encapsulation. When you use this encapsulation type, you can configure the inet family only.</p> <p>atm-mlppp-llc—For ATM2 IQ interfaces only, use Multilink PPP (MLPPP) over AAL5 LLC. For this encapsulation type, your router must be equipped with a Link Services or Voice Services PIC. MLPPP over ATM encapsulation is not supported on ATM2 IQ OC48 interfaces.</p> <p>atm-nlpid—Use ATM NLPID encapsulation. When you use this encapsulation type, you can configure the inet family only.</p> <p>atm-ppp-llc—For ATM2 IQ interfaces only, use PPP over AAL5 LLC encapsulation.</p> <p>atm-ppp-vc-mux—For ATM2 IQ interfaces only, use PPP over ATM AAL5 multiplex encapsulation.</p> <p>atm-snap—Use ATM subnetwork attachment point (SNAP) encapsulation.</p> <p>atm-tcc-snap—Use ATM SNAP encapsulation on translational cross-connect (TCC) circuits.</p> <p>atm-tcc-vc-mux—Use ATM VC multiplex encapsulation on TCC circuits. When you use this encapsulation type, you can configure the tcc family only.</p> <p>atm-vc-mux—Use ATM VC multiplex encapsulation. When you use this encapsulation type, you can configure the inet family only.</p> <p>ether-over-atm-llc—For interfaces that carry IPv4 traffic, use Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces.</p>

ether-vpls-over-atm-llc—For ATM2 IQ interfaces only, use the Ethernet virtual private LAN service (VPLS) over ATM LLC encapsulation to bridge Ethernet interfaces and ATM interfaces over a VPLS routing instance (as described in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*). Packets from the ATM interfaces are converted to standard ENET2/802.3 encapsulated Ethernet frames with the frame check sequence (FCS) field removed.

ether-vpls-over-fr—For E1, T1, E3, T3, and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over Frame Relay encapsulation to support Bridged Ethernet over Frame Relay encapsulated TDM interfaces for VPLS applications, as per *RFC 2427 (1490)*.

ether-vpls-over-ppp—For E1, T1, E3, T3 and SONET interfaces only, use the Ethernet virtual private LAN service (VPLS) over point-to-point-protocol (PPP) encapsulation to support Bridged Ethernet over PPP encapsulated TDM interfaces for VPLS applications.

ethernet—Use Ethernet II encapsulation (as described in RFC 894, *A Standard for the Transmission of IP Datagrams over Ethernet Networks*).

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard Tag Protocol ID (TPID) values.

extended-vlan-vpls—Use extended virtual LAN (VLAN) VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901.



NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-ppp—Use PPP over Frame Relay circuits. When you use this encapsulation type, you can configure the **ppp** family only. J Series Routers do not support frame-relay-ppp encapsulation.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits for connecting unlike media. When you use this encapsulation type, you can configure the **tcc** family only.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay. The physical interface must be configured with flexible-frame-relay encapsulation.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect unlike media. The physical interface must be configured with flexible-frame-relay encapsulation.

multilink-frame-relay-end-to-end—Use MLFR FRF.15 encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

multilink-ppp—Use MLPPP encapsulation. This encapsulation is used only on multilink, link services, and voice services interfaces and their constituent T1 or E1 interfaces.

ppp-over-ether—For underlying Ethernet interfaces on J Series Services Routers only, use PPP over Ethernet encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

ppp-over-ether-over-atm-llc—For underlying ATM interfaces on J Series Services Routers only, use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead, configure the interface address on the PPP interface.

vlan-bridge—Use Ethernet VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q tagging, flexible-ethernet-services, and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

vlan-ccc—Use Ethernet virtual LAN (VLAN) encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

vlan-tcc—Use Ethernet VLAN encapsulation on TCC circuits. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-vpls—Use Ethernet VLAN encapsulation on VPLS circuits.

Usage Guidelines See “Configuring Interface Encapsulation on Logical Interfaces” on page 160, “Configuring Circuit and Translational Cross-Connects” on page 223, “Identifying the Access Concentrator” on page 791, “Configuring ATM Interface Encapsulation” on page 330, “Configuring VLAN Encapsulation” on page 609, “Configuring Extended VLAN Encapsulation” on page 610, “Configuring ISDN Logical Interface Properties” on page 823, “Configuring ATM-to-Ethernet Interworking” on page 229, and the *JUNOS Services Interfaces Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

encapsulation (Physical Interface)

Syntax	encapsulation (atm-ccc-cell-relay atm-pvc cisco-hdlc cisco-hdlc-ccc cisco-hdlc-tcc ethernet-bridge ethernet-ccc ethernet-over-atm ethernet-tcc ethernet-vpls extended-frame-relay-ccc extended-frame-relay-ether-type-tcc extended-frame-relay-tcc extended-vlan-bridge extended-vlan-ccc extended-vlan-tcc extended-vlan-vpls flexible-ethernet-services flexible-frame-relay frame-relay frame-relay-ccc frame-relay-ether-type frame-relay-ether-type-tcc frame-relay-port-ccc frame-relay-tcc multilink-frame-relay-uni-nni ppp ppp-ccc ppp-tcc vlan-ccc vlan-vci-ccc vlan-vpls);
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Physical link-layer encapsulation type.
Default	PPP encapsulation.
Options	<p>atm-ccc-cell-relay—Use ATM cell-relay encapsulation.</p> <p>atm-pvc—Use ATM PVC encapsulation.</p> <p>cisco-hdlc—Use Cisco-compatible High-Level Data Link Control (HDLC) framing.</p> <p>cisco-hdlc-ccc—Use Cisco-compatible HDLC framing on CCC circuits.</p> <p>cisco-hdlc-tcc—Use Cisco-compatible HDLC framing on TCC circuits for connecting unlike media.</p> <p>ethernet-bridge—Use Ethernet bridge encapsulation on Ethernet interfaces that have bridging enabled and that must accept all packets.</p> <p>ethernet-ccc—Use Ethernet CCC encapsulation on Ethernet interfaces that must accept packets carrying standard Tag Protocol ID (TPID) values.</p> <p>ethernet-over-atm—For interfaces that carry IPv4 traffic, use Ethernet over ATM encapsulation. When you use this encapsulation type, you cannot configure multipoint interfaces. As defined in RFC 1483, <i>Multiprotocol Encapsulation over ATM Adaptation Layer 5</i>, this encapsulation type allows ATM interfaces to connect to devices that support only bridged protocol data units (BPDUs). The JUNOS Software does not completely support bridging, but accepts BPDU packets as a default gateway. If you use the router as an edge device, then the router acts as a default gateway. It accepts Ethernet LLC/SNAP frames with IP or ARP in the payload, and drops the rest. For packets destined to the Ethernet LAN, a route lookup is done using the destination IP address. If the route lookup yields a full address match, the packet is encapsulated with an LLC/SNAP and MAC header, and the packet is forwarded to the ATM interface.</p> <p>ethernet-tcc—For interfaces that carry IPv4 traffic, use Ethernet TCC encapsulation on interfaces that must accept packets carrying standard TPID values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, TCC is not supported.</p>

ethernet-vpls—Use Ethernet VPLS encapsulation on Ethernet interfaces that have VPLS enabled and that must accept packets carrying standard TPID values.

extended-frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to CCC.

extended-frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect unlike media. This encapsulation type allows you to dedicate DLCIs 1 through 1022 to TCC.

extended-vlan-bridge—Use extended VLAN bridge encapsulation on Ethernet interfaces that have IEEE 802.1Q VLAN tagging and bridging enabled and that must accept packets carrying TPID 0x8100 or a user-defined TPID.

extended-vlan-ccc—Use extended VLAN encapsulation on CCC circuits with Gigabit Ethernet and 4-port Fast Ethernet interfaces that must accept packets carrying 802.1Q values. For 8-port, 12-port, and 48-port Fast Ethernet PICs, extended VLAN CCC is not supported. For 4-port Gigabit Ethernet PICs, extended VLAN CCC is not supported.

extended-vlan-tcc—For interfaces that carry IPv4 traffic, use extended VLAN encapsulation on TCC circuits with Gigabit Ethernet interfaces on which you want to use 802.1Q tagging. For 4-port Gigabit Ethernet PICs, extended VLAN TCC is not supported.

extended-vlan-vpls—Use extended VLAN VPLS encapsulation on Ethernet interfaces that have VLAN 802.1Q tagging and VPLS enabled and that must accept packets carrying TPIDs 0x8100, 0x9100, and 0x9901.



NOTE: The built-in Gigabit Ethernet PIC on an M7i router does not support extended VLAN VPLS encapsulation.

flexible-ethernet-services—For Gigabit Ethernet IQ interfaces and Gigabit Ethernet PICs with small form-factor pluggable transceivers (SFPs) (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), use flexible Ethernet services encapsulation when you want to configure multiple per-unit Ethernet encapsulations. Aggregated Ethernet bundles can use this encapsulation type. This encapsulation type allows you to configure any combination of route, TCC, CCC, Layer 2 virtual private networks (VPNs), and VPLS encapsulations on a single physical port. If you configure flexible Ethernet services encapsulation on the physical interface, VLAN IDs from 1 through 511 are no longer reserved for normal VLANs.

flexible-frame-relay—For IQ interfaces only, use flexible Frame Relay encapsulation when you want to configure multiple per-unit Frame Relay encapsulations. This encapsulation type allows you to configure any combination of TCC, CCC, and standard Frame Relay encapsulations on a single physical port. Also, each logical interface can have any DLCI value from 1 through 1022.

frame-relay—Use Frame Relay encapsulation.

frame-relay-ccc—Use Frame Relay encapsulation on CCC circuits.

frame-relay-port-ccc—Use Frame Relay port CCC encapsulation to transparently carry all the DLCIs between two customer edge (CE) routers without explicitly configuring each DLCI on the two provider edge (PE) routers with Frame Relay transport. When you use this encapsulation type, you can configure the **ccc** family only.

frame-relay-tcc—Use Frame Relay encapsulation on TCC circuits to connect unlike media.

frame-relay-ether-type—Use Frame Relay ether type encapsulation for compatibility with Cisco Frame Relay.

frame-relay-ether-type-tcc—Use Frame Relay ether type TCC for Cisco-compatible Frame Relay on TCC circuits to connect unlike media.

extended-frame-relay-ether-type-tcc—Use extended Frame Relay ether type TCC for Cisco-compatible Frame Relay for DLCIs 1 through 1022. This encapsulation is used for circuits with different media on either side of the connection.

multilink-frame-relay-uni-nni—Use MLFR UNI NNI encapsulation. This encapsulation is used only on link services and voice services interfaces functioning as FRF.16 bundles and their constituent T1 or E1 interfaces.

ppp—Use serial PPP encapsulation.

ppp-ccc—Use serial PPP encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only.

ppp-tcc—Use serial PPP encapsulation on TCC circuits for connecting unlike media. When you use this encapsulation type, you can configure the **tcc** family only.

vlan-ccc—Use Ethernet VLAN encapsulation on CCC circuits.

vlan-vci-ccc—Use ATM-to-Ethernet interworking encapsulation on CCC circuits. When you use this encapsulation type, you can configure the **ccc** family only. All logical interfaces configured on the Ethernet interface must also have the encapsulation type set to **vlan-vci-ccc**.

vlan-vpls—Use VLAN VPLS encapsulation on Ethernet interfaces with VLAN tagging and VPLS enabled. Interfaces with VLAN VPLS encapsulation accept packets carrying standard TPID values only.

Usage Guidelines See “Configuring Interface Encapsulation on Physical Interfaces” on page 106, “Defining the Encapsulation for Switching Cross-Connects” on page 225, “Configuring ATM Interface Encapsulation” on page 330, “Configuring VLAN Encapsulation” on page 609, “Configuring ATM-to-Ethernet Interworking” on page 229, and “Configuring Extended VLAN Encapsulation” on page 610.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

encoding

Syntax	encoding (nrz nrzi);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For serial interfaces, set the line encoding format.
Default	The default line encoding is non-return to zero (NRZ).
Options	nrz—Use NRZ line encoding. nrzi—Use non-return to zero inverted (NRZI) line encoding.
Usage Guidelines	See “Configuring Serial Line Encoding” on page 277.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

epd-threshold

See the following sections:

- `epd-threshold` (Logical Interface) on page 1008
- `epd-threshold` (Physical Interface) on page 1009

epd-threshold (Logical Interface)

Syntax	<code>epd-threshold cells plp1 cells;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, define the early packet discard (EPD) threshold on a VC. The EPD threshold is a limit on the number of transmit packets that can be queued. Packets that exceed the limit are discarded. For interfaces configured in trunk mode, you can also configure dual EPD thresholds depending on the packet loss priorities (PLPs). For more information, see “Configuring Two EPD Thresholds per Queue” on page 328.
Default	Approximately 1 percent of the available cell buffers. If shaping is enabled, the default EPD threshold is proportional to the shaping rate according to the following formula: $\text{default epd-threshold} = \text{number of buffers} * \text{shaping rate} / \text{line rate}$ <p>The minimum EPD threshold value is 48 cells. If the default EPD threshold formula results in an EPD threshold of less than 48 cells, the result will be ignored, and the minimum value of 48 cells will be used.</p>
Options	<i>cells</i> —Maximum number of cells. Range: For 1-port and 2-port OC12 interfaces, 48 through 425,984 cells.
Usage Guidelines	See “Configuring the ATM2 IQ EPD Threshold” on page 326.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

epd-threshold (Physical Interface)

Syntax	<code>epd-threshold <i>cells</i> plp1 <i>cells</i>;</code>
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> atm-options scheduler-maps <i>map-name</i> forwarding-class <i>class-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, define the EPD threshold on a VC. The EPD threshold is a limit on the number of transmit packets that can be queued. Packets that exceed the limit are discarded.
Default	If you do not include either the epd-threshold or the linear-red-profile statement in the forwarding class configuration, the JUNOS Software uses an EPD threshold based on the available bandwidth and other parameters.
Options	<p><i>cells</i>—Maximum number of cells.</p> <p>Range: For 1-port and 2-port OC12 interfaces, 48 through 425,984 cells. For 1-port OC48 interfaces, 48 through 425,984 cells. For 2-port OC3, DS3, and E3 interfaces, 48 through 212,992 cells. For 4-port DS3 and E3 interfaces, 48 through 106,496 cells.</p> <p>The plp1 statement is explained separately.</p>
Usage Guidelines	See “Configuring an ATM Scheduler Map” on page 341.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	linear-red-profile

es-options

Syntax	<pre>es-options { backup-interface <i>interface-name</i>; }</pre>
Hierarchy Level	[edit interfaces <i>es-fpc/pic/port</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>On ES interfaces, configure ES interface-specific interface properties.</p> <p>The backup-interface statement is explained separately.</p>
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

ethernet

```

Syntax ethernet {
    connectivity-fault-management {
        action-profile profile-name {
            default-action {
                interface-down;
            }
        }
    }
    performance-monitoring {
        hardware-assisted-timestamping;
    }
    linktrace {
        age (30m | 10m | 1m | 30s | 10s);
        path-database-size path-database-size;
    }
    maintenance-domain domain-name {
        level number;
        name-format (character-string | none | dns | mac+2octet);
        maintenance-association ma-name {
            short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
            continuity-check {
                hold-interval (OAM) minutes;
                interval (10m | 10s | 1m | 1s | 100ms);
                loss-threshold number;
            }
            mep mep-id {
                action-profile profile-name;
                auto-discovery;
                direction (up | down);
                interface interface-name;
                priority number;
                remote-mep mep-id {
                    action-profile profile-name;
                }
            }
        }
    }
}

evcs evc-id {
    evc-protocol cfm management-domain domain-id (<management-association
        association-id> | vpls (routing-instance instance-id);
    remote-uni-count count;
    multipoint-to-multipoint;
}

link-fault-management {
    action-profile profile-name {
        action {
            syslog;
            link-down;
            send-critical-event;
        }
        event {

```

```

        link-adjacency-loss;
        link-event-rate {
            frame-error count;
            frame-period count;
            frame-period-summary count;
            symbol-period count;
        }
        protocol-down;
    }
}
interface interface-name {
    apply-action-profile;
    link-discovery (active | passive);
    pdu-interval interval;
    pdu-threshold threshold-value;
    remote-loopback;
    event-thresholds {
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
    }
    negotiation-options {
        allow-remote-loopback;
        no-allow-link-events;
    }
}
}
lmi (Ethernet OAM) {
    status-counter count;
    polling-verification-timer value;
    interface name {
        uni-id uni-name;
        status-counter number;
        polling-verification-timer value;
        evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
        evc evc-name {
            default-evc;
            vlan-list vlan-id-list;
        }
    }
}
}
}

```

Hierarchy Level [edit protocols oam]

Release Information Statement introduced in JUNOS Release 8.2.

Description For Ethernet interfaces on M320, MX Series, and T Series routers, provide fault signaling and detection for 802.3ah Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately.

Usage Guidelines See “Enabling IEEE 802.3ah OAM Support” on page 746.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

ethernet-policer-profile

Syntax

```
ethernet-policer-profile {
  input-priority-map {
    ieee802.1p premium [ values ];
  }
  output-priority-map {
    classifier {
      premium {
        forwarding-class class-name {
          loss-priority (high | low);
        }
      }
    }
  }
  policer cos-policer-name {
    aggregate {
      bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
      burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
    }
    premium {
      bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
      burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
    }
  }
}
```

Hierarchy Level [edit interfaces *interface-name* gigether-options ethernet-switch-profile],
 [edit interfaces *interface-name* aggregated-ether-options ethernet-switch-profile]

Release Information Statement introduced before JUNOS Release 7.4.

Description For Gigabit Ethernet IQ, 10-Gigabit Ethernet, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), configure a class of service (CoS)-based policer. Policing applies to the inner VLAN identifiers, not to the outer tag. For Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), the **premium** policer is not supported.

The statements are explained separately.

Usage Guidelines See “Configuring Gigabit Ethernet Policers” on page 757.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

ethernet-ring

Syntax ethernet-ring *ring-name* (
 east-interface {
 control-channel *channel-name* {
 vlan *number*;
 }
 }
 guard-interval *number*;
 node-id *mac-address*;
 restore-interval *number*;
 ring-protection-link-owner;
 west-interface {
 control-channel *channel-name* {
 vlan *number*;
 }
 }
)

Hierarchy Level [edit protocols protection-group]

Description For Ethernet PICs on MX Series routers, specify the Ethernet ring in an Ethernet ring protection switching configuration.

Options The statement options are described separately.

Usage Guidelines See “Configuring Ethernet Ring Protection Switching” on page 799.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

ethernet-switch-profile

Syntax

```

ethernet-switch-profile {
  ethernet-policer-profile {
    input-priority-map {
      ieee802.1p premium [ values ];
    }
    output-priority-map {
      classifier {
        premium {
          forwarding-class class-name {
            loss-priority (high | low);
          }
        }
      }
    }
  }
  policer cos-policer-name {
    aggregate {
      bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
      burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
    }
    premium {
      bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
      burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
    }
  }
  tag-protocol-id tpid;
}
(mac-learn-enable | no-mac-learn-enable);

```

Hierarchy Level [edit interfaces *interface-name* *gether-options*],
[edit interfaces *interface-name* *aggregated-ether-options*]

Release Information Statement introduced before JUNOS Release 7.4.

Description For Gigabit Ethernet IQ, 10-Gigabit Ethernet IQ2 and IQ2-E, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, aggregated Ethernet with Gigabit Ethernet IQ interfaces, and the built-in Gigabit Ethernet port on the M7i router), configure VLAN tag and MAC address accounting and filtering properties.

The statements are explained separately.



NOTE: When you gather interfaces into a bridge domain, the `no-mac-learn-enable` statement at the [edit interfaces *interface-name* *gether-options* ethernet-switch-profile] hierarchy level is not supported. You must use the `no-mac-learning` statement at the [edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name*] hierarchy level to disable MAC learning on an interface in a bridge domain. For information on disabling MAC learning for a bridge domain, see the *MX Series Layer 2 Configuration Guide*.

Default	If the <code>ethernet-switch-profile</code> statement is not configured, Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router) behave like Gigabit Ethernet interfaces.
Usage Guidelines	See “Configuring Gigabit Ethernet Policers” on page 757, “Configuring MAC Address Filtering” on page 761, and “Configuring the Management Ethernet Interface” on page 775.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

evcs

Syntax	<pre>evcs evc-id { evc-protocol cfm management-domain <i>domain-id</i> (<management-association association-id> vpls (routing-instance <i>instance-id</i>); remote-uni-count <i>count</i>; multipoint-to-multipoint; }</pre>
Hierarchy Level	[edit protocols oam ethernet]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	On MX Series routers with <code>ge</code> , <code>xe</code> , or <code>ae</code> interfaces, configure an OAM Ethernet virtual connection.
Options	<p><code>evc-protocol cfm vpls</code>—Specify connectivity fault management (CFM) or virtual private LAN service (VPLS) as the Ethernet Virtual Connection (EVC) protocol.</p> <p><code>management-domain <i>domain-id</i></code>—(Optional) For CFM, specify the CFM management domain.</p> <p><code>management-association <i>association-id</i></code>—(Optional) For CFM, specify the CFM management association.</p> <p><code>routing-instance <i>instance-id</i></code>—(Optional) For VPLS, specify the VPLS routing instance.</p> <p><code>remote-uni-count <i>count</i></code>—(Optional) Specify the number of remote UNIs in the EVC configuration, the default is 1.</p> <p><code>multipoint-to-multipoint</code>—(Optional) Specify multiple points in the EVC configuration, the default is point-to-point if <code>remote-uni-count</code> is 1.</p>
Usage Guidelines	See “Configuring Ethernet Local Management Interface” on page 690.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	Imi (Ethernet OAM)

event

Syntax event {
 link-adjacency-loss;
 link-event-rate {
 frame-error *count*;
 frame-period *count*;
 frame-period-summary *count*;
 symbol-period *count*;
 }
 protocol-down;
 }

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile]

Release Information Statement introduced in JUNOS Release 8.5.

Description Configure threshold values for link events in an action profile.

 The remaining statements are explained separately.

Usage Guidelines See “Monitoring Protocol Status” on page 750.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

event-thresholds

Syntax event-thresholds {
 frame-error *count*;
 frame-period *count*;
 frame-period-summary *count*;
 symbol-period *count*;
 }

Hierarchy Level [edit protocols oam link-fault-management interface *interface-name*]

Release Information Statement introduced in JUNOS Release 8.4.

Description Configure threshold limit values for link events in periodic OAM PDUs.

 The remaining statements are explained separately.

Usage Guidelines See “Configuring Threshold Values for Local Fault Events on an Interface” on page 747.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

eui-64

Syntax	eui-64;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>number</i> family ipv6 address <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For interfaces that carry IP version 6 (IPv6) traffic, automatically generate the host number portion of interface addresses.
Usage Guidelines	See “Configuring the Interface Address” on page 174.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

facility-override

Syntax	facility-override <i>facility-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options syslog host <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Override default facility for system log reporting.
Options	<i>facility-name</i> —Name of facility that overrides the default assignment.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

failover-delay

Syntax	failover-delay <i>milliseconds</i> ;
Hierarchy Level	[edit protocols vrrp]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	Configure the failover delay for VRRP and VRRP for IPv6 operations.
Options	<i>milliseconds</i> —Specify the failover delay time, in milliseconds. Range: 50 through 2000
Usage Guidelines	See “Configuring VRRP and VRRP for IPv6” on page 753.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

family

Syntax family *family* {
 accounting {
 destination-class-usage;
 source-class-usage {
 direction;
 }
 }
 address *address* {
 destination *address*;
 }
 bundle *interface-name*;
 filter {
 dialer *filter-name*;
 input *filter-name*;
 output *filter-name*;
 group *filter-group-number*;
 }
 interface-mode (access | trunk);
 ipsec-sa *sa-name*;
 keep-address-and-control;
 llc2 {
 ack-delay-time *time*;
 ack-max *count*;
 idle-time *time*;
 local-window *count*;
 max-retry *count*;
 p-bit-timeout *time*;
 redundancy-group *group-number* {
 advertise-interval *seconds*;
 map {
 local-mac *mac-address* request *mac-address*;
 }
 preempt hold-time *seconds*;
 no-preempt;
 priority *priority*;
 track {
 dls {
 peer *ip-address* priority-cost *priority*;
 destination *mac-address* priority-cost *priority*;
 }
 interface *interface-name* priority-cost *priority*;
 }
 }
 t1-time *time*;
 t2-time *time*;
 trej-time *time*;
 }
 mac-validate (loose | strict);
 mtu *bytes*;
 multicast-only;
 negotiate-address;

```

no-redirects;
policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
}
primary;
protocols [inet iso mpls];
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check <fail-filter filter-name>;
sampling {
    direction;
}
service {
    input {
        service-set service-set-name <service-filter filter-name>;
        post-service-filter filter-name;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
vlan-id number;
vlan-id-list (Interface in Bridge Domain) [number number-number];
unnumbered-address interface-name destination address destination-profile
    profile-name;
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    destination-profile name;
    eui-64;
    multipoint-destination address dlci dlci-identifier;
    multipoint-destination address {
        epd-threshold cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (disable | seconds);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
                rate burst length);
            queue-length number;
        }
        vci vpi-identifier.vci-identifier;
    }
    primary;
    preferred;
    (vrrp-group | vrrp-inet6-group) group-number {

```

```

    (accept-data | no-accept-data);
    advertise-interval seconds;
    authentication-type authentication;
    authentication-key key;
    fast-interval milliseconds;
    (preempt | no-preempt) {
        hold-time seconds;
    }
    priority-number number;
    track {
        priority-cost seconds;
        priority-hold-time interface-name {
            interface priority;
            bandwidth-threshold bits-per-second {
                priority;
            }
        }
        route ip-address/mask routing-instance instance-name priority-cost cost;
    }
    virtual-address [ addresses ];
}
}
}

```

Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure protocol family information for the logical interface.

Options *family*—Protocol family:

- *any*—Protocol-independent family used for Layer 2 packet filtering
- *bridge*—(M Series and T Series routers only) Configure only when the physical interface is configured with **ethernet-bridge** type encapsulation or when the logical interface is configured with **vlan-bridge** type encapsulation
- *ccc*—Circuit cross-connect protocol suite
- *inet*—Internet Protocol version 4 suite
- *inet6*—Internet Protocol version 6 suite
- *iso*—International Organization for Standardization Open Systems Interconnection (ISO OSI) protocol suite
- *mlfr-end-to-end*—Multilink Frame Relay FRF.15
- *mlfr-uni-nni*—Multilink Frame Relay FRF.16
- *multilink-ppp*—Multilink Point-to-Point Protocol
- *mpls*—Multiprotocol Label Switching (MPLS)
- *tcc*—Translational cross-connect protocol suite
- *tnp*—Trivial Network Protocol
- *vpls*—(M Series and T Series routers only) Virtual private LAN service

The remaining statements are explained separately.

Usage Guidelines See “Configuring the Protocol Family” on page 172 and the *JUNOS Services Interfaces Configuration Guide*.

Required Privilege Level *interface*—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

fastether-options

Syntax

```
fastether-options {
  802.3ad {
    aex (primary | backup);
    lacp {
      port-priority;
    }
  }
  (flow-control | no-flow-control);
  ignore-l3-incompletes;
  ingress-rate-limit rate;
  (loopback | no-loopback);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  source-address-filter {
    mac-address;
  }
  (source-filtering | no-source-filtering);
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.


Description Configure Fast Ethernet-specific interface properties.

The statements are explained separately.

Usage Guidelines See “Configuring Ethernet Interfaces” on page 585.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

fcs

Syntax	<code>fcs (16 32);</code>
Hierarchy Level	<code>[edit interfaces e1-fpc/pic/port],</code> <code>[edit interfaces t1-fpc/pic/port],</code> <code>[edit interfaces interface-name ds0-options],</code> <code>[edit interfaces interface-name e1-options],</code> <code>[edit interfaces interface-name e3-options],</code> <code>[edit interfaces interface-name sonet-options],</code> <code>[edit interfaces interface-name t1-options],</code> <code>[edit interfaces interface-name t3-options]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For E1/E3, SONET/SDH, and T1/T3 interfaces, configure the frame checksum (FCS) on the interface. The checksum must be the same on both ends of the interface.</p> <p>On a channelized OC12 interface, the SONET/SDH fcs statement is not supported. To configure FCS on each DS3 channel, you must include the t3-options fcs statement in the configuration for each channel. For SONET/SDH, the channelized OC12 interface supports DS3 to STS-1 to OC12. For SDH, the channelized OC12 interface supports NxDS3 to NxVC3 to AU3 to STM.</p>
	<p>NOTE: When configuring E1 or T1 interfaces on 10-port Channelized E1/T1 IQE PICs, the fcs statement must be included at the <code>[edit interfaces e1-fpc/pic/port]</code> or <code>[edit interfaces t1-fpc/pic/port]</code> hierarchy level as appropriate.</p>
Options	<p>16—Use a 16-bit frame checksum on the interface.</p> <p>32—Use a 32-bit frame checksum on the interface. Using a 32-bit checksum provides more reliable packet verification, but some older equipment might not support 32-bit checksums.</p> <p>Default: 16</p>
Usage Guidelines	See “Configuring the E1 Frame Checksum” on page 545, “Configuring the E3 Frame Checksum” on page 554, “Configuring the SONET/SDH Frame Checksum” on page 851, “Configuring the T1 Frame Checksum” on page 563, and “Configuring the T3 Frame Checksum” on page 574.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

feac-loop-respond

Syntax	(feac-loop-respond no-feac-loop-respond);
Hierarchy Level	[edit interfaces <i>interface-name</i> t3-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For T3 interfaces only, configure the router so a remote CSU can place the local router into loopback.</p> <p>If you configure remote or local loopback with the T3 loopback statement, the router does not respond to FEAC requests from the CSU even if you include the feac-loop-respond statement in the configuration. For the router to respond, you must delete the loopback statement from the configuration.</p>
Default	The router does not respond to FEAC requests.
Usage Guidelines	See “Configuring the T3 FEAC Response” on page 575.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	loopback, remote-loopback-respond

filter

Syntax	<pre>filter { group filter-group-number; input filter-name; input-list [filter-names]; output filter-name; output-list [filter-names]; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Apply a filter to an interface. You can also use filters for encrypted traffic. When you configure filters, you can configure the family <i>inet</i> , <i>inet6</i> , <i>mpls</i> , or <i>vpls</i> only.
Options	<p>group <i>filter-group-number</i>—Define an interface to be part of a filter group. The default filter group number is 0. Range: 0 through 255</p> <p>input <i>filter-name</i>—Name of one filter to evaluate when packets are received on the interface.</p> <p>output <i>filter-name</i>—Name of one filter to evaluate when packets are transmitted on the interface.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Applying a Filter to an Interface” on page 203 and the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	<i>interface</i> —To view this statement in the configuration. <i>interface-control</i> —To add this statement to the configuration.
Related Topics	<i>JUNOS Policy Framework Configuration Guide</i> , <i>JUNOS System Basics Configuration Guide</i>

flexible-vlan-tagging

Syntax flexible-vlan-tagging;
 unit *logical-unit-number* {
 vlan-id *number*;
 family *family* {
 address *address*;
 }
 }
 unit *logical-unit-number* {
 vlan-tags (Stacked VLAN Tags) inner *tpid.vlan-id* outer *tpid.vlan-id*;
 family *family* {
 address *address*;
 }
 }

Hierarchy Level [edit interfaces *ge-fpc/pic/port*]

Release Information Statement introduced in JUNOS Release 8.1.

Description On M Series and T Series routers, for Fast Ethernet and Gigabit Ethernet interfaces only on Gigabit Ethernet IQ2 and IQ2-E, IQ, and IQE PICs, and for aggregated Ethernet interfaces with member links in IQ2 and IQ2-E PICs or in MX Series DPCs, simultaneously support transmission of 802.1Q VLAN single-tag and dual-tag frames on logical interfaces on the same Ethernet port.

The statements are explained separately.

Usage Guidelines See “Configuring Mixed Tagging” on page 602.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

flow-control

Syntax	(flow-control no-flow-control);
Hierarchy Level	[edit interfaces <i>interface-name</i> aggregated-ether-options], [edit interfaces <i>interface-name</i> fastether-options], [edit interfaces <i>interface-name</i> ggether-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet interfaces only, explicitly enable flow control, which regulates the flow of packets from the router to the remote side of the connection. Enabling flow control is useful when the remote device is a Gigabit Ethernet switch. Flow control is not supported on the 4-port Fast Ethernet PIC.
Default	Flow control is the default behavior.
Usage Guidelines	See “Configuring Flow Control” on page 594.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

f-max-period

Syntax	f-max-period <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression rtp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For all adaptive services interfaces and for ISDN interfaces on J Series Services Routers. Specify the maximum number of compressed packets allowed between the transmission of full headers in a compressed Real-Time Transport Protocol (RTP) traffic stream.
Options	<i>number</i> —Maximum number of packets. The value can be from 1 through 65535.
Usage Guidelines	See “Configuring Bandwidth on Demand” on page 829 and the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

force

Syntax	force (protect working);
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Perform a forced switch between the protect and working circuits. This statement is honored only if there are no higher-priority reasons to switch. It can be overridden by a signal failure on the protect circuit, thus causing a switch to the working circuit.
Options	protect—Request the circuit to become the protect circuit. working—Request the circuit to become the working circuit.
Usage Guidelines	See “Configuring Switching Between the Working and Protect Circuits” on page 866.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	request

forwarding-class

See the following sections:

- forwarding-class (ATM2 IQ Scheduler Maps) on page 1030
- forwarding-class (Gigabit Ethernet IQ Classifier) on page 1031

forwarding-class (ATM2 IQ Scheduler Maps)

Syntax forwarding-class *class-name* {
 epd-threshold *cells* plp1 *cells*;
 linear-red-profile *profile-name*;
 priority (high | low);
 transmit-weight (cells *number* | percent *number*);
 }

Hierarchy Level [edit interfaces *at-fpc/pic/port* atm-options scheduler-maps *map-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description For ATM2 IQ interfaces only, define forwarding class name and option values.

Options *class-name*—Name of forwarding class.

The statements are explained separately.

Usage Guidelines See “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics forwarding-class statement in the *JUNOS Class of Service Configuration Guide*

forwarding-class (Gigabit Ethernet IQ Classifier)

Syntax	forwarding-class <i>class-name</i> { loss-priority (high low); }
Hierarchy Level	[edit interfaces <i>interface-name</i> gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier premium]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Gigabit Ethernet IQ interfaces only, define forwarding class name and option values.
Options	<i>class-name</i> —Name of forwarding class. The statements are explained separately.
Usage Guidelines	See “Specifying an Output Priority Map” on page 759.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	input-priority-map, forwarding-class statement in the <i>JUNOS Class of Service Configuration Guide</i>

fragment-threshold

Syntax	fragment-threshold <i>bytes</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For multilink, link services, and voice services interfaces, and for J Series Services Routers ISDN interfaces, set the fragmentation threshold.
Options	<i>bytes</i> —Maximum size, in bytes, for multilink packet fragments. Any nonzero value must be a multiple of 64 bytes. Range: 128 through 16,320 bytes Default: 0 bytes (no fragmentation)
Usage Guidelines	See “Configuring ISDN Logical Interface Properties” on page 823 and the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

frame-error

Syntax	<code>frame-error count;</code>
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile event link-event-rate], [edit protocols oam link-fault-management interface <i>interface-name</i> event-thresholds]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	<p>Threshold for sending frame error events or taking the action specified in the action profile.</p> <p>A frame error is any frame error on the underlying physical layer. The threshold is reached when the number of frame errors reaches the configured value within the window. The default window is 1 second and is not configurable.</p>
Options	<p><i>count</i>—Threshold count for frame error events.</p> <p>Range: 1 through 100</p>
Usage Guidelines	See “Configuring Threshold Values for Local Fault Events on an Interface” on page 747 and “Configuring Threshold Values for Fault Events in an Action Profile” on page 750.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

frame-period

Syntax	frame-period <i>count</i> ;
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile event link-event-rate], [edit protocols oam link-fault-management interface <i>interface-name</i> event-thresholds]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	<p>Threshold for sending frame period error events or taking the action specified in the action profile.</p> <p>A frame error is any frame error on the underlying physical layer. The frame period threshold is reached when the number of frame errors reaches the configured value within the period window. The default period window is the number of minimum-size frames that can be transmitted on the underlying physical layer in 1 second. The window is not configurable.</p>
Options	<p><i>count</i>—Threshold count for frame period error events.</p> <p>Range: 1 through 100</p>
Usage Guidelines	See “Configuring Threshold Values for Local Fault Events on an Interface” on page 747 and “Configuring Threshold Values for Fault Events in an Action Profile” on page 750.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

frame-period-summary

Syntax	frame-period-summary <i>count</i> ;
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile event link-event-rate], [edit protocols oam link-fault-management interface <i>interface-name</i> event-thresholds]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	<p>Threshold for sending frame period summary error events or taking the action specified in the action profile.</p> <p>An errored frame second is any 1-second period that has at least one errored frame. This event is generated if the number of errored frame seconds is equal to or greater than the specified threshold for that period window. The default window is 60 seconds. The window is not configurable.</p>
Options	<p><i>count</i>—Threshold count for frame period summary error events.</p> <p>Range: 1 through 100</p>
Usage Guidelines	See “Configuring Threshold Values for Local Fault Events on an Interface” on page 747 and “Configuring Threshold Values for Fault Events in an Action Profile” on page 750.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

framing

See the following sections:

- framing (E1, E3, and T1 Interfaces) on page 1036
- framing (10-Gigabit Ethernet Interfaces) on page 1037
- framing (SONET and SDH Interfaces) on page 1037

framing (E1, E3, and T1 Interfaces)

Syntax framing (g704 | g704-no-crc4 | g.751 | g.832 | unframed | sf | esf);

Hierarchy Level [edit interfaces ce1-*fpc/pic/port*],
[edit interfaces ct1-*fpc/pic/port*],
[edit interfaces at-*fpc/pic/port* e3-options],
[edit interfaces e1-*fpc/pic/port* e1-options],
[edit interfaces t1-*fpc/pic/port* t1-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the framing format.



NOTE: When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the framing statement must be included at the [edit interfaces ce1-*fpc/pic/port*] or [edit interfaces ct1-*fpc/pic/port*] hierarchy level as appropriate.

Default esf for T1 interfaces; g704 for E1 interfaces. There is no default value for E3 over ATM interfaces.

Options esf—Extended superframe (ESF) mode for T1 interfaces.

g704—G.704 framing format for E1 interfaces.

g704-no-crc4—G.704 framing with no cyclic redundancy check 4 (CRC4) for E1 interfaces.

g.751—G.751 framing format for E3 over ATM interfaces.

g.832—G.832 framing format for E3 over ATM interfaces.

sf—Superframe (SF) mode for T1 interfaces.

unframed—Unframed mode for E1 interfaces.

Usage Guidelines See “Configuring E1 Framing” on page 546, “Configuring E3 and T3 Parameters on ATM Interfaces” on page 337, and “Configuring T1 Framing” on page 564.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

framing (10-Gigabit Ethernet Interfaces)

Syntax	framing (lan-phy wan-phy);
Hierarchy Level	[edit interfaces <i>xe-fpc/pic/port</i>]
Release Information	Statement introduced in JUNOS Release 8.0.
Description	For M120, M320, and T Series routers using the 10-Gigabit Ethernet IQ2 and IQ2-E PIC, configure the framing format.
Default	Operates in LAN PHY mode.
Options	<p>lan-phy—10GBASE-R interface framing format that bypasses the WIS sublayer to directly stream block-encoded Ethernet frames on a 10-Gigabit Ethernet serial interface.</p> <p>wan-phy—10GBASE-W interface framing format that allows 10-Gigabit Ethernet wide area links to use fiber-optic cables and SONET devices.</p>
Usage Guidelines	See “Configuring 10-Gigabit Ethernet Framing” on page 781 and “Configuring SONET Options for 10-Gigabit Ethernet Interfaces” on page 872.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

framing (SONET and SDH Interfaces)

Syntax	framing (sdh sonet);
Hierarchy Level	[edit interfaces <i>so-fpc/pic/port</i>]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	For the 4-port OC48 PIC with SFP installed and the 4-port OC192 PIC in T Series and M Series routers, configure SONET or SDH framing on a per-port basis. This functionality allows you to mix SONET and SDH modes on interfaces on the same PIC.
Options	<p>sdh—SDH framing.</p> <p>sonet—SONET framing.</p>
Usage Guidelines	See “Configuring SONET/SDH Framing” on page 846.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

gigether-options

```

Syntax  gigether-options {
            802.3ad {
                aex (primary | backup);
                lacp {
                    port-priority;
                }
            }
            (asynchronous-notification | no-asynchronous-notification);
            (auto-negotiation | no-auto-negotiation) remote-fault <local-interface-online |
            local-interface-offline>;
            (flow-control | no-flow-control);
            ignore-l3-incompletes;
            (loopback | no-loopback);
            mpls {
                pop-all-labels {
                    required-depth number;
                }
            }
            source-address-filter {
                mac-address;
            }
            (source-filtering | no-source-filtering);
            speed (MX Series DPC)
            ethernet-switch-profile {
                (mac-learn-enable | no-mac-learn-enable);
                tag-protocol-id [ tpids ];
                ethernet-policer-profile {
                    input-priority-map {
                        ieee802.1p premium [ values ];
                    }
                    output-priority-map {
                        classifier {
                            premium {
                                forwarding-class class-name {
                                    loss-priority (high | low);
                                }
                            }
                        }
                    }
                }
            }
            policer cos-policer-name {
                aggregate {
                    bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
                    burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
                }
                premium {
                    bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
                    burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
                }
            }
        }
    
```

}

Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure Gigabit Ethernet-specific interface properties. The statements are explained separately.
Usage Guidelines	See “Configuring Ethernet Interfaces” on page 585.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

gratuitous-arp-reply

Syntax	(gratuitous-arp-reply no-gratuitous-arp-reply);
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Ethernet interfaces, enable updating of the ARP cache for replies received in response to gratuitous ARP requests.
Default	Updating of the ARP cache is disabled on all Ethernet interfaces.
Options	gratuitous-arp-reply—Update the ARP cache. no-gratuitous-arp-reply—Do not update the ARP cache.
Usage Guidelines	See “Configuring Gratuitous ARP” on page 596.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	no-gratuitous-arp-request

guard-interval

Syntax	guard-interval <i>number</i> ;
Hierarchy Level	[edit protocols protection-group ethernet-ring <i>ring-name</i>]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	Specify the guard timer interval in 10 milliseconds (ms) intervals.
Options	Range: 10 through 2000 ms Default: 500 ms
Usage Guidelines	See “Configuring Ethernet Ring Protection Switching” on page 799.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

hardware-assisted-timestamping

Syntax	hardware-assisted-timestamping;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management performance-monitoring]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	<p>For Ethernet interfaces on Enhanced and Enhanced Queuing Dense Port Concentrators (DPCs) in MX Series routers only, enable hardware-assisted timestamping support for Ethernet frame delay measurement.</p> <p>By default, the ETH-DM feature calculates frame delays using software-based timestamping of the ETH-DM PDU frames sent and received by the MEPs in the session. As an option that can increase the accuracy of ETH-DM calculations when the DPC is loaded with heavy traffic in the receive direction, you can enable hardware-assisted timestamping of session frames in the receive direction.</p>
Usage Guidelines	See “Ethernet Frame Delay Measurements Overview” on page 711, “Guidelines for Configuring Routers to Support an ETH-DM Session” on page 717, and “Enabling the Hardware-Assisted Timestamping Option” on page 726.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

hello-timer

Syntax	hello-timer <i>milliseconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services and voice services interfaces only, configure the rate at which hello messages are sent. A hello message is transmitted after a period defined in milliseconds has elapsed.
Options	<i>milliseconds</i> —The rate at which hello messages are sent. Range: 1 through 180 milliseconds Default: 10 milliseconds
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	address, acknowledge-timer

high-plp-max-threshold

Syntax	high-plp-max-threshold <i>percent</i> ;
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> atm-options linear-red-profiles <i>profile-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, define the drop profile fill-level for the high PLP CoS VC. When the fill level exceeds the defined percentage, all packets are dropped.
Options	<i>percent</i> —Fill-level percentage when linear random early discard (RED) is applied to cells with PLP.
Usage Guidelines	See “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	low-plp-max-threshold, low-plp-threshold, queue-depth

high-plp-threshold

Syntax	<code>high-plp-threshold percent;</code>
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> atm-options linear-red-profiles <i>profile-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, define CoS VC drop profile fill-level percentage when linear RED is applied to cells with high PLP. When the fill level exceeds the defined percentage, packets with high PLP are randomly dropped by RED. This statement is mandatory.
Options	<i>percent</i> —Fill-level percentage when linear RED is applied to cells with PLP.
Usage Guidelines	See “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	high-plp-max-threshold, low-plp-max-threshold, low-plp-threshold, queue-depth

hierarchical-policer

Syntax hierarchical-policer *name* {
 aggregate {
 if-exceeding {
 bandwidth-limit *bandwidth*;
 burst-size-limit *burst*;
 }
 then {
 discard;
 }
 }
 premium {
 if-exceeding {
 bandwidth-limit *bandwidth*;
 burst-size-limit *burst*;
 }
 then {
 discard;
 }
 }
}

Hierarchy Level [edit firewall]

Release Information Statement introduced in JUNOS Release 9.5.

Description On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, to specify a hierarchical policer, use the **hierarchical-policer** statement at the [edit firewall] hierarchy level.

Options Options are described separately.

Usage Guidelines See “Applying Policers” on page 194 and the *JUNOS Class of Service Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

hold-interval

See the following sections:

- hold-interval (OAM) on page 1044
- hold-interval (Protection Group) on page 1044

hold-interval (OAM)

Syntax	hold-interval <i>minutes</i> ;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> maintenance-association <i>ma-name</i> continuity-check]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	The time to wait before flushing the maintenance end point (MEP) database, if no updates occur.
Options	<i>minutes</i> —Time to wait, in minutes.
Usage Guidelines	See “Configuring the Continuity Check Hold Interval” on page 685.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

hold-interval (Protection Group)

Syntax	hold-interval <i>number</i> ;
Hierarchy Level	[edit protocols protection-group ethernet-ring <i>name</i>]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	Specify the hold-off timer interval <i>for all rings</i> in 100 millisecond (ms) increments.
Options	Range: 0 through 10000 ms Default: 100 ms
Usage Guidelines	See “Configuring Ethernet Ring Protection Switching” on page 799.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

hold-time

See the following sections:

- hold-time (APS) on page 1045
- hold-time (DLSw) on page 1046
- hold-time (Physical Interface) on page 1047
- hold-time (SONET/SDH Defect Triggers) on page 1048



NOTE: For information about the hold-time statement at the [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-number* preempt] and [edit logical-systems *logical-system-name* interface *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-number* preempt], see the *JUNOS High Availability Configuration Guide*.

hold-time (APS)

Syntax	hold-time <i>milliseconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Hold-time value to use to determine whether a neighbor APS router is operational.
Options	<i>milliseconds</i> —Hold-time value. Range: 1 through 65,534 milliseconds Default: 3000 milliseconds (3 times the advertisement interval)
Usage Guidelines	See “Configuring APS Timers” on page 868.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	advertise-interval

hold-time (DLSw)

Syntax	hold-time <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> preempt], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> preempt]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Hold time before a higher-priority backup router preempts the master router.
Default	DLSw preemption is not timed.
Options	<i>seconds</i> —Hold-time period. Range: 0 through 3600 Default: 0 seconds (DLSw preemption is not timed.)
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 181.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

hold-time (Physical Interface)

Syntax hold-time up *milliseconds* down *milliseconds*;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Hold-time value to use to damp interface transitions. When an interface goes from up to down, it is not advertised to the rest of the system as being down until it has remained down for the hold-time period. Similarly, an interface is not advertised as being up until it has remained up for the hold-time period.



NOTE: The hold-time option is not available for controller interfaces.

Default Interface transitions are not damped.

Options down *milliseconds*—Hold time to use when an interface transitions from up to down. The JUNOS Software advertises the transition within 100 milliseconds of the time value you specify.

Range: 0 through 4,294,967,295 milliseconds

Default: 0 milliseconds (interface transitions are not damped)

up *milliseconds*—Hold time to use when an interface transitions from down to up. The JUNOS Software advertises the transition within 100 milliseconds of the time value you specify.

Range: 0 through 4,294,967,295 milliseconds

Default: 0 milliseconds (interface transitions are not damped)

Usage Guidelines See “Damping Interface Transitions” on page 138.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics advertise-interval

hold-time (SONET/SDH Defect Triggers)

Syntax `hold-time up milliseconds down milliseconds;`

Hierarchy Level [edit interfaces *interface-name* sonet-options trigger defect]

Release Information Statement introduced before JUNOS Release 7.4.

Description For ATM over SONET/SDH and SONET/SDH interfaces only, apply up and down hold times to SONET/SDH defect triggers. When you apply a down hold time to a defect, the defect must remain present for at least the hold-time period before the interface is marked down. When you apply an up hold time to a defect, the defect must remain absent for at least the hold-time period before the interface is marked up, assuming no other defect is outstanding.



NOTE: On M Series and T Series platforms with Channelized SONET IQ PICs and Channelized SONET IQE PICs, the SONET defect alarm trigger **hold-time** statement is not supported.

Default If you do not include this statement, when a defect is detected the interface is marked down immediately, and when the defect becomes absent the interface is marked up immediately.

Options `down milliseconds`—Hold time to wait before the interface is marked down.

Range: 1 through 65,534 milliseconds

Default: No hold time

`up milliseconds`—Hold time to wait before the interface is marked up.

Range: 1 through 65,534 milliseconds

Default: No hold time

Usage Guidelines See “Configuring SONET/SDH Defect Hold Times” on page 856.

Required Privilege Level `interface`—To view this statement in the configuration.
 `interface-control`—To add this statement to the configuration.

host

Syntax `host hostname {
 facility-override;
 log-prefix prefix-number;
 services priority-level;
 }`

Hierarchy Level [edit interfaces *interface-name* services-options syslog]

Release Information Statement introduced before JUNOS Release 7.4.

Description Specify hostname for system logging utility.

Options *hostname*—Name of system logging utility host machine.

The remaining statements are explained separately.

Usage Guidelines See the *JUNOS Services Interfaces Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

idle-cycle-flag

Syntax `idle-cycle-flag value;`

Hierarchy Level `[edit interfaces e1-fpc/pic/port],`
`[edit interfaces t1-fpc/pic/port],`
`[edit interfaces interface-name ds0-options],`
`[edit interfaces interface-name e1-options],`
`[edit interfaces interface-name e3-options],`
`[edit interfaces interface-name serial-options],`
`[edit interfaces interface-name t1-options],`
`[edit interfaces interface-name t3-options]`

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the value that the DS0, E1, E3, T1, or T3 interface transmits during idle cycles.



NOTE: When configuring E1 or T1 interfaces on 10-port Channelized E1/T1 IQE PICs, the `idle-cycle-flag` statement must be included at the `[edit interfaces e1-fpc/pic/port]` or `[edit interfaces t1-fpc/pic/port]` hierarchy level as appropriate.

Options *value*—Value to transmit in the idle cycles:

- `flags`—Transmit the value 0x7E.
- `ones`—Transmit the value 0xFF (all ones).

Default: `Flags`

Usage Guidelines See “Configuring the E1 Idle Cycle Flag” on page 546, “Configuring the E3 Idle Cycle Flag” on page 555, “Configuring the T1 Idle Cycle Flag” on page 566, and “Configuring the T3 Idle Cycle Flag” on page 575.

Required Privilege Level `interface`—To view this statement in the configuration.
`interface-control`—To add this statement to the configuration.

idle-time

Syntax	<code>idle-time time;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series Services Routers only. On Ethernet interfaces configured for DLSw, configure the time for which a TCP connection between DLSw peers will stay up without any circuit using the connection.
Options	<i>time</i> —Number of seconds. Range: 1 through 60000 seconds Default: 10 seconds
Usage Guidelines	See “Configuring LLC2 Options” on page 180.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

idle-timeout

Syntax	<code>idle-timeout seconds;</code>
Hierarchy Level	[edit interfaces <i>dln</i> unit <i>logical-unit-number</i> dialer-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	On J Series Services Routers with ISDN interfaces, configure the number of seconds the link is idle before losing connectivity.
Options	<i>seconds</i> —Time for which the connection can remain idle. For interfaces configured to use a filter for traffic, the idle timeout is based on traffic. Range: 1 through 429497295 Default: 120 seconds
Usage Guidelines	See “Configuring ISDN Logical Interface Properties” on page 823.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

ieee802.1p

Syntax	ieee802.1p premium [<i>values</i>];
Hierarchy Level	[edit interfaces <i>interface-name</i> gigether-options ethernet-switch-profile ethernet-policer-profile input-priority-map]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, configure premium priority values for IEEE 802.1p input traffic.
Options	<i>values</i> —Define IEEE 802.1p priority values to be treated as premium. Range: 0 through 7
Usage Guidelines	See “Specifying an Input Priority Map” on page 758.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

if-exceeding

Syntax	if-exceeding { bandwidth-limit <i>bandwidth</i> ; burst-size-limit <i>burst</i> ; }
Hierarchy Level	[edit firewall hierarchical-policer aggregate], [edit firewall hierarchical-policer premium]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, to specify bandwidth and burst limits for an aggregate level of a hierarchical policer, use the if-exceeding statement at the [edit firewall hierarchical-policer aggregate] or [edit firewall hierarchical-policer premium] hierarchy level.
Options	Options are described separately.
Usage Guidelines	See “Applying Policers” on page 194 and the <i>JUNOS Class of Service Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ignore

Syntax	ignore;
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options trigger defect]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM over SONET/SDH and SONET/SDH interfaces only, ignore a specific SONET/SDH defect trigger.
Default	If you do not include this statement, all defects are honored with no hold time.
Usage Guidelines	See “Configuring SONET/SDH Defect Triggers to Be Ignored” on page 855.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	hold-time

ignore-all

Syntax	ignore-all;
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options dce-options], [edit interfaces <i>interface-name</i> serial-options dte-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Ignore all control leads. You can include the ignore-all statement in the configuration only if you do not explicitly enable other signal handling options at the dte-options hierarchy level.
Usage Guidelines	See “Configuring the Serial Signal Handling” on page 271.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ignore-l3-incompletes

Syntax	ignore-l3-incompletes;
Hierarchy Level	[edit interfaces <i>interface-name</i> fastether-options], [edit interfaces <i>interface-name</i> gigether-options]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Ignore the counting of Layer 3 incomplete errors on Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces.
Usage Guidelines	See “Ignoring Layer 3 Incomplete Errors” on page 594
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ilmi

Syntax	ilmi;
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> atm-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable the router to communicate with directly attached ATM switches and routers. The router uses the VC 0.16 to communicate with the ATM switch or router. Once configured, you can display the IP address and port number of an ATM switch or router using the show interfaces <i>interface-name</i> switch-id command.
Usage Guidelines	See “Configuring Communication with Directly Attached ATM Switches and Routers” on page 291.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	show ilmi and show ilmi statistics commands in the <i>JUNOS Interfaces Command Reference</i> .

inactivity-timeout

Syntax	<code>inactivity-timeout seconds;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For adaptive services interfaces, configure the inactivity timeout period for established flows. The timeout configured in the application protocol definition overrides this value.
Options	<i>seconds</i> —Timeout period, in seconds. Range: 4 through 86,400 seconds Default: 30 seconds
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

incoming-called-number

Syntax	<code>incoming-called-number number <reject>;</code>
Hierarchy Level	[edit interfaces <i>br-pim/0/port</i> isdn-options]
Release Information	Statement introduced on JUNOS Release 7.5.
Description	On J Series Services Routers with interfaces configured for ISDN, screen incoming calls. If the incoming number is configured, the call is accepted. If the reject option is specified with the number, the call is rejected. If no numbers are configured, all calls are accepted.
Options	<i>number</i> —(Optional) Incoming caller number. Multiple numbers can be configured, up to a maximum of 30 entries. Only a precise match is a valid match. For example, the configured caller number 1-222-333-4444 or 222-333-4444 will match the incoming caller number 1-222-333-4444. <i>reject</i> —(Optional) Rejects the incoming number.
Usage Guidelines	See “Configuring an ISDN Interface to Screen Incoming Calls” on page 823.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Services Interfaces Configuration Guide</i>

incoming-map

Syntax incoming-map {
 caller *caller-number* | accept-all;
 }

Hierarchy Level [edit interfaces *dlm* unit *logical-unit-number* dialer-options],
 [edit logical-systems *logical-system-name* interfaces *dlm* unit *logical-unit-number*
 dialer-options]

Release Information Statement introduced in JUNOS Release 7.5.

Description On J Series Services Routers with interfaces configured for ISDN, specify the dialer to accept incoming calls.

The statements are explained separately.



NOTE: The `incoming-map` statement is mandatory for the router to accept any incoming ISDN calls.

Usage Guidelines See “Configuring Dial-In and Callback” on page 832.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics *J-series Services Router Basic LAN and WAN Access Configuration Guide*

indication

Syntax	indication (ignore normal require);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options dce-options], [edit interfaces <i>interface-name</i> serial-options dte-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For X.21 interfaces only, configure the from-DCE signal indication.
Options	ignore—The from-DCE signal is ignored. normal—Normal indication signal handling as defined by ITU-T Recommendation X.21. require—The from-DCE signal must be asserted. Default: normal
Usage Guidelines	See “Configuring the Serial Signal Handling” on page 271.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

indication-polarity

Syntax	indication-polarity (negative positive);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For X.21 interfaces only, configure the indication signal polarity.
Options	positive—Positive signal polarity. negative—Negative signal polarity. Default: positive
Usage Guidelines	See “Configuring Serial Signal Polarities” on page 274.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ingress-rate-limit

Syntax	ingress-rate-limit <i>rate</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> fastether-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Perform port-based rate limiting on ingress traffic arriving on Fast Ethernet 8-port, 12-port, and 48-port PICs.
Options	<i>rate</i> —Traffic rate, in megabits per second (Mbps). Range: 1 through 100 Mbps.
Usage Guidelines	See “Configuring the Ingress Rate Limit” on page 597.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

init-command-string

Syntax	<code>init-command-string <i>initialization-command-string</i>;</code>
Hierarchy Level	[edit interfaces umd0 modem-options]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	<p>For J Series Services Routers, configure the command string used to initialize the USB modem.</p> <p>When you connect the USB modem to the USB port on a Services Router, the router applies the modem AT commands configured in the <code>init-command-string</code> command to the initialization commands on the modem.</p> <p>For example, the initialization command string <code>ATSO = 2\n</code> configures the USB modem to pick up a call after 2 rings.</p> <p>If you do not include the <code>init-command-string</code> statement, the router applies the default initialization string to the modem.</p>
Options	<p><i>initialization-command-string</i>—Specify an initialization command string using the following AT command values:</p> <ul style="list-style-type: none"> ■ <code>%C0</code>—Disables data compression. ■ <code>&C1</code>—Disables reset of the modem when it loses the carrier signal. ■ <code>&Q8</code>—Enables Microcom Networking Protocol (MNP) error control mode. ■ <code>AT</code>—Attention. Informs the modem that a command follows. ■ <code>E0</code>—Disables the display on the local terminal of commands issued to the modem from the local terminal. ■ <code>Q0</code>—Enables the display of result codes. ■ <code>S0=0</code>—Disables the auto-answer feature, whereby the modem automatically answers calls. ■ <code>S7=45</code>—Instructs the modem to wait 45 seconds for a telecommunications service provider (carrier) signal before terminating the call. ■ <code>V1</code>—Displays result codes as words. <p>Default: <code>AT S7 = 45 S0 = 0 V1 X4 &C1 E0 Q0 &Q8 %C0</code></p>
Usage Guidelines	See “Specifying a USB Modem Interface on J Series Routers” on page 93.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

initial-route-check

Syntax	initial-route-check <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>dln</i> unit <i>logical-unit-number</i> dialer-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	On J Series Services Routers with ISDN interfaces, allows the router to check whether the primary route is up after the initial startup of the router is complete and the timer expires.
Options	<i>seconds</i> —How long to wait to check if the primary interface is up after the router comes up. Range: 1 through 300 seconds Default: 120 seconds
Usage Guidelines	See “Configuring ISDN Interfaces” on page 819.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

inner-tag-protocol-id

Syntax	inner-tag-protocol-id <i>tpid</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, configure the IEEE 802.1Q TPID value to rewrite for the inner tag. All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces <i>interface-name</i> <i>giether-options</i> ethernet-switch-profile tag-protocol-id [<i>tpids</i>]] hierarchy level.
Default	If the inner-tag-protocol-id statement is not configured, the TPID value is 0x8100.
Usage Guidelines	See “Configuring Inner and Outer TPIDs and VLAN IDs” on page 645.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

inner-vlan-id

Syntax	<code>inner-vlan-id <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	<p>For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, specify the VLAN ID to rewrite for the inner tag of the final packet.</p> <p>You cannot include the <code>inner-vlan-id</code> statement with the <code>swap</code> statement, <code>swap-push</code> statement, <code>push-push</code> statement, or <code>push-swap</code> statement and the <code>inner-vlan-id</code> statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map] hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the <code>inner-vlan-id</code> statement you include at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p>
Options	<p><i>number</i>—VLAN ID number.</p> <p>Range: 0 through 4094</p>
Usage Guidelines	See “Configuring Inner and Outer TPIDs and VLAN IDs” on page 645.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

inner-vlan-id-range

Syntax	<code>inner-vlan-id-range start <i>start-id</i> end <i>end-id</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>],
Release Information	Statement introduced in JUNOS Release 9.0.
Description	The range of VLAN IDs to be used in the ATM-to-Ethernet interworking cross-connect. Specify the starting VLAN ID and ending VLAN ID.
Options	<code>start-id</code> —The lowest VLAN ID to be used. <code>end-id</code> —The highest VLAN ID to be used. Range: 32 through 4094
Usage Guidelines	See “Configuring ATM-to-Ethernet Interworking” on page 229.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.

input

Syntax	<code>input { service-set <i>service-set-name</i> <service-filter <i>filter-name</i>>; post-service-filter <i>filter-name</i>; }</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define one or more input service sets and filters, and one postservice filter to be applied to traffic.
Options	The remaining statements are explained separately.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.

input-list

Syntax	input-list [<i>filter-names</i>];
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> filter], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> filter]
Release Information	Statement introduced in JUNOS Release 7.6.
Description	Apply a group of filters to evaluate when packets are received on an interface.
Options	[<i>filter-names</i>]—Name of a filter to evaluate when packets are received on the interface. Up to 16 filters can be included in a filter input list.
Usage Guidelines	See “Applying a Filter to an Interface” on page 203.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	output-list, <i>JUNOS Policy Framework Configuration Guide</i> , <i>JUNOS System Basics Configuration Guide</i>

input-policer

Syntax	input-policer <i>policer-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Associate a Layer 2 policer with a logical interface. The input-policer and input-three-color statements are mutually exclusive.
Options	<i>policer-name</i> —Name of the policer that you define at the [edit firewall] hierarchy level.
Usage Guidelines	See “Applying a Policer” on page 759.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	output-policer

input-priority-map

Syntax	input-priority-map { ieee802.1p premium [<i>values</i>]; }
Hierarchy Level	[edit interfaces <i>interface-name</i> together-options ethernet-switch-profile ethernet-policer-profile]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the input policer priority map to be applied to incoming frames on this interface. The statements are explained separately.
Usage Guidelines	See “Specifying an Input Priority Map” on page 758.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	output-priority-map

input-three-color

Syntax	input-three-color <i>policer-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Associate a Layer 2, three-color policer with a logical interface. The input-three-color and input-policer statements are mutually exclusive.
Options	<i>policer-name</i> —Name of the three-color policer that you define at the [edit firewall] hierarchy level.
Usage Guidelines	See “Applying a Policer” on page 759.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	output-three-color

input-vlan-map

See the following sections:

- input-vlan-map (Gigabit Ethernet IQ) on page 1065
- input-vlan-map (Aggregated Ethernet) on page 1066

input-vlan-map (Gigabit Ethernet IQ)

Syntax input-vlan-map {
 (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
 inner-tag-protocol-id *tpid*;
 inner-vlan-id *number*;
 tag-protocol-id *tpid*;
 vlan-id *number*;
 }

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before JUNOS Release 7.4.
 pop-pop, pop-swap, push-push, swap-push, and swap-swap statements introduced in JUNOS Release 8.1.

Description For Gigabit Ethernet IQ interfaces only, define the rewrite profile to be applied to incoming frames on this logical interface.

The statements are explained separately.

Usage Guidelines See “Stacking a VLAN Tag” on page 648.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics output-vlan-map

input-vlan-map (Aggregated Ethernet)

Syntax	input-vlan-map { (pop push swap); tag-protocol-id <i>tpid</i> ; vlan-id <i>number</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	For aggregated Ethernet interfaces using Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces only, define the rewrite profile to be applied to incoming frames on this logical interface. The statements are explained separately.
Usage Guidelines	See “Stacking a VLAN Tag” on page 648.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	output-vlan-map

instance

Syntax	instance <i>vpls-instance-name</i> ;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>name</i>]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	Specify the VPLS instance of the default maintenance domain.
Usage Guidelines	See “Configuring Maintenance Intermediate Points” on page 683.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	maintenance-domain

interface

See the following sections:

- interface (DLSw Ethernet Redundancy) on page 1067
- interface (Hierarchical CoS Schedulers) on page 1068
- interface (IEEE 802.1x) on page 1069
- interface (IEEE 802.1ag OAM Connectivity-Fault Management) on page 1070
- interface (OAM Link-Fault Management) on page 1071



NOTE: For information about the interface statement available at the [edit interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-number* track] and [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-number* track] hierarchy levels, see the *JUNOS High Availability Configuration Guide*.

interface (DLSw Ethernet Redundancy)

Syntax	interface <i>interface-name</i> priority-cost <i>cost</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	For J Series Services Routers only, on Ethernet interfaces configured for DLSw Ethernet redundancy, enable tracking options for an interface.
Options	<i>interface-name</i> —Interface name. Include the logical portion of the name, which corresponds to the logical unit number. The statements are explained separately. <i>priority-cost cost</i> —Cost value that is subtracted from the priority value when remote peer connectivity is lost. Specify a value from 1 through 254.
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 181.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Services Interfaces Configuration Guide</i>

interface (Hierarchical CoS Schedulers)

Syntax	interface <i>interface-name</i> ;
Hierarchy Level	[edit interfaces interface-set (Ethernet Interfaces) <i>interface-set-name</i>]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Interface that is a member of the interface set. Supported on Ethernet interfaces and IP demux interfaces on an MX Series router.
Usage Guidelines	See the <i>JUNOS Class of Service Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

interface (IEEE 802.1x)

Syntax interface *interface-ids* {
 supplicant *single*;
 retries *integer*;
 quiet-period *seconds*;
 transmit-period *seconds*;
 reauthentication (disable | interval *seconds*);
 supplicant-timeout *seconds*;
 server-timeout *seconds*;
 maximum-requests *seconds*;
 }

Hierarchy Level [edit protocols dot1x authenticator]

Release Information Statement introduced in JUNOS Release 9.3.

Description Use this statement to configure the 802.1x Port-Based Network Access Control protocol-specific Ethernet interface options.

Default The default values are provided for the options below on the respective statement pages.

Options maximum-requests—Specify the maximum number of retransmission times for an EAPOL Request packet to the client before it times out the authentication session.

quiet-period—Specify the number of seconds the port remains in the wait state following a failed authentication exchange with the client, before reattempting the authentication.

reauthentication—Includes two options:

- disable—Periodic reauthentication of the client is disabled.
- interval—Specify the periodic reauthentication time interval.

retries—Specify the number of tries after which the port remains in the wait state for quiet-period seconds before reattempting the authentication.

server-timeout—Specify the number of seconds the port waits for a reply when relaying a response from the client to the authentication server before timing out and invoking the server-fail action.

supplicant (*single*)—Specify supplicant single mode. See the usage guidelines to configure other modes.

supplicant-timeout—Specify the number of seconds the port waits for a response when relaying a request from the authentication server to the client before resending the request.

transmit-period—Specify the number of seconds the port waits before retransmitting the initial EAPOL PDUs to the client.

Usage Guidelines	See “Configuring IEEE 802.1x Port-Based Network Access Control” on page 741.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	authenticator, dot1x

interface (IEEE 802.1ag OAM Connectivity-Fault Management)

Syntax	interface (<i>interface-name</i> ((ge- xe-) (<i>fpc/pic/port</i> <i>fpc/pic/port.unit-number</i> <i>fpc/pic/port.unit-number vlan vlan-id</i>)));
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> maintenance-association <i>ma-name</i> mep <i>mep-id</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	For Ethernet interfaces on M320, MX Series, and T Series routers, configure IEEE 802.1ag Operation, Administration, and Management (OAM) support. For Gigabit Ethernet interfaces and 10-Gigabit Ethernet interfaces on MX Series routers, configure IEEE 802.1ag Connectivity Fault Management (CFM) support on trunk interface ports.
Options	interface-name —Interface to which the MEP is attached. It could be a physical Ethernet interface, logical Ethernet interface, or on a specific VLAN of a trunk port interface (MX Series only).
Usage Guidelines	See “Configuring the Maintenance End Point Interface” on page 687.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

interface (OAM Link-Fault Management)

Syntax interface *interface-name* {
 apply-action-profile *profile-name*;
 link-discovery (active | passive);
 pdu-interval *interval*;
 pdu-threshold *threshold-value*;
 remote-loopback;
 event-thresholds {
 frame-error *count*;
 frame-period *count*;
 frame-period-summary *count*;
 symbol-period *count*;
 }
 negotiation-options {
 allow-remote-loopback;
 no-allow-link-events;
 }
 }

Hierarchy Level [edit protocols oam ethernet link-fault-management]

Release Information Statement introduced in JUNOS Release 8.2.

Description For Ethernet interfaces on M320, MX Series, and T Series routers, configure IEEE 802.3ah Operation, Administration, and Management (OAM) support.

Options interface *interface-name*—Interface to be enabled for IEEE 802.3ah link fault management OAM support.

Range: 1 through 10 interfaces can be tracked.

The remaining statements are described separately.

Usage Guidelines See “Enabling IEEE 802.3ah OAM Support” on page 746.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

interfaces

Syntax	interfaces { ... }
Hierarchy Level	[edit]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure interfaces on the router.
Default	The management and internal Ethernet interfaces are automatically configured. You must configure all other interfaces.
Usage Guidelines	See individual chapters.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

interface-down

Syntax	interface-down;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management action-profile <i>profile-name</i> default-action]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Bring the interface down when a remote MEP connectivity failure is detected.
Usage Guidelines	See “Configuring a CFM Action Profile Action” on page 688.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

interface-mode

Syntax	interface-mode (access trunk);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Determines whether the logical interface accepts or discards packets based on VLAN tags. Specify the trunk option to accept packets with a VLAN ID that matches the list of VLAN IDs specified in the vlan-id-list statement, then forward the packet within the bridge domain configured with the matching VLAN ID. Specify the access option to accept packets with no VLAN ID, then forward the packet within the bridge domain configured with the VLAN ID that matches the VLAN ID specified in the vlan-id statement.
Options	<p>access—Configure a logical interface to accept untagged packets. Specify the VLAN to which this interface belongs using the vlan-id statement.</p> <p>trunk—Configure a single logical interface to accept packets tagged with any VLAN ID specified with the vlan-id-list statement.</p>
Usage Guidelines	See “Configuring a Logical Interface for Access Mode” on page 619 and “Configuring a Logical Interface for Trunk Mode” on page 620.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

interface-set

See the following sections:

- interface-set (Ethernet Interfaces) on page 1074
- interface-set (IP Demux Interfaces) on page 1074

interface-set (Ethernet Interfaces)

Syntax interface-set *interface-set-name* {
 interface *ethernet-interface-name* {
 (unit *unit-number* | vlan-tags-outer *vlan-tag*);
 }
 }

Hierarchy Level [edit interfaces]

Release Information Statement introduced in JUNOS Release 8.5.

Description The set of interfaces used to configure hierarchical CoS schedulers on Ethernet interfaces on the MX Series router.

The remaining statements are described separately.

Usage Guidelines See the *JUNOS Class of Service Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

interface-set (IP Demux Interfaces)

Syntax interface-set *interface-set-name* {
 interface *interface-name* {
 unit *unit-number*;
 }
 }

Hierarchy Level [edit interfaces]

Release Information Statement introduced in JUNOS Release 9.2.

Description The set of interfaces used to configure hierarchical CoS schedulers for subscribers on IP demux interfaces on the MX Series router.

The remaining statements are described separately.

Usage Guidelines See the *JUNOS Subscriber Access Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

interface-status-tlv

Syntax	interface-status-tlv <down lower-layer-down>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management action-profile (Defining for CFM) <i>tlv-action</i> event]
Release Information	Statement introduced in JUNOS Release 9.6.
Description	Defines an action-profile consisting of various events and the action. Based on values of interface-status-tlv in the received CCM packets, specific action such as <i>interface-down</i> can be taken using action-profile (Defining for CFM) options.
Options	<p>down—When the incoming CCM packet contains interface status TLV with value down, the action will be triggered for this action-profile.</p> <p>lower-layer-down—When the incoming CCM packet contains interface status TLV with value lower-layer-down, the action will be triggered for this action-profile.</p>
Usage Guidelines	See “Configuring Remote MEP Action Profile Support” on page 707.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

interface-switch

Syntax	<pre>interface-switch <i>connection-name</i> { interface <i>interface-name.unit-number</i>; interface <i>interface-name.unit-number</i>; }</pre>
Hierarchy Level	[edit protocols connections]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Configure Layer 2 switching cross-connects. The cross-connect is bidirectional, so packets received on the first interface are transmitted out the second interface, and those received on the second interface are transmitted out the first interface.</p> <p>For Layer 2 switching cross-connects to work, you must also configure MPLS.</p>
Options	interface <i>interface-name.unit-number</i> —Interface name. Include the logical portion of the name, which corresponds to the logical unit number.
Usage Guidelines	See “Defining the Connection for Switching Cross-Connects” on page 228.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS MPLS Applications Configuration Guide</i>

interface-type

Syntax	interface-type (bc coc1 ct1 ct3 dc ds so t1 t3);
Hierarchy Level	[edit interfaces <i>interface-name</i> no-partition], [edit interfaces <i>interface-name</i> partition <i>partition-number</i>], [edit interfaces <i>interface-name</i> partition <i>partition-number</i> oc-slice <i>oc-slice-range</i>], [edit interfaces <i>interface-name</i> partition <i>partition-number</i> timeslot <i>timeslot-range</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For IQ and IQE interfaces only, configure the sublevel interface type.
Options	<p>bc—Dual—Port Channelized E1 and T1 ISDN PRI interface type. You can specify this interface type at the [edit interfaces <i>interface-name</i> partition <i>partition-number</i> timeslot <i>timeslot-range</i>] hierarchy level to create a bearer (B) channel <i>bc-pim/O/port:channel</i> interface for each time you want to function as an ISDN PRI B-channel.</p> <p>coc1—Channelized OC1 interface type. You can specify this interface type at the [edit interfaces <i>interface-name</i> partition <i>partition-number</i> oc-slice <i>oc-slice-range</i> interface-type <i>coc1-fpc/pic/port</i>] hierarchy level.</p> <p>ct1—Channelized T1 interface type. You can specify this interface type at the [edit interfaces <i>interface-name</i> partition <i>partition-number</i> interface-type <i>ct3-fpc/pic/port<:channel></i>] hierarchy level.</p> <p>ct3—Channelized T3 interface type. You can specify this interface type at the [edit interfaces <i>interface-name</i> partition <i>partition-number</i> oc-slice <i>oc-slice-range</i> interface-type <i>coc1-fpc/pic/port:channel</i> no-partition] hierarchy level.</p> <p>dc—Dual-Port Channelized E1 and T1 ISDN PRI interface type. You can specify this interface type at the [edit interfaces <i>interface-name</i> partition <i>partition-number</i> timeslot <i>timeslot-range</i>] hierarchy level to create a (D) channel <i>dc-pim/O/port</i> to control the B-channels.</p> <p>ds—DS0 interface type. You can specify this interface type at the [edit interfaces <i>interface-name</i> partition <i>partition-number</i> interface-type (<i>ce1-fpc/pic/port</i> <i>ct1-fpc/pic/port<:channel></i>)] hierarchy level.</p> <p>so—SONET/SDH interface type. You can specify this interface type at the [edit interfaces <i>interface-name</i> partition <i>partition-number</i> oc-slice <i>oc-slice-range</i> interface-type <i>coc12-fpc/pic/port</i>] hierarchy level.</p> <p>t1—T1 interface type. You can specify this interface type at the [edit interfaces <i>interface-name</i> partition <i>partition-number</i> oc-slice <i>oc-slice-range</i> interface-type (<i>coc12-fpc/pic/port</i> <i>coc1-fpc/pic/port</i>)] hierarchy level.</p> <p>t3—T3 interface type. You can specify this interface type at the [edit interfaces <i>interface-name</i> partition <i>partition-number</i> oc-slice <i>oc-slice-range</i> interface-type (<i>coc12-fpc/pic/port</i> <i>coc1-fpc/pic/port:channel</i> no-partition)] hierarchy level.</p>

Usage Guidelines See “Configuring Channelized E1 Interfaces” on page 501, “Configuring Channelized OC12/STM4 Interfaces” on page 423, and “Configuring Channelized T3 Interfaces” on page 479.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

interleave-fragments

Syntax interleave-fragments;

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description For link services interfaces only, interleave long packets with high-priority packets.

Allows small delay-sensitive packets, such as Voice over IP (VoIP) packets, to interleave with long fragmented packets. This minimizes the latency of delay-sensitive packets.

Usage Guidelines See the *JUNOS Services Interfaces Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.


interval

Syntax	interval (10m 10s 1m 1s 100ms 10ms);
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> maintenance-association <i>ma-name</i> continuity-check]
Release Information	Statement introduced in JUNOS Release 8.4. Ten milliseconds option introduced in JUNOS Release 9.1.
Description	The time between continuity check messages.
Options	10m—10 minutes. 10s—10 seconds. 1m—1 minute. 1s—1 second. 100ms—100 milliseconds. 10ms—10 milliseconds.
Usage Guidelines	See “Configuring the Continuity Check Interval” on page 686.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

inverse-arp

Syntax	inverse-arp;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> multipoint-destination <i>destination</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> multipoint-destination <i>destination</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM encapsulation, enable responses to received inverse ATM ARP requests. For Frame Relay encapsulation, enable responses to received inverse Frame Relay ARP requests.
Default	Inverse ARP is disabled on all ATM and Frame Relay interfaces.
Usage Guidelines	See “Configuring Inverse ATM1 or ATM2 ARP” on page 318 or “Configuring Inverse Frame Relay ARP” on page 379.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

invert-data

Syntax	invert-data;
Hierarchy Level	[edit interfaces <i>e1-fpc/pic/port</i>], [edit interfaces <i>t1-fpc/pic/port</i>], [edit interfaces <i>interface-name</i> ds0-options], [edit interfaces <i>interface-name</i> e1-options], [edit interfaces <i>interface-name</i> t1-options], [edit interfaces <i>interface-name</i> e3-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Invert the transmission of unused data bits on the DS0, E1, E3, and T1 interface.
	NOTE: When configuring E1 or T1 interfaces on 10-port Channelized E1/T1 IQE PICs, the invert-data statement must be included at the [edit interfaces <i>e1-fpc/pic/port</i>] or [edit interfaces <i>t1-fpc/pic/port</i>] hierarchy level as appropriate.
Usage Guidelines	See “Configuring E1 Data Inversion” on page 546, “Configuring E3 Data Inversion” on page 555, and “Configuring T1 Data Inversion” on page 563.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ipsec-sa

Syntax	ipsec-sa <i>sa-name</i> ;
Hierarchy Level	[edit interfaces <i>es-fpc/pic/port</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>es-fpc/pic/port</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the IP Security (IPsec) security association (SA) name associated with the interface.
Options	<i>sa-name</i> —IPsec security association name.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS System Basics Configuration Guide</i>

isdn-options

Syntax	<pre> isdn-options { bchannel-allocation (ascending descending); calling-number <i>number</i>; incoming-called-number <i>number</i> <reject>; spid1 <i>spid-string</i>; spid2 <i>spid-string</i>; static-tei-val <i>value</i>; switch-type (att5e etsi ni1 ntdms100 ntt); t310 <i>seconds</i>; tei-option (first-call power-up); } </pre>
Hierarchy Level	<pre> [edit interfaces <i>br-pim</i>/0/<i>port</i>], [edit interfaces <i>ct1-pim</i>/0/<i>port</i>], [edit interfaces <i>ce1-pim</i>/0/<i>port</i>] </pre>
Release Information	<p>Statement introduced before JUNOS Release 7.4.</p> <p>bchannel-allocation option added in JUNOS Release 8.3.</p>
Description	<p>For J Series Services Routers only. Specify the ISDN options for configuring ISDN interfaces for group and user sessions.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Configuring ISDN Physical Interface Properties” on page 821 and “Allocating B-Channels for Dialout” on page 513.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

keep-address-and-control

Syntax	keep-address-and-control;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ccc], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ccc]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For interfaces with encapsulation type PPP CCC, do not remove the address and control bytes before encapsulating the packet into a tunnel.
Default	If you do not include this statement, address and control bytes are removed before encapsulating the packet into a tunnel.
Usage Guidelines	See “Disabling the Removal of Address and Control Bytes” on page 192.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

keepalives

Syntax	keepalives <interval seconds> <down-count <i>number</i> > <up-count <i>number</i> >;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Enable the sending of keepalives on a physical interface configured with PPP, Frame Relay, or Cisco HDLC encapsulation.</p> <p>For ATM2 IQ interfaces only, you can enable keepalives on a logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:</p> <ul style="list-style-type: none"> ■ atm-ppp-llc—PPP over AAL5 LLC encapsulation. ■ atm-ppp-vc-mux—PPP over AAL5 multiplex encapsulation.
Default	Sending of keepalives is enabled by default. The default keepalive interval is 10 seconds for PPP, Frame Relay, or Cisco HDLC. The default down-count is 3 and the default up-count is 1 for PPP or Cisco HDLC.
Options	<p>down-count <i>number</i>—The number of keepalive packets a destination must fail to receive before the network takes down a link. Range: 1 through 255 Default: 3</p> <p>interval <i>seconds</i>—The time in seconds between successive keepalive requests. Range: 1 through 32767 seconds Default: 10 seconds</p> <p>up-count <i>number</i>—The number of keepalive packets a destination must receive to change a link's status from down to up. Range: 1 through 255 Default: 1</p>
Usage Guidelines	See “Configuring Keepalives” on page 126 or “Configuring Frame Relay Keepalives” on page 378.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

key

Syntax	<code>key number;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Adaptive Services PICs on M Series routers (except the M320 and M120 routers), identify an individual traffic flow within a tunnel, as defined in RFC 2890, <i>Key and Sequence Number Extensions to GRE</i> .
Options	<i>number</i> —Value of the key. Range: 0 through 4,294,967,295
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

l2tp-interface-id

Syntax	<code>l2tp-interface-id name;</code> <code>(dedicated shared);</code>
Hierarchy Level	[edit interfaces <i>sp-fpc/pic/port</i> unit <i>logical-unit-number</i> dialer-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>sp-fpc/pic/port</i> unit <i>logical-unit-number</i> dialer-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the L2TP options for configuring logical interfaces for group and user sessions.
Options	<code>(dedicated shared)</code> —Specifies whether a logical interface can host one (dedicated) or multiple (shared) sessions at one time. <i>name</i> —Interface identifier that must be replicated at the [edit access profile <i>name</i>] hierarchy level.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

lACP

See the following sections:

- lACP (802.3ad) on page 1085
- lACP (Aggregated Ethernet) on page 1086

lACP (802.3ad)

Syntax

```
lACP {
    traceoptions {
        file lACPd;
        flag all;
    }
    ppm (centralized | distributed);
}
```

Hierarchy Level [edit interfaces *interface-name* fastether-options 802.3ad],
[edit interfaces *interface-name* gigether-options 802.3ad]

Release Information Statement introduced in JUNOS Release 9.3.
The ppm (centralized | distributed) option introduced in JUNOS Release 9.4.

Description For aggregated Ethernet interfaces only, configure the Link Aggregation Control Protocol (LACP).

On MX Series routers you can specify distributed or centralized periodic packet management (PPM).

Default If you do not specify lACP as either **active** or **passive**, LACP remains passive.
If you do not specify ppm as either **centralized** or **distributed**, PPM will be distributed.

Options

- **active**—Initiate transmission of LACP packets.
- **passive**—Respond to LACP packets.
- **ppm**—Set PPM to centralized or distributed.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Aggregated Ethernet LACP” on page 627.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

lacp (Aggregated Ethernet)

Syntax lacp {
 (active | passive);
 link-protection {
 disable;
 (revertive | non-revertive);
 periodic *interval*;
 system-priority *priority*;
 }

Hierarchy Level [edit interfaces aex aggregated-ether-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description For aggregated Ethernet interfaces only, configure Link Aggregation Control Protocol (LACP).

Default If you do not specify **lacp** as either active or passive, LACP remains passive.

Options ■ active—Initiate transmission of LACP packets.
 ■ passive—Respond to LACP packets.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Aggregated Ethernet LACP” on page 627.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

layer2-policer

Syntax	<pre>layer2-policer { input-policer <i>policer-name</i>; input-three-color <i>policer-name</i>; output-policer <i>policer-name</i>; output-three-color <i>policer-name</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	<p>For 1-Gigabit Ethernet and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces on M Series and T Series routers, apply Layer 2 logical interface policers. The following policers are supported:</p> <ul style="list-style-type: none"> ■ Two-color ■ Single-rate tricolor marking (srTCM) ■ Two-rate tricolor marking (trTCM) <p>Two-color and tricolor policers are configured at the [edit firewall] hierarchy level.</p>
Options	<p>input-policer <i>policer-name</i>—Two-color input policer to associate with the interface. This statement is mutually exclusive with the input-three-color statement.</p> <p>input-three-color <i>policer-name</i>—Tricolor input policer to associate with the interface. This statement is mutually exclusive with the input-policer statement.</p> <p>output-policer <i>policer-name</i>—Two-color output policer to associate with the interface. This statement is mutually exclusive with the output-three-color statement.</p> <p>output-three-color <i>policer-name</i>—Tricolor output policer to associate with the interface. This statement is mutually exclusive with the output-policer statement.</p>
Usage Guidelines	See “Configuring Gigabit Ethernet Two-Color and Tricolor Policers” on page 763.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Class of Service Configuration Guide, JUNOS Policy Framework Configuration Guide</i>

lcp-max-conf-req

Syntax	<code>lcp-max-conf-req <i>number</i></code>
Hierarchy Level	[edit interfaces <i>so-fpc/pic/port</i> unit <i>number</i> ppp-options]
Release Information	Statement introduced in JUNOS Release 9.6.
Description	Set the maximum number of LCP Configure-Requests to be sent, after which the router goes to LCP down state.
Options	<i>number</i> —From 0 to 65,535, where 0 means send infinite LCP Configure-Requests, and any other value specifies the maximum number LCP Configure-Requests to send and then stop sending.
Usage Guidelines	See “Configuring the LCP Configure-Request Maximum Sent” on page 161.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	ppp-options

lcp-restart-timer

Syntax	<code>lcp-restart-timer <i>milliseconds</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	For interfaces with PPP, PPP TCC, PPP over Ethernet, PPP over ATM, and PPP over Frame Relay encapsulations, configure a restart timer for the Link Control Protocol (LCP) component of a PPP session.
Options	<i>milliseconds</i> —The time, in milliseconds, between successive LCP configuration requests. Range: 20 through 10000 milliseconds Default: 3 seconds
Usage Guidelines	See “Configuring the PPP Restart Timers” on page 162.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

level

Syntax	<code>level <i>number</i>;</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	A number used in CFM messages to identify the maintenance association.
Options	<p><i>number</i>—A number used to identify the maintenance domain to which the CFM message belongs.</p> <p>Range: 0 through 7</p>
Usage Guidelines	See “Configuring the Maintenance Domain Level” on page 682.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

linear-red-profile

Syntax	<code>linear-red-profile <i>profile-name</i>;</code>
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> atm-options scheduler-maps <i>map-name</i> forwarding-class <i>class-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, assign a linear RED profile to a specified forwarding class. To define the linear RED profiles, include the linear-red-profiles statement at the [edit interfaces <i>at-fpc/pic/port</i> atm-options] hierarchy level.
Default	If you do not include either the epd-threshold or the linear-red-profile statement in the forwarding class configuration, the JUNOS Software uses an EPD threshold based on the available bandwidth and other parameters.
Options	<i>profile-name</i> —Name of the linear RED profile.
Usage Guidelines	See “Configuring an ATM Scheduler Map” on page 341.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	epd-threshold, linear-red-profiles

linear-red-profiles

Syntax linear-red-profiles *profile-name* {
 high-plp-threshold *percent*;
 low-plp-threshold *percent*;
 queue-depth *cells*;
 }

Hierarchy Level [edit interfaces *at-fpc/pic/port* atm-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description For ATM2 IQ interfaces only, define CoS virtual circuit drop profiles for RED. When a packet arrives, RED checks the queue fill level. If the fill level corresponds to a nonzero drop probability, the RED algorithm determines whether to drop the arriving packet.


Options *profile-name*—Name of the drop profile.

The statements are explained separately.

Usage Guidelines See “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

line-encoding

Syntax	line-encoding (ami b8zs);
Hierarchy Level	[edit interfaces <i>ct1-fpc/pic/port</i>], [edit interfaces <i>interface-name</i> t1-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the line encoding format on the T1 interface.
	NOTE: When configuring CT1 interfaces on the 10-port Channelized E1/T1 IQE PIC, the <code>line-encoding</code> statement must be included at the [edit interfaces <i>ct1-fpc/pic/port</i>] hierarchy level.
Default	The default line encoding is B8ZS.
Options	ami—Use Alternate Mark Inversion (AMI) line encoding. b8zs—Use bipolar with 8-zeros substitution (B8ZS) line encoding.
Usage Guidelines	See “Configuring T1 Line Encoding” on page 564.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

line-protocol

Syntax	line-protocol <i>protocol</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For serial interfaces only, configure the line protocol.
Options	<i>protocol</i> —You can specify the one of the following line protocols: <ul style="list-style-type: none"> ■ eia530—Line protocol EIA-530 ■ v.35—Line protocol V.35 ■ x.21—Line protocol X.21
Usage Guidelines	See “Configuring the Serial Line Protocol” on page 265.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

line-rate

Syntax	line-rate <i>line-rate</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> shdsl-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> shdsl-options]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series Services Routers only, configure the SHDSL line rate.
Options	<p><i>line-rate</i>—SHDSL line rate, in Kbps. Possible values are:</p> <p>2-wire (Kbps): 192, 256, 320, 384, 448, 512, 576, 640, 704, 768, 832, 896, 960, 1024, 1088, 1152, 1216, 1280, 1344, 1408, 1472, 1536, 1600, 1664, 1728, 1792, 1856, 1920, 1984, 2048, 2112, 2176, 2240, 2304, auto</p> <p>4-wire (Kbps): 384, 512, 640, 768, 896, 1024, 1152, 1280, 1408, 1536, 1664, 1792, 1920, 2048, 2176, 2304, 2432, 2560, 2688, 2816, 2944, 3072, 3200, 3328, 3456, 3584, 3712, 3840, 3968, 4096, 4224, 4352, 4480, 4608</p> <p>Default: For 2-wire mode, auto; for 4-wire mode, 4608 Kbps</p>
Usage Guidelines	See “Configuring ATM-over-SHDSL Interfaces” on page 361.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

link-adjacency-loss

Syntax	link-adjacency-loss;
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile event]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Loss of adjacency with IEEE 802.3ah link-fault management peer event. When included, the loss-of-adjacency event triggers the action specified under the <i>action</i> statement.
Usage Guidelines	See “Monitoring the Loss of Link Adjacency” on page 750.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

link-down

Syntax	link-down;
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile action]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Mark the interface down for transit traffic.
Usage Guidelines	See “Specifying the Actions to Be Taken for Link-Fault Management Events” on page 749.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

link-discovery

Syntax	link-discovery (active passive);
Hierarchy Level	[edit protocols oam ethernet link-fault-management interface <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	For Ethernet interfaces on M320, M120, MX Series, and T Series routers, specify the discovery mode used for IEEE 802.3ah Operation, Administration, and Management (OAM) support. The discovery process is triggered automatically when OAM 802.3ah functionality is enabled on a port. Link monitoring is done when the interface sends periodic OAM PDUs.
Options	(active passive)—Passive or active mode. In active mode, the interface discovers and monitors the peer on the link if the peer also supports IEEE 802.3ah OAM functionality. In passive mode, the peer initiates the discovery process. Once the discovery process is initiated, both sides participate in discovery.
Usage Guidelines	See “Configuring Link Discovery” on page 746.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

link-event-rate

Syntax link-event-rate {
 frame-error *count*;
 frame-period *count*;
 frame-period-summary *count*;
 symbol-period *count*;
 }

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event]

Release Information Statement introduced in JUNOS Release 8.5.

Description Configure the number of link-fault management events per second.

Usage Guidelines See “Configuring Threshold Values for Fault Events in an Action Profile” on page 750.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

link-fault-management

Syntax

```

link-fault-management {
  action-profile profile-name {
    action {
      syslog;
      link-down;
      send-critical-event;
    }
    event {
      link-adjacency-loss;
      link-event-rate {
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
      }
      protocol-down;
    }
  }
  interface interface-name {
    apply-action-profile profile-name;
    link-discovery (active | passive);
    pdu-interval interval;
    pdu-threshold threshold-value;
    remote-loopback;
    event-thresholds {
      frame-error count;
      frame-period count;
      frame-period-summary count;
      symbol-period count;
    }
    negotiation-options {
      allow-remote-loopback;
      no-allow-link-events;
    }
  }
}

```

Hierarchy Level [edit protocols oam ethernet]

Release Information Statement introduced in JUNOS Release 8.2.

Description For Ethernet interfaces on M320, M120, MX Series, and T Series routers, specify fault signaling and detection for IEEE 802.3ah Operation, Administration, and Management (OAM) support.

The remaining statements are explained separately.

Usage Guidelines See “Enabling IEEE 802.3ah OAM Support” on page 746.


Required Privilege Level interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

link-layer-overhead

Syntax	link-layer-overhead <i>percent</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For AS PIC or MultiServices PIC link services IQ interfaces (lsq) only, configure the percentage of total bundle bandwidth to be set aside for link-layer overhead.
Options	<i>percent</i> —Percentage of total bundle bandwidth to be set aside for link-layer overhead. Range: 0 through 50 percent Default: 4 percent
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

link-mode

Syntax	link-mode (full-duplex half-duplex);
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>ge-pim</i> /0/0 switch-options switch-port <i>port-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the device's link connection characteristic.
Options	<p>full-duplex—Connection is full duplex.</p> <p>half-duplex—Connection is half duplex.</p> <p>Default: Fast Ethernet interfaces, except the J Series ePIM Fast Ethernet interfaces, can operate in either full-duplex or half-duplex mode. The router's management Ethernet interface, fxp0 or em0, the built-in Fast Ethernet interfaces on the FIC (M7i router), and the Gigabit Ethernet ports on J Series Services Routers with uPIMs installed and configured for access switching mode autonegotiate whether to operate in full-duplex or half-duplex mode. Unless otherwise noted here, all other interfaces operate only in full-duplex mode.</p>
	<p>NOTE: On J Series ePIM Fast Ethernet interfaces, if you specify half-duplex (or if full-duplex mode is not autonegotiated), the following message is written to the system log: "Half-duplex mode not supported on this PIC, forcing full-duplex mode."</p>
Usage Guidelines	See "Configuring the Link Characteristics on Ethernet Interfaces" on page 595.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

link-protection

Syntax	link-protection { disable (Link Protection); (revertive non-revertive); }
Hierarchy Level	[edit interfaces aex aggregated-ether-options]
Release Information	Statement introduced in JUNOS Release 8.3. Support for disable , revertive , and non-revertive statements added in JUNOS Release 9.3.
Description	For aggregated Ethernet interfaces only, configure link protection. In addition to enabling link protection, a primary and a secondary (backup) link must be configured to specify what links egress traffic should traverse. To configure primary and secondary links, include the primary and secondary statements at the [edit interfaces <i>ge-fpc/pic/port</i> gigether-options 802.3ad aex] hierarchy level or the [edit interfaces <i>fe-fpc/pic/port</i> fastether-options 802.3ad aex] hierarchy level.
Options	The statements are explained separately.
Usage Guidelines	See “Configuring Aggregated Ethernet Link Protection” on page 626.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

link-speed

See the following sections:

- link-speed (Aggregated Ethernet) on page 1099
- link-speed (Aggregated SONET/SDH) on page 1100

link-speed (Aggregated Ethernet)

Syntax link-speed *speed*;

Hierarchy Level [edit interfaces aex aggregated-ether-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description For aggregated Ethernet interfaces only, set the required link speed.

Options *speed*—For aggregated Ethernet links, you can specify *speed* in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000).

Aggregated Ethernet links on the M120 router can have one of the following speed values,

- 100m—Links are 100 Mbps.
- 10g—Links are 10 Gbps.
- 1g—Links are 1 Gbps.
- OC192—Links are OC192 or STM64c.

Usage Guidelines See “Configuring Aggregated Ethernet Link Speed” on page 634.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

link-speed (Aggregated SONET/SDH)

Syntax	link-speed (<i>speed</i> mixed);
Hierarchy Level	[edit interfaces asx aggregated-sonet-options]
Release Information	Statement introduced before JUNOS Release 7.4. mixed option added in Release 8.0.
Description	For aggregated SONET/SDH interfaces only, set the required link speed.
Options	<p><i>speed</i>—Aggregated SONET/SDH links can have one of the following speed values.</p> <ul style="list-style-type: none"> ■ <i>oc3</i>—Links are OC3c or STM1c. ■ <i>oc12</i>—Links are OC12c or STM4c. ■ <i>oc48</i>—Links are OC48c or STM16c. ■ <i>oc192</i>—Links are OC192c or STM64c. ■ <i>oc768</i>—Links are OC768c or STM256c. <p><i>mixed</i>—For aggregated SONET/SDH links on T Series routers, you can mix interface speeds in SONET/SDH aggregation bundles. Interface speeds from OC3 through OC768 are supported.</p>
Usage Guidelines	See “Configuring Aggregated Ethernet Link Speed” on page 634 and “Configuring Aggregated SONET/SDH Link Speed” on page 883.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

linktrace

Syntax	<pre>linktrace { age (30m 10m 1m 30s 10s); path-database-size <i>path-database-size</i>; }</pre>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Configure connectivity fault management linktrace parameters.
Usage Guidelines	See “Linktrace Protocol” on page 682.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

llc2

Syntax llc2 {
 ack-delay-time *time*;
 ack-max *count*;
 idle-time *time*;
 local-window *count*;
 max-retry *count*;
 p-bit-timeout *time*;
 redundancy-group *group-number* {
 advertise-interval *seconds*;
 map {
 local-mac *mac-address* request *mac-address*;
 }
 preempt hold-time *seconds*;
 no-preempt;
 priority *priority*;
 track {
 dls {
 peer *ip-address* priority-cost *priority*;
 destination *mac-address* priority-cost *priority*;
 }
 interface *interface-name* priority-cost *priority*;
 }
 }
 t1-time *time*;
 t2-time *time*;
 trej-time *time*;
 }

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family]

Release Information Statement introduced in JUNOS Release 7.4.

Description For J Series Services Routers only, configure the data link-layer protocol logical link control 2 (LLC2) used on a LAN. LLC2 provides connection-oriented data transfer for Ethernet interfaces configured for DLSw.

The statements are explained separately.

Usage Guidelines See “Configuring LLC2 Options” on page 180.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics *JUNOS Services Interfaces Configuration Guide*

Imi

See the following sections:

- Imi (Frame Relay) on page 1103
- Imi (Ethernet OAM) on page 1104

Imi (Frame Relay)

Syntax	<pre> Imi { lmi-type (ansi itu); n391dte number; n392dce seconds; n392dte number; n393dce number; n393dte number; t391dte number; t392dce seconds; } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set Frame Relay keepalive parameters.
Options	<p>n391dte—DTE full status polling interval. Range: 1 through 255 Default: 6</p> <p>n392dce—DCE error threshold, in number of errors. Range: 1 through 10 Default: 3</p> <p>n392dte—DTE error threshold, in number of errors. Range: 1 through 10 Default: 3</p> <p>n393dce—DCE monitored event-count. Range: 1 through 10 Default: 4</p> <p>n393dte—DTE monitored event-count. Range: 1 through 10 Default: 4</p> <p>t391dte—DTE polling timer. Range: 5 through 30 seconds Default: 10 seconds</p> <p>t392dce—DCE polling timer. Range: 5 through 30 seconds Default: 15 seconds</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring Tunable Keepalives for Frame Relay LMI” on page 378.
Required Privilege Level	interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Related Topics lmi-type

lmi (Ethernet OAM)

Syntax lmi {
 status-counter *count*;
 polling-verification-timer *value*;
 interface *name*; {
 uni-id *uni-name*;
 status-counter *number*;
 polling-verification-timer *value*;
 evc-map-type (all-to-one-bundling | bundling | service-multiplexing);
 evc *evc-name* {
 default-evc;
 vlan-list *vlan-id-list*;
 }
 }
 }

Hierarchy Level [edit protocols oam ethernet]

Release Information Statement introduced in JUNOS Release 9.5.

Description On routers with *ge*, *xe*, or *ae* interfaces, configure an OAM Ethernet local management interface.

Options status-counter *count*—Status counter (N393), defaults to 4.
 interface *name*—Polling verification timer (T392), defaults to 15 seconds.
 uni-id *uni-name*—(Optional) Defaults to the physical interface name.
 status-counter *number*—(Optional) Defaults to a global value.
 polling-verification-timer *value*—(Optional) Defaults to a global value.
 evc-map-type (<all-to-one-bundling | bundling | service-multiplexing>)—Specify the Ethernet virtual connection (EVC) map type.
 evc *evc-name*—Specify the name of the EVC.
 default-evc—Set the specified EVC as the default EVC.
 vlan-list *vlan-id-list*—Specify a group of VLANs to assign to the EVC.

Usage Guidelines See “Configuring Ethernet Local Management Interface” on page 690.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics evcs

lmi-type

Syntax	lmi-type (ansi itu);
Hierarchy Level	[edit interfaces <i>interface-name</i> lmi (Frame Relay)], [edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set Frame Relay Local Management Interface (LMI) type.
Options	ansi—Use ANSI T1.167 Annex D LMIs. itu—Use ITU Q933 Annex A LMIs. Default: ansi
Usage Guidelines	See “Configuring Tunable Keepalives for Frame Relay LMI” on page 378 and the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

load-interval

Syntax	load-interval <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>dln</i> unit <i>logical-unit-number</i> dialer-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>dln</i> unit <i>logical-unit-number</i> dialer-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	On J Series Services Routers with ISDN logical interfaces, specify the interval used to calculate the average load on the network. By default, the average interface load is calculated every 60 seconds.
Options	<i>seconds</i> —Number of seconds at which the average load calculation is triggered. Range: 20 through 180, in 10-second intervals Default: 60 seconds
Usage Guidelines	See “Configuring ISDN Logical Interface Properties” on page 823.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

load-threshold

Syntax	<code>load-threshold percent;</code>
Hierarchy Level	[edit interfaces <i>dln</i> unit <i>logical-unit-number</i> dialer-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>dln</i> unit <i>logical-unit-number</i> dialer-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	On J Series Services Routers with ISDN logical interfaces, specify the bandwidth threshold percentage used for adding interfaces. Another link is added to the multilink bundle when the load reaches the threshold value you set. Specify a percentage between 0 and 100.
Options	<i>percent</i> —Bandwidth threshold percentage used for adding interfaces. When set to 0, all available channels are dialed. Range: 0 through 100 Default: 100 seconds
Usage Guidelines	See “Configuring Bandwidth on Demand” on page 829.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

local-mac

Syntax	<code>local-mac mac-address remote-mac mac-address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> map]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	For J Series Services Routers only. On Ethernet interfaces configured for DLSw, specify the local MAC address to be mapped to a remote destination MAC address.
Options	<i>mac-address</i> —MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i> . For example, 0011.2233.4455 or 00:11:22:33:44:55.
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 181.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Services Interfaces Configuration Guide</i>

local-name

Syntax	local-name <i>name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> ppp-options chap], [edit interfaces <i>interface-name</i> ppp-options pap], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options chap], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options pap], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options chap], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options pap]
Release Information	Statement introduced before JUNOS Release 7.4. Support for PAP added in JUNOS Release 8.3.
Description	<p>For CHAP authentication, the value sent in CHAP challenge and response packets on a per interface basis. For PAP authentication, the local hostname for sending PAP authentication requests.</p> <p>For ATM2 IQ interfaces only, you can configure a CHAP local name on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:</p> <ul style="list-style-type: none"> ■ atm-ppp-llc—PPP over AAL5 LLC encapsulation. ■ atm-ppp-vc-mux—PPP over AAL5 multiplex encapsulation.
Default	For CHAP authentication, if you do not include the local-name statement in the configuration, the interface sends the router's system hostname in CHAP challenge and response packets.
Usage Guidelines	See “Configuring the PPP Challenge Handshake Authentication Protocol” on page 112 and “Configuring the PPP Password Authentication Protocol” on page 114.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS System Basics Configuration Guide</i>

local-password

Syntax	<code>local-password password;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> ppp-options pap], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options pap], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options pap]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Configure the host password for sending PAP requests.
Usage Guidelines	See “Configuring the Local Password” on page 165.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	Configuring the PPP Password Authentication Protocol on page 114

local-window

Syntax	<code>local-window count;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series Services Routers only. On Ethernet interfaces configured for DLSw, configure the maximum number of Information-frames (I-frames) to send before waiting for acknowledgment.
Options	<i>count</i> —Number of I-frames. Range: 1 through 127 I-frames Default: 7 I-frames
Usage Guidelines	See “Configuring LLC2 Options” on page 180.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

lockout

Syntax	lockout;
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure a lockout of protection, forcing the use of the working circuit and locking out the protect circuit regardless of anything else.
Usage Guidelines	See “Configuring Switching Between the Working and Protect Circuits” on page 866.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

logical-interface-policer

Syntax	logical-interface-policer;
Hierarchy Level	[edit firewall policer <i>policer-template-name</i>], [edit firewall three-color-policer <i>policer-name</i>],
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Apply a policer to a logical interface in the ingress or egress direction as part of a configuration to discard high loss priority traffic, or configure an aggregate policer.
Usage Guidelines	See “Configuring Gigabit Ethernet Two-Color and Tricolor Policers” on page 763, See “Applying a Policer” on page 765, and the <i>JUNOS Class of Service Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	action (Policer)

logical-systems

Syntax	<code>logical-systems <i>logical-system-name</i>;</code>
Hierarchy Level	[edit]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure a logical system.
Options	<i>logical-system-name</i> —Name of the logical system.
Usage Guidelines	See “Configuring Logical System Interface Properties” on page 155.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

log-prefix

Syntax	<code>log-prefix <i>prefix-number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options syslog host <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the system logging prefix value.
Options	<i>prefix-number</i> —System logging prefix value.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

long-buildout

Syntax	(long-buildout no-long-buildout);
Hierarchy Level	[edit interfaces <i>interface-name</i> t3-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Configure the T3 line buildout. A T3 interface has two settings for the T3 line buildout: a short setting, which is less than 255 feet (68 meters), and a long setting, which is greater than 255 feet and shorter than 450 feet (137 meters).</p> <p>This statement applies to copper-cable-based T3 interfaces only. You cannot configure a line buildout for a DS3 channel on a channelized OC12 interface, which runs over fiber-optic cable.</p>
Default	A T3 interface uses the short line buildout setting (no-long-buildout) for wires shorter than 255 feet (68 meters).
Usage Guidelines	See “Configuring the T3 Line Buildout” on page 575.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

loopback

See the following sections:

- loopback (ADSL, DS0, E1/E3, SONET/SDH, SHDSL, and T1/T3) on page 1113
- loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet) on page 1114
- loopback (Serial) on page 1115

loopback (ADSL, DS0, E1/E3, SONET/SDH, SHDSL, and T1/T3)

Syntax loopback (local | payload | remote);

Hierarchy Level [edit interfaces *ce1-fpc/pic/port*],
 [edit interfaces *ct1-fpc/pic/port*],
 [edit interfaces *t1-fpc/pic/port*],
 [edit interfaces *interface-name* ds0-options],
 [edit interfaces *interface-name* dsl-options],
 [edit interfaces *interface-name* e1-options],
 [edit interfaces *interface-name* e3-options],
 [edit interfaces *interface-name* shdsl-options],
 [edit interfaces *interface-name* sonet-options],
 [edit interfaces *interface-name* t1-options],
 [edit interfaces *interface-name* t3-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a loopback connection. To turn off the loopback capability, remove the **loopback** statement from the configuration.



NOTE: When configuring CE1 or CT1 interfaces on 10-port Channelized E1/T1 IQE PICs, the **loopback** statement must be included with the **local** or **remote** option at the [edit interfaces *ce1-fpc/pic/port*] or [edit interfaces *ct1-fpc/pic/port*] hierarchy level as appropriate.

When configuring T1 interfaces on 10-port Channelized E1/T1 IQE PICs, the **loopback** statement must be included with the **payload** option at the [edit interfaces *t1-fpc/pic/port*] hierarchy level.

Options **local**—Loop packets, including both data and timing information, back on the local router's PIC. NxDS0 IQ interfaces do not support local loopback.

payload—For channelized T3, T1, and NxDS0 IQ interfaces only, loop back data only (without clocking information) on the remote router's PIC. With payload loopback, overhead is recalculated. Neither ATM-over-asymmetrical digital subscriber line (ADSL) interfaces nor ATM-over-SHDSL interfaces support payload loopback.

remote—Loop packets, including both data and timing information, back on the remote router's interface card. NxDS0 IQ interfaces do not support remote loopback. ATM-over-ADSL interfaces do not support payload loopback.

Usage Guidelines See “Configuring E3 and T3 Parameters on ATM Interfaces” on page 337, “Configuring E1 Loopback Capability” on page 547, “Configuring E3 Loopback Capability” on page 555, “Configuring Channelized IQ and IQE SONET/SDH Loop Timing” on page 852, “Configuring Channelized IQ and IQE SONET/SDH Loop Timing” on page 852, “Configuring SHDSL Operating Mode on an ATM Physical Interface” on page 364, “Configuring T1 Loopback Capability” on page 565, and “Configuring T3 Loopback Capability” on page 576.

To configure loopback on channelized IQ and IQE PICs, SONET/SDH level, use the `sonet-options loopback` statement `local` and `remote` options at the controller interface (`coc48`, `cstm16`, `coc12`, `cstm4`, `coc3`, `cstm1`). It is ignored for path-level interfaces `so-fpc/pic/port` or `so-fpc/pic/port:channel`.

Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	feac-loop-respond

loopback (Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet)

Syntax	(loopback no-loopback);
Hierarchy Level	[edit interfaces <i>interface-name</i> aggregated-ether-options], [edit interfaces <i>interface-name</i> fastether-options], [edit interfaces <i>interface-name</i> ggether-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces, enable or disable loopback mode.



NOTE: By default, local aggregated Ethernet, Fast Ethernet, Tri-Rate Ethernet copper, Gigabit Ethernet, and 10-Gigabit Ethernet interfaces connect to a remote system.

Usage Guidelines	See “Configuring Ethernet Loopback Capability” on page 593.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

loopback (Serial)

Syntax	<code>loopback mode;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure a loopback connection.
Default	If you do not include this statement, there is no loopback connection.
Options	<i>mode</i> —You can specify the one of the following loopback modes: <ul style="list-style-type: none"> ■ <i>dce-local</i>—For EIA-530 interfaces only, loop packets back on the local DCE. ■ <i>dce-remote</i>—For EIA-530 interfaces only, loop packets back on the remote DCE. ■ <i>local</i>—Loop packets back on the local router's PIC. ■ <i>remote</i>—Loop packets back on the line interface unit (LIU).
Usage Guidelines	See “Configuring Serial Loopback Capability” on page 275.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

loopback-clear-timer

Syntax	<code>loopback-clear-timer seconds;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	For interfaces with PPP, PPP TCC, PPP over Ethernet, PPP over ATM, and PPP over Frame Relay encapsulations, configure a loop detection clear timer for the Link Control Protocol (LCP) component of a PPP session.
Options	<i>seconds</i> —The time in seconds to wait before the loop detection flag is cleared if it is not cleared by the protocol. Range: 1 through 60 seconds. Default: 9 seconds
Usage Guidelines	See “Configuring the PPP Clear Loop Detected Timer” on page 162.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

loop-timing

Syntax	(loop-timing no-loop-timing);
Hierarchy Level	[edit interfaces <i>ct3-fpc/pic/port</i> t3-options], [edit interfaces <i>e1-fpc/pic/port:0</i> sonet-options], [edit interfaces <i>stm1-fpc/pic/port</i> sonet-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For channelized IQ interfaces and non-IQ channelized STM1 interfaces only, configure the SONET/SDH or DS3-level clocking source.
Options	loop-timing—Configure loop timing (external) clocking. no-loop-timing—Configure line timing (internal) clocking. Default: no-loop-timing
Usage Guidelines	See “Configuring Channelized IQ and IQE SONET/SDH Loop Timing” on page 852 and “Configuring the Channelized T3 Loop Timing” on page 576.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	clocking

loss-priority

Syntax	loss-priority (high low);
Hierarchy Level	[edit interfaces <i>interface-name</i> gigheter-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier premium forwarding-class <i>class-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the packet loss priority value.
Options	high—Packet has high loss priority. low—Packet has low loss priority.
Usage Guidelines	See “Specifying an Output Priority Map” on page 759.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

loss-threshold

Syntax	loss-threshold <i>number</i> ;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> maintenance-association <i>ma-name</i> continuity-check]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	The number of continuity check messages lost before marking the remote MEP as down.
Options	<i>number</i> —Specify how many continuity check messages can be lost before the remote MEP is considered down.
Usage Guidelines	See “Configuring the Continuity Check Loss Threshold” on page 686.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

low-plp-max-threshold

Syntax	low-plp-max-threshold <i>percent</i> ;
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> atm-options linear-red-profiles <i>profile-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, define the drop profile fill-level for the low PLP CoS VC. When the fill level exceeds the defined percentage, all packets are dropped.
Options	<i>percent</i> —Fill-level percentage when linear RED is applied to cells with PLP.
Usage Guidelines	See “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	high-plp-max-threshold, low-plp-threshold, queue-depth

low-plp-threshold

Syntax	low-plp-threshold <i>percent</i> ;
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> atm-options linear-red-profiles <i>profile-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, define the CoS VC drop profile fill-level percentage when linear RED is applied to cells with low PLP. When the fill level exceeds the defined percentage, packets with low PLP are randomly dropped by RED. This statement is mandatory.
Options	<i>percent</i> —Fill-level percentage when linear RED is applied to cells with low PLP.
Usage Guidelines	See “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	high-plp-max-threshold, high-plp-threshold, low-plp-max-threshold, queue-depth

lsq-failure-options

Syntax	lsq-failure-options { no-termination-request; [trigger-link-failure <i>interface-name</i>]; }
Hierarchy Level	[edit interfaces <i>lsq-fpc/pic/port</i>]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For AS PIC or MultiServices PIC link services IQ (lsq) interfaces only, define the failure recovery option settings.
Options	The remaining statements are explained separately.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mac

Syntax	<code>mac mac-address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Set the MAC address of the interface. You can configure the MAC address on the management Ethernet interface (fxp0 or em0) only.
Options	<i>mac-address</i> —MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i> . For example, 0011.2233.4455 or 00:11:22:33:44:55.
Usage Guidelines	See “Configuring the MAC Address on the Management Ethernet Interface” on page 777.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mac-address

Syntax	<code>mac-address mac-address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> accept-source-mac], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> accept-source-mac]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), specify a remote MAC address on which to count incoming and outgoing packets.
Options	<i>mac-address</i> —MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i> . For example, 0011.2233.4455 or 00:11:22:33:44:55.
Usage Guidelines	See “Configuring MAC Address Filtering” on page 761.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mac-learn-enable

Syntax	(mac-learn-enable no-mac-learn-enable);
Hierarchy Level	[edit interfaces <i>interface-name</i> gigether-options ethernet-switch-profile], [edit interfaces <i>interface-name</i> aggregated-ether-options ethernet-switch-profile]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), configure whether source and destination MAC addresses are dynamically learned:</p> <ul style="list-style-type: none"> ■ mac-learn-enable—Allow the interface to dynamically learn source and destination MAC addresses. ■ no-mac-learn-enable—Prohibit the interface from dynamically learning source and destination MAC addresses.

MAC address learning is based on source addresses. You can start accounting for traffic after there has been traffic sent from the MAC address. Once the MAC address is learned, the frames and bytes transmitted to or received from the MAC address can be tracked.



NOTE: When you gather interfaces into a bridge domain, the **no-mac-learn-enable** statement at the [edit interfaces *interface-name* gigether-options ethernet-switch-profile] hierarchy level is not supported. You must use the **no-mac-learning** statement at the [edit bridge-domains *bridge-domain-name* bridge-options interface *interface-name*] hierarchy level to disable MAC learning on an interface in a bridge domain. For information on disabling MAC learning for a bridge domain, see *MX Series Layer 2 Configuration Guide*.

Usage Guidelines	See “Configuring MAC Address Filtering” on page 761.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mac-validate

Syntax	mac-validate (loose strict);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Enable IP and MAC address validation for static Ethernet and IP demux interfaces. Supported on MX Series routers only.
Options	<p>loose—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the IP source address matches one of the trusted tuples, but the MAC address does not match the MAC address of the tuple. Continues to forward incoming packets when the source address of the incoming packet does not match any of the trusted IP addresses.</p> <p>strict—Forwards incoming packets when both the IP source address and the MAC source address match one of the trusted address tuples. Drops packets when the MAC address does not match the tuple's MAC source address, or when IP source address of the incoming packet does not match any of the trusted IP addresses.</p>
Usage Guidelines	See “Configuring MAC Address Validation on Static Ethernet Interfaces” on page 675 and “Configuring MAC Address Validation on Static Demux Interfaces” on page 254.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

maintenance-association

Syntax maintenance-association *ma-name* {
 short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
 continuity-check {
 hold-interval (OAM) *minutes*;
 interval (10m | 10s | 1m | 1s| 100ms);
 loss-threshold *number*;
 }
 mep *mep-id* {
 auto-discovery;
 direction (up | down);
 interface *interface-name*;
 priority *number*;
 remote-mep *mep-id* {
 action-profile *profile-name*;
 }
 }
 }
 }

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain
 domain-name]

Release Information Statement introduced in JUNOS Release 8.4.

Description Configure the name of the maintenance association in IEEE-compliant format.

Options ma-name—The name of the maintenance association within the maintenance domain.
 The remaining statements are explained separately.

Usage Guidelines See “Creating the Maintenance Association” on page 684.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

maintenance-domain

Syntax `maintenance-domain domain-name {`
 `bridge-domain name;`
 `instance vpls-instance;`
 `level number;`
 `mip-half-function (none | default | explicit);`
 `name-format (character-string | none | dns | mac+2oct);`
 `maintenance-association ma-name {`
 `short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);`
 `continuity-check {`
 `hold-interval (OAM) minutes;`
 `interval (10m | 10s | 1m | 1s| 100ms);`
 `loss-threshold number`
 `}`
 `mep mep-id {`
 `action-profile profile-name;`
 `auto-discovery;`
 `direction (up | down);`
 `interface interface-name;`
 `priority number;`
 `remote-mep mep-id {`
 `action-profile profile-name;`
 `}`
 `}`
 `mip-half-function(none | default | explicit);`
 `}`
 `routing-instance name {`
 `bridge-domain name;`
 `}`
 `virtual-switch name bridge-domain name;`
 `}`

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management]

Release Information Statement introduced in JUNOS Release 8.4.

Description Configure the name of the maintenance domain in IEEE-compliant format.

Options `domain-name`—The name for the maintenance domain.

The remaining statements are explained separately.

Usage Guidelines See “Creating the Maintenance Domain” on page 682.

Required Privilege Level `interface`—To view this statement in the configuration.
 `interface-control`—To add this statement to the configuration.

map

Syntax map {
 local-mac *mac-address* remote-mac *mac-address*;
 }

Hierarchy Level [edit interfaces interface-name unit logical-unit-number family llc2 redundancy-group group-number],
 [edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family llc2 redundancy-group group-number]

Release Information Statement introduced in JUNOS Release 7.5.

Description For J Series Services Routers only. On Ethernet interfaces configured for DLSw Ethernet redundancy, map a local peer MAC address to a remote peer MAC address.

The statements are explained separately



NOTE: For DLSw configurations, you must specify the **local-mac** statement and the **remote-mac** statement when you configure the **map** statement. If you do not include the MAC translation information with the **map** statement, the commit operation will fail.

Usage Guidelines See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 181.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics *JUNOS Services Interfaces Configuration Guide*

master-only

Syntax	master-only;
Hierarchy Level	[edit groups rex interfaces (fxp0 em0) unit <i>logical-unit-number</i> family <i>family</i> address], [edit groups rex logical-systems <i>logical-system-name</i> interfaces fxp0 unit <i>logical-unit-number</i> family <i>family</i> address] [edit interfaces (fxp0 em0) unit <i>logical-unit-number</i> family <i>family</i> address], [edit logical-systems <i>logical-system-name</i> interfaces fxp0 unit <i>logical-unit-number</i> family <i>family</i> address]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the IP address to be used when the Routing Engine is the current master.
Usage Guidelines	See “Configuring a Consistent Management IP Address” on page 775.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	For information about the groups statement, see the <i>JUNOS CLI User Guide</i> .

maximum-contexts

Syntax	maximum-contexts <i>number</i> <force>;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression rtp], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> compression rtp]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Specify the maximum number of RTP contexts to accept during negotiation.
Options	<i>number</i> —Maximum number of contexts. <i>force</i> —(Optional) Requires the PIC to use the value specified for maximum RTP contexts, regardless of the negotiated value. This option allows the software to interoperate with JUNOS releases that base the RTP context value on link speed.
Usage Guidelines	See <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

maximum-vcs

Syntax	<code>maximum-vcs <i>maximum-vcs</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>at-fpc/pic/port</i> atm-options vpi <i>vpi-identifier</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For ATM1 interfaces, configure the maximum number of virtual circuits (VCs) allowed on a virtual path (VP). When configuring ATM1 interfaces on the router, you must include this statement.</p> <p>For a configured virtual path identifier (VPI), valid virtual channel identifier (VCI) numbers are from 0 through (<i>maximum-vcs</i> value – 1). VCI numbers 0 through 31 are reserved by the ATM Forum. It is recommended that you use a VCI number higher than 31 when connecting to an ATM switch.</p>
Options	<p><i>maximum-vcs</i>—Maximum number of VCs on the VP.</p> <p>Range: 1 through 4090</p>
Usage Guidelines	See “Configuring the Maximum Number of ATM1 VCs on a VP” on page 300.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	multipoint-destination, promiscuous-mode, vci

maximum-requests

Syntax	<code>maximum-requests <i>times</i>;</code>
Hierarchy Level	<code>[edit protocols dot1x authenticator interface <i>interface-id</i>]</code>
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specify the maximum number of retransmission times of an EAPOL Request packet to the client before it times out the authentication session.
Options	<p><i>times</i>—Specify the maximum number of retransmission times.</p> <p>Range: 1 through 10 times</p> <p>Default: 2 times</p>
Usage Guidelines	See “Configuring IEEE 802.1x Port-Based Network Access Control” on page 741.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	dot1x, authenticator, interface (IEEE 802.1x)

max-retry

Syntax	<code>max-retry count;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series Services Routers only. On Ethernet interfaces configured for DLSw, configure the number of retries the router should attempt when waiting for a response.
Options	<i>count</i> —Number of retries. Range: 1 through 127 Default: 10
Usage Guidelines	See “Configuring LLC2 Options” on page 180.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

member-interface-speed

Syntax	<code>member-interface-speed speed;</code>
Hierarchy Level	[edit interfaces container-options member-interface-type]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Specify container-interface member-interface speed options.
Options	<i>speed</i> —Set interface speed to OC3, OC12, OC48, OC192, OC768, or mixed.
Usage Guidelines	See “Configuring Container Interfaces” on page 863.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	container-options

member-interface-type

Syntax	<pre>member-interface-type sonet { member-interface-speed [<i>speed</i>]; }</pre>
Hierarchy Level	[edit interfaces container-options]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Specify container-interface member-interface type as sonet and speed options.
Options	<p>sonet—Protocol type of the container interface, specify sonet.</p> <p>speed—Set interface speed to OC3, OC12, OC48, OC192, OC768, or mixed.</p>
Usage Guidelines	See “Configuring Container Interfaces” on page 863.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	container-options

mep

Syntax `mep mep-id {
 auto-discovery;
 direction (up | down);
 interface interface-name;
 priority number;
 remote-mep mep-id {
 action-profile profile-name;
 }
 }`

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain
 domain-name maintenance-association *ma-name*]

Release Information Statement introduced in JUNOS Release 8.4.

Description The numeric identifier of the maintenance association end point (MEP) within the maintenance association.

Options mep-id—Specify the numeric identifier of the MEP.
 Range: 1 through 8191

The remaining statements are explained separately.

Usage Guidelines See “Configuring a Maintenance End Point” on page 686.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

minimum-links

Syntax	minimum-links <i>number</i> ;
Hierarchy Level	[edit interfaces aex aggregated-ether-options], [edit interfaces asx aggregated-sonet-options], [edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For aggregated Ethernet, SONET/SDH, multilink, link services, and voice services interfaces only, set the minimum number of links that must be up for the bundle to be labeled up.
Options	<i>number</i> —Number of links. Range: 1 through 8 (1 through 16 for Ethernet and SONET interfaces on the MX Series, M320, M120, T Series or TX Matrix routers) Default: 1
Usage Guidelines	See “Configuring Aggregated Ethernet Minimum Links” on page 635, “Configuring Aggregated SONET/SDH Minimum Links” on page 883, or the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mip-half-function

Syntax mip-half-function (none | default | explicit);

Hierarchy Level [edit protocols oam ethernet connectivity-fault-management maintenance-domain
 md-name]

Description Specify the OAM Ethernet CFM maintenance domain MIP half functions.



NOTE: Whenever a MIP is configured and a bridge domain is mapped to multiple maintenance domains (MD) or maintenance associations (MA), it is essential that the **mip-half-function** value for all MDs and MAs are the same.

Options none—Specify to not use the mip-half-function.

 default—Specify to use the default mip-half-function.

 explicit—Specify an explicit mip-half-function.

Usage Guidelines See “Creating the Maintenance Domain” on page 682.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics See maintenance-domain.

mlfr-uni-nni-bundle-options

Syntax mlfr-uni-nni-bundle-options {
 acknowledge-retries *number*;
 acknowledge-timer *milliseconds*;
 action-red-differential-delay (disable-tx | remove-link);
 drop-timeout *milliseconds*;
 fragment-threshold *bytes*;
 hello-timer *milliseconds*;
 link-layer-overhead *percent*;
 lmi-type (ansi | itu);
 minimum-links *number*;
 mrru *bytes*;
 n391 *number*;
 n392 *number*;
 n393 *number*;
 red-differential-delay *milliseconds*;
 t391 *seconds*;
 t392 *number*;
 yellow-differential-delay *milliseconds*;
 }

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure link services and voice services interface management properties.

The statements are explained separately.

Usage Guidelines See the *JUNOS Services Interfaces Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

mode

Syntax	mode loose;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) rpf-check], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family (inet inet6) rpf-check]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Check whether the packet has a source address with a corresponding prefix in the routing table. If a corresponding prefix is not found, unicast reverse path forwarding (RPF) loose mode does not accept the packet. Unlike strict mode, loose mode does not check whether the interface expects to receive a packet with a specific source address prefix.
Default	If you do not include this statement, unicast RPF is in strict mode.
Usage Guidelines	See “Configuring Unicast RPF Strict Mode” on page 209 and “Configuring Unicast RPF Loose Mode” on page 210.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

modem-options

Syntax	modem-options { dialin (console routable); init-command-string <i>initialization-command-string</i> ; }
Hierarchy Level	[edit interfaces umd0]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	For J Series Services Routers, configure a USB port to act as a USB modem. The remaining statement is explained separately.
Usage Guidelines	See “Specifying a USB Modem Interface on J Series Routers” on page 93.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

monitor-session

Syntax	monitor-session (<i>interface-name</i> all);
Hierarchy Level	[edit protocols ppp]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	Monitor PPP packet exchanges. When monitoring is enabled, packets exchanged during a session are logged to the default log of <code>/var/log/pppd</code> .
Default	If you do not include this statement, no PPPD-specific monitoring operations are performed.
Options	all—Monitor PPP packet exchanges on all sessions. <i>interface-name</i> —Logical interface name on which to enable session monitoring.
Usage Guidelines	See “Monitoring a PPP Session” on page 118.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

mpls

Syntax	mpls { pop-all-labels { required-depth <i>number</i> ; } }
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options], [edit interfaces <i>interface-name</i> sonet-options], [edit interfaces <i>interface-name</i> fastether-options], [edit interfaces <i>interface-name</i> gigether-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For passive monitoring on ATM and SONET/SDH interfaces and 10-Gigabit Ethernet interfaces in WAN PHY mode, process incoming IP packets that have MPLS labels. The remaining statements are explained separately.
Usage Guidelines	See “Removing MPLS Labels from Incoming Packets” on page 294 and “Removing MPLS Labels from Incoming Packets” on page 874.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Services Interfaces Configuration Guide</i>

mrru

Syntax	<code>mrru bytes;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For multilink, link services, voice services, and J Series Services Routers ISDN interfaces only, set the maximum received reconstructed unit (MRRU). The MRRU is similar to the MTU, but is specific to multilink interfaces.
Options	<i>bytes</i> —MRRU size. Range: 1500 through 4500 bytes Default: 1500 bytes
Usage Guidelines	See “Configuring the Dialer Interface” on page 830 and the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	mtu

mtu

Syntax `mtu bytes;`**Hierarchy Level** `[edit interfaces interface-name],`
`[edit interfaces interface-name unit logical-unit-number family family],`
`[edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number`
`family family]`**Release Information** Statement introduced before JUNOS Release 7.4.**Description** Maximum transmission unit (MTU) size for the media or protocol. The default MTU size depends on the device type. Changing the media MTU or protocol MTU causes an interface to be deleted and added again.**NOTE:** Not all devices allow you to set an MTU value, and some devices have restrictions on the range of allowable MTU values. You cannot configure an MTU for management Ethernet (**fxp0** or **em0**) interfaces or for loopback, multilink, and multicast tunnel devices.**Options** `bytes`—MTU size.**Range:** 0 through 9192 bytes**Default:** 1500 bytes (INET, INET6, and ISO families), 1448 bytes (MPLS)**Usage Guidelines** See “Configuring the Media MTU” on page 98 and “Setting the Protocol MTU” on page 191.**Required Privilege Level** `interface`—To view this statement in the configuration.
`interface-control`—To add this statement to the configuration.

multicast-dlci

Syntax	<code>multicast-dlci <i>dlci-identifier</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For point-to-multipoint Frame Relay, link services, and voice services interfaces only, enable multicast support on the interface. You can configure multicast support on the interface if the Frame Relay switch performs multicast replication.
Options	<i>dlci-identifier</i> —DLCI identifier, a number from 16 through 1022 that defines the Frame Relay DLCI over which the switch expects to receive multicast packets for replication.
Usage Guidelines	See “Configuring a Multicast-Capable Frame Relay Connection” on page 381 and the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	dlci, multipoint-destination

multicast-vci

Syntax	<code>multicast-vci vpi-identifier.vci-identifier;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM encapsulation only, and for point-to-multipoint ATM logical interfaces only, enable the support of multicast on the interface. You can configure multicast support on the interface if the ATM switch performs multicast replication.
Options	<i>vci-identifier</i> —ATM virtual circuit identifier. Range: 0 through 16384 <i>vpi-identifier</i> —ATM virtual path identifier. Range: 0 through 255 Default: 0
Usage Guidelines	See “Configuring a Multicast-Capable ATM1 or ATM2 IQ Connection” on page 318.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	multipoint-destination, vci

multicast-only

Syntax	<code>multicast-only;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the unit and family so that it can transmit and receive multicast traffic only. You can configure this property on the IP family only.
Usage Guidelines	See “Configuring the Protocol Family” on page 172.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	tunnel, <i>JUNOS Services Interfaces Configuration Guide</i>

multilink-max-classes

Syntax	multilink-max-classes <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Adaptive Services (AS) PIC link services IQ interfaces (lsq) only, configure the number of multilink classes to be negotiated when a link joins the bundle.
Options	<i>number</i> —The number of multilink classes to be negotiated when a link joins the bundle. Range: 1 through 8 Default: None
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	multipoint

multipoint

Syntax	multipoint;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the interface unit as a multipoint connection.
Default	If you omit this statement, the interface unit is configured as a point-to-point connection.
Usage Guidelines	See “Configuring a Multipoint Connection” on page 157.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	point-to-point

multipoint-destination

Syntax	<pre> multipoint-destination address dlcidlcid-identifier; multipoint-destination address { epd-threshold cells; inverse-arp; oam-liveness { up-count cells; down-count cells; } oam-period (disable seconds); shaping { (cbr rate rtvbr peak rate sustained rate burst length vbr peak rate sustained rate burst length); queue-length number; } vci vpi-identifier.vci-identifier; } </pre>
Hierarchy Level	<pre> [edit interfaces interface-name unit logical-unit-number family family address address], [edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number family family address address] </pre>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For point-to-multipoint Frame Relay or ATM interfaces only, enable the support of multicast on the interface. You can configure multicast support on the interface if the Frame Relay or ATM switch performs multicast replication.
Options	<p>address—Address of the remote side of the point-to-multipoint connection.</p> <p>dlci-identifier—For Frame Relay interfaces, the data-link connection identifier. Range: 0 through 0xFFFFFFF (24 bits)</p> <p>vci-identifier—For ATM interfaces, the virtual circuit identifier. Range: 0 through 16384</p> <p>vpi-identifier—For ATM interfaces, the virtual path identifier. Range: 0 through 255 Default: 0</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring a Point-to-Point ATM1 or ATM2 IQ Connection” on page 316, and “Configuring a Point-to-Multipoint Frame Relay Connection” on page 381.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	dlci, encapsulation

multiservice-options

Syntax	multiservice-options { (syslog no-syslog); (core-dump no-core-dump); (dump-on-flow-control); }
Hierarchy Level	[edit interfaces <i>mo-fpc/pic/port</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For monitoring services interfaces only, configure multiservice-specific interface properties. The statements are explained separately.
Usage Guidelines	See “Configuring Multiservice Physical Interface Properties” on page 138 and the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	passive-monitor-mode

n391

Syntax	n391 <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services and voice services interfaces only, set the Frame Relay full status polling interval.
Options	number—Number of polling interval. Range: 1 through 255 Default: 6
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	n392, n393, timeslots, and t392

n392

Syntax	n392 <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services and voices interfaces only, set the Frame Relay error threshold, in number of errors.
Options	number—Error threshold. Range: 1 through 10 Default: 3
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	n391, n393, timeslots, t392

n393

Syntax	n393 <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services and voices interfaces only, set the Frame Relay monitored event count.
Options	<i>number</i> —Number of event count. Range: 1 through 255 Default: 6
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	n391, n392, timeslots, t392

name-format

Syntax	name-format (character-string none dns mac+2oct);
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Specify the format of the maintenance domain name.
Options	<p>character-string—The name is an ASCII character string.</p> <p>none—Name format none means that maintenance domain name is not used.</p> <p>dns—Name is in domain name service (DNS) format. For example: www.juniper.net.</p> <p>mac+2octet—Name is the MAC address plus a two-octet maintenance association identifier. For example: 08:00:22:33:44:55.100</p> <p>Default: character-string</p>
Usage Guidelines	See “Configuring the Maintenance Association Short Name Format” on page 685.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

native-vlan-id

Syntax	<code>native-vlan-id <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>ge-fpc/pic/port</i>]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	<p>For 1-, 4-, and 8-port Gigabit Ethernet IQ2 and IQ2-E PICs, for 1-port 10-Gigabit Ethernet IQ2 and IQ2-E PICs configured for 802.1Q flexible VLAN tagging, for all Ethernet interfaces on MX Series routers, and for aggregated Ethernet interfaces on IQ2 and IQ2-E PICs or MX Series DPCs, configure mixed tagging support for untagged packets on a port. When the native-vlan-id statement is included with the flexible-vlan-tagging statement, untagged packets are accepted on the same mixed VLAN-tagged port.</p> <p>The logical interface on which untagged packets will be received must be configured with the same native VLAN ID as that configured on the physical interface. To configure the logical interface, include the vlan-id statement (matching the native-vlan-id statement on the physical interface) at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.</p> <p>When the native-vlan-id statement is included with the interface-mode the statement, untagged packets are accepted and forwarded within the bridge domain that is configured with the matching VLAN ID.</p>
Options	<i>number</i> —VLAN ID number.
Usage Guidelines	See “Configuring Mixed Tagging Support for Untagged Packets” on page 603, flexible-vlan-tagging , and “Configuring a Logical Interface for Access Mode” on page 619.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ncp-max-conf-req

Syntax	<code>ncp-max-conf-req <i>number</i></code>
Hierarchy Level	[edit interfaces <i>so-fpc/pic/port</i> unit <i>number</i> ppp-options]
Release Information	Statement introduced in JUNOS Release 9.6.
Description	Set the maximum number of NCP Configure-Requests to be sent, after which the router goes to NCP down state.
Options	<p>Range:</p> <p><i>number</i>—From 0 to 65535, where 0 means send infinite NCP Configure-Requests and any other value specifies the maximum number NCP Configure-Requests to send and then stop sending.</p>
Usage Guidelines	See “Configuring the NCP Configure-Request Maximum Sent” on page 161.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	ppp-options

ncp-restart-timer

Syntax	<code>ncp-restart-timer <i>milliseconds</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	For interfaces with PPP and PPP TCC encapsulations and on multilink PPP bundle interfaces, configure a restart timer for the Network Control Protocol (NCP) component of a PPP session.
Options	<p><i>milliseconds</i>—The time in milliseconds between successive NCP configuration requests.</p> <p>Range: 500 through 10000 milliseconds.</p> <p>Default: 3 seconds</p>
Usage Guidelines	See “Configuring the PPP Restart Timers” on page 162.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

negotiate-address

Syntax	negotiate-address;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For interfaces with PPP encapsulation, enable the interface to be assigned an IP address by the remote end.
Usage Guidelines	See “Configuring IPCP Options” on page 177.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	address, unnumbered-address, <i>JUNOS System Basics Configuration Guide</i>

negotiation-options

Syntax	negotiation-options { allow-remote-loopback; no-allow-link-events; }
Hierarchy Level	[edit protocols oam link-fault-management interface <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Enable and disable IEEE 802.3ah Operation, Administration, and Management (OAM) features for Ethernet interfaces. The statements are explained separately.
Usage Guidelines	See “Configuring IEEE 802.3ah OAM Link-Fault Management” on page 745.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

neighbor

Syntax	neighbor <i>address</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>If you are configuring one router to be the working router and a second to be the protect router, configure the address of the remote interface. You configure this on one or both of the interfaces.</p> <p>The address you specify for the neighbor must never be routed through the interface on which APS is configured, or instability will result. We strongly recommend that you directly connect the working and protect routers and that you configure the interface address of this shared network as the neighbor address.</p>
Options	<i>address</i> —Neighbor's address.
Usage Guidelines	See "Configuring Basic APS Support" on page 861.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

no-allow-link-events

Syntax	no-allow-link-events;
Hierarchy Level	[edit protocols oam ethernet link-fault-management interface <i>interface-name</i> negotiation-options]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Disable the sending of link event TLVs.
Usage Guidelines	See "Disabling the Sending of Link Event TLVs" on page 748.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

no-asynchronous-notification

See asynchronous-notification

no-auto-mdix

See speed (MX Series DPC).

no-auto-negotiation

See auto-negotiation

no-cbit-parity

See cbit-parity

no-core-dump

See core-dump

no-feac-loop-respond

See feac-loop-respond

no-flow-control

See flow-control

no-gratuitous-arp-reply

See gratuitous-arp-reply

no-gratuitous-arp-request

Syntax	no-gratuitous-arp-request;
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 9.6 for EX Series switches.
Description	For Ethernet interfaces, do not respond to gratuitous ARP requests.
Default	Gratuitous ARP responses are enabled on all Ethernet interfaces.
Usage Guidelines	See “Configuring Gratuitous ARP” on page 596.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	gratuitous-arp-reply

no-keepalives

Syntax	no-keepalives;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Disable the sending of keepalives on a physical interface configured with PPP, Frame Relay, or Cisco HDLC encapsulation. The default keepalive interval is 10 seconds.</p> <p>For ATM2 IQ interfaces only, you can disable keepalives on a logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:</p> <ul style="list-style-type: none"> ■ atm-ppp-llc—PPP over AAL5 LLC encapsulation. ■ atm-ppp-vc-mux—PPP over AAL5 multiplex encapsulation.
Usage Guidelines	See “Configuring Keepalives” on page 126 “Disabling the Sending of PPPoE Keepalive Messages” on page 796, and “Configuring Frame Relay Keepalives” on page 378.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

no-long-buildout

See long-buildout

no-loopback

See loopback

no-mac-learn-enable

See mac-learn-enable

node-id

Syntax node-id *mac-address*;

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in JUNOS Release 9.4.

Description Optionally specify the MAC address of a node in the protection group. If this statement is not included, the router assigns the node's MAC address by default.

Usage Guidelines See “Configuring Ethernet Ring Protection Switching” on page 799.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

non-revertive

Syntax non-revertive;

Hierarchy Level [edit interfaces aeX aggregated-ether-options lacp link-protection]

Release Information Statement introduced in JUNOS Release 9.3.

Description Disable the ability to switch to a better priority link (if one is available) once a link is established as active and collection distribution is enabled.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

no-partition

See the following sections:

- no-partition (Channelized E1 IQ Interfaces) on page 1151
- no-partition (Channelized OC1 IQ Interfaces) on page 1152
- no-partition (Channelized OC12 IQ Interfaces) on page 1152
- no-partition (Channelized STM1 IQ Interfaces) on page 1153
- no-partition (Channelized T3 IQ Interfaces) on page 1153

no-partition (Channelized E1 IQ Interfaces)

Syntax	no-partition interface-type e1;
Hierarchy Level	[edit interfaces ce1-fpc/pic/port]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Channelized E1 IQ PICs only, configure the channelized E1 interface as an unpartitioned, clear channel.
Default	If you do not include either this statement or the partition statement, the Channelized IQ PIC is not partitioned, and no data channels are configured.
Usage Guidelines	See “Configuring Channelized E1 Interfaces” on page 501.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	partition

no-partition (Channelized OC1 IQ Interfaces)

Syntax	no-partition interface-type (ct3 t3);
Hierarchy Level	[edit interfaces coc1-fpc/pic/port:channel]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For the Channelized OC12 PIC only, convert the channelized OC1 IQ interface into a channelized T3 interface or a T3 interface. You perform this configuration task for C-bit parity and M13-mapped configurations.
Default	If you do not include either this statement or the partition statement, the Channelized IQ PICs not partitioned, and no data channels are configured.
Options	ct3—Channelized T3 interface type. t3—T3 interface type.
Usage Guidelines	See “Configuring Channelized OC12/STM4 Interfaces” on page 423.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	partition

no-partition (Channelized OC12 IQ Interfaces)

Syntax	no-partition interface-type so;
Hierarchy Level	[edit interfaces coc12-fpc/pic/port]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Channelized OC12 IQ PICs only, configure the channelized OC12 interface as an unpartitioned, clear channel.
Default	If you do not include either this statement or the partition statement, the Channelized IQ PICs not partitioned, and no data channels are configured.
Usage Guidelines	See “Configuring an OC12/STM4 Interface” on page 427.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	partition

no-partition (Channelized STM1 IQ Interfaces)

Syntax	no-partition interface-type (cau4 so);
Hierarchy Level	[edit interfaces cstm1-fpc/pic/port]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For the Channelized STM1 PIC only, convert the channelized STM1 IQ interface into a channelized Administrative Unit 4 (AU-4) interface or a SONET/SDH STM1 interface.
Default	If you do not include either this statement or the partition statement, the Channelized IQ PICs not partitioned, and no data channels are configured.
Options	cau4—Channelized AU-4 interface type. so—SONET/SDH STM1 interface type.
Usage Guidelines	See “Configuring Channelized STM1 IQ and IQE Interfaces” on page 465.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	partition

no-partition (Channelized T3 IQ Interfaces)

Syntax	no-partition interface-type t3;
Hierarchy Level	[edit interfaces ct3-fpc/pic/port]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For the Channelized DS3 PIC only, configure the channelized T3 interface as an unpartitioned, clear channel.
Default	If you do not include either this statement or the partition statement, the Channelized IQ PIC is not partitioned, and no data channels are configured.
Usage Guidelines	See “Configuring T3 IQ Interfaces” on page 479.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	partition

no-payload-scrambler

See payload-scrambler

no-preempt

See preempt

no-redirects

Syntax no-redirects;

Hierarchy Level [edit interfaces *interface-name* unit *number* family *family*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Do not send protocol redirect messages on the interface.

To disable the sending of protocol redirect messages for the entire router, include the `no-redirects` statement at the [edit system] hierarchy level.

Default Interfaces send protocol redirect messages.

Usage Guidelines See “Disabling the Transmission of Redirect Messages on an Interface” on page 192.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics *JUNOS System Basics Configuration Guide*

no-source-filtering

See source-filtering

no-syslog

See syslog

no-termination-request

Syntax	no-termination-request;
Hierarchy Level	[edit interfaces <i>interface-name</i> ppp-options], [edit interfaces <i>lsq-fpc/pic/port</i> lsq-failure-options]
Release Information	Statement introduced in JUNOS Release 7.4. Support at the [edit interfaces <i>interface-name</i> ppp-options] hierarchy level added in JUNOS Release 8.3.
Description	For LSQ PICs or link PICs in redundant LSQ configurations, you can inhibit the router from sending PPP termination-request messages to the remote host if the PIC fails.
Usage Guidelines	See “Configuring Link PIC Failover on Channelized OC3 IQ and IQE Interfaces” on page 462 for Channelized OC3 IQ PICs, “Configuring Link PIC Failover on Channelized OC12/STM4 IQ and IQE Interfaces” on page 447 for OC12 IQ PICs, “Configuring Link PIC Failover on Channelized STM1 Interfaces” on page 475 for Channelized STM1 IQ PICs, and the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

no-translate-discard-eligible

See translate-discard-eligible

no-translate-fecn-and-becn

See translate-fecn-and-becn

no-unframed

See unframed

no-z0-increment

See z0-increment

oam

```

Syntax  oam {
            ethernet {
                connectivity-fault-management {
                    action-profile profile-name {
                        default-action {
                            interface-down;
                        }
                    }
                }
                performance-monitoring {
                    hardware-assisted-timestamping;
                }
                linktrace {
                    age (30m | 10m | 1m | 30s | 10s);
                    path-database-size path-database-size;
                }
                maintenance-domain domain-name {
                    level number;
                    name-format (character-string | none | dns | mac+2octet);
                    maintenance-association ma-name {
                        short-name-format (character-string | vlan | 2octet | rfc-2685-vpn-id);
                        continuity-check {
                            hold-interval (OAM) minutes;
                            interval (10m | 10s | 1m | 1s| 100ms);
                            loss-threshold number;
                        }
                        mep mep-id {
                            action-profile profile-name;
                            auto-discovery;
                            direction (up | down);
                            interface interface-name;
                            priority number;
                            remote-mep mep-id {
                                action-profile profile-name;
                            }
                        }
                    }
                }
            }
            link-fault-management {
                action-profile profile-name {
                    action {
                        syslog;
                        link-down;
                        send-critical-event;
                    }
                    event {
                        link-adjacency-loss;
                        link-event-rate {
                            frame-error count;
                            frame-period count;
                            frame-period-summary count;
                        }
                    }
                }
            }
        }

```

```

        symbol-period count;
    }
    protocol-down;
}
}
interface interface-name {
    apply-action-profile
    link-discovery (active | passive);
    pdu-interval interval;
    pdu-threshold threshold-value;
    remote-loopback;
    event-thresholds {
        frame-error count;
        frame-period count;
        frame-period-summary count;
        symbol-period count;
    }
    negotiation-options {
        allow-remote-loopback;
        no-allow-link-events;
    }
}
}
}
}
}

```

Hierarchy Level	[edit protocols]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	<p>For Ethernet interfaces on M320, M120, MX Series, and T Series routers, provide IEEE 802.3ah Operation, Administration, and Management (OAM) support.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring IEEE 802.3ah OAM Link-Fault Management” on page 745.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

oam-liveness

Syntax	oam-liveness { down-count <i>cells</i> ; up-count <i>cells</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM encapsulation only, configure Operation, Administration, and Maintenance (OAM) F5 loopback cell count thresholds. Not supported on ATM-over-SHDSL interfaces. For ATM2 IQ PICs only, configure OAM F4 loopback cell count thresholds at the [edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>] hierarchy level.
Options	down-count <i>cells</i> —Minimum number of consecutive OAM F4 or F5 loopback cells lost before a VC is declared down. Range: 1 through 255 Default: 5 cells up-count <i>cells</i> —Minimum number of consecutive OAM F4 or F5 loopback cells received before a VC is declared up. Range: 1 through 255 Default: 5 cells
Usage Guidelines	See “Configuring the ATM OAM F5 Loopback Cell Threshold” on page 329.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

oam-period

Syntax	oam-period (disable seconds);
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM encapsulation only, configure the OAM F5 loopback cell period. Not supported on ATM-over-SHDSL interfaces. For ATM2 IQ PICs only, configure OAM F4 loopback cell period at the [edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>] hierarchy level.
Default	If you omit this statement, OAM F5 loopback cells are not originated, but the interface still responds if it receives OAM F5 loopback cells.
Options	disable—Disable OAM loopback cell transmit feature. seconds—OAM loopback cell period. Range: 1 through 900 seconds
Usage Guidelines	See “Defining the ATM OAM F5 Loopback Cell Period” on page 329.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

oc-slice

Syntax	<code>oc-slice <i>oc-slice-range</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> partition <i>partition-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For channelized OC12 IQ interfaces only, configure the range of SONET/SDH slices.
Default	If you do not include either this statement or the <code>no-partition</code> statement, the Channelized OC12 IQ PICs not partitioned, and no data channels are configured.
Options	<p><i>oc-slice-range</i>—Range of SONET/SDH slices. OC3 interfaces must occupy three consecutive OC slices per interface, in the form 1–3, 4–6, 7–9, or 10–12. The T3, T1, and DS0 interface types each occupy one OC slice per interface.</p> <p>Range: For OC3 interfaces, 1–3, 4–6, 7–9, or 10–12; for SONET/SDH and T3 interfaces, 1–12.</p> <p>The remaining statement is explained separately.</p>
Usage Guidelines	See “Configuring Channelized OC12/STM4 Interfaces” on page 423.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

open-timeout

Syntax	<code>open-timeout <i>seconds</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure timeout period for Transmission Control Protocol (TCP) session establishment.
Options	<p><i>seconds</i>—Timeout period in seconds.</p> <p>Range: 4 through 86,400 seconds</p> <p>Default: 30 seconds</p>
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

operating-mode

Syntax	<code>operating-mode mode;</code>
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> dsl-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For J Series Services Routers only, modify the operating mode of the digital subscriber line for an ATM interface.
Options	<p><i>mode</i>—Operating mode for ATM-over-ADSL interfaces. The mode can be one of the following:</p> <ul style="list-style-type: none"> ■ <i>adsl2plus</i>—Set the ADSL line to train in the ITU G.992.5 mode. ■ <i>ansi-dmt</i>—Set the ADSL line to train in the ANSI T1.413 Issue 2 mode. ■ <i>auto</i>—Set the ADSL line to autonegotiate the setting to match the setting of the DSL access multiplexer (DSLAM) located at the central office. The ADSL line trains in the ANSI T1.413 Issue 2 (<i>ansi-dmt</i>) or ITU G.992.1 (<i>itu-dmt</i>) mode. ■ <i>etsi</i>—Set the ADSL line to train in the ETSI TS 101 388 V1.3.1 mode. ■ <i>itu-annexb-ur2</i>—Set the ADSL line to train in the ITU G.992.1 UR-2 mode. ■ <i>itu-annexb-non-ur2</i>—Set the ADSL line to train in the ITU G.992.1 non-UR-2 mode. ■ <i>itu-dmt</i>—Set the ADSL line to train in the ITU G.992.1 mode. ■ <i>itu-dmt-bis</i>—Set the ADSL line to train in the ITU G.992.3 mode. <p>Default: <i>auto</i></p>
Usage Guidelines	See “Configuring ATM-over-ADSL Interfaces” on page 355.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

optics-options

Syntax	optics-options { wavelength <i>nm</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For 10-Gigabit Ethernet dense wavelength-division multiplexing (DWDM) interfaces only, configure full C-band International Telecommunication Union (ITU)-Grid tunable optics.
Options	The wavelength statement is explained separately.
Usage Guidelines	See “Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength” on page 779.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

otn-options

Syntax otn-options {
 apply-groups *group-name*;
 apply-groups-except *exception-group-name*;
 fec (efec | gfec | none);
 (laser-enable | no-laser-enable);
 (line-loopback } no-line-loopback);
 pass-thru;
 rate (fixed-stuff-bytes | no-fixed-stuff-bytes | pass-thru);
 trigger (oc-lof | oc-lom | oc-los | oc-wavelength-lock | odu-ais | odu-bbe-th | odu-bdi |
 odu-es-th | odu-lck | odu-oci | odu-sd | odu-ses-th | odu-ttim | odu-uas-th | opu-ptm |
 otu-ais | otu-bbe-th | otu-bdi | otu-es-th | otu-fec-deg | otu-fec-exe | otu-iae | otu-sd |
 otu-ses-th | otu-ttim | otu-uas-th);
 tti;
 }

Hierarchy Level [edit interfaces *ge-fpc/pic/port*]

Release Information Statement introduced in JUNOS Release 9.4.

Description Specify the Gigabit Ethernet Optical Transport Network (OTN) interface and the Xenpac interface OTN options.

Options apply-groups—Groups from which to inherit configuration data.

 apply-groups-except—Don't inherit configuration data from these groups.

 fec—Enable Forward Error Correction (FEC) mode.

 laser-enable—Enable laser.

 line-loopback—Enable line loopback.

 no-laser-enable—Don't enable laser.

 no-line-loopback—Don't enable line loopback.

 rate (options)—OTN mode, select from the following options:

- fixed-stuff-bytes—Fixed stuff bytes 11.0957 Gbps
- no-fixed-stuff-bytes—No fixed stuff bytes 11.0491 Gbps
- pass-through—Enable OTN passthrough mode.
- no-pass-through—Do not enable OTN passthrough mode.

 trigger—Defect triggers, specify from the following possible completions:

- oc-lof—OC Loss of Frame defect trigger.
- oc-lom—OC Loss of Multiframe defect trigger.

- oc-los—OC Loss of Signal defect trigger.
- oc-wavelength-lock—OC Wavelength Lock defect trigger.
- odu-ais—ODU Alarm Indication Signal defect trigger.
- odu-bbe-th—ODU Background Block Error Threshold defect trigger.
- odu-bdi—ODU Backward Defect Indication defect trigger.
- odu-es-th—ODU Errored Seconds Threshold defect trigger.
- odu-lck—ODU Locked defect trigger.
- odu-oci—ODU Open Connection Indication defect trigger.
- odu-sd—ODU Signal Degrade defect trigger.
- odu-ses-th—ODU Severely Errored Seconds Threshold defect trigger.
- odu-ttim—ODU Trail Trace Identifier Mismatch defect trigger.
- odu-uas-th—ODU Unavailable Seconds Threshold defect trigger.
- opu-ptm—OPU Payload Type Mismatch defect trigger.
- otu-ais—OTU Alarm Indication Signal defect trigger.
- otu-bbe-th—OTU Background Block Error Threshold defect trigger.
- otu-bdi—OTU Backward Defect Indication defect trigger.
- otu-es-th—OTU Errored Seconds Threshold defect trigger.
- otu-fec-deg—OTU FEC Degrade defect trigger.
- otu-fec-exe—OTU FEC Excessive Error defect trigger.
- otu-iae—OTU Incoming Alignment defect trigger.
- otu-sd—OTU Signal Degrade defect trigger.
- otu-ses-th—OTU Severely Errored Seconds Threshold defect trigger.
- otu-ttim—OTU Trail Trace Identifier Mismatch defect trigger.
- otu-uas-th—OTU Unavailable Seconds Threshold defect trigger.

tti—Trace identifier, select from the following options:

- odu-dapi—ODU Destination Access Point Identifier.
- odu-expected-receive-dapi—ODU Expected Receive Destination Access Point Identifier.
- odu-expected-receive-sapi—ODU Expected Receive Source Access Point Identifier.
- odu-sapi—ODU Source Access Point Identifier.
- otu-dapi—OTU Destination Access Point Identifier.
- otu-expected-receive-dapi—OTU Expected Receive Destination Access Point Identifier.
- otu-expected-receive-sapi—OTU Expected Receive Source Access Point Identifier.
- otu-sapi—OTU Source Access Point Identifier.

Usage Guidelines See “Configuring Gigabit Ethernet OTN Options” on page 773.

Required Privilege Level interfaces—To view this statement in the configuration.
interfaces-control—To add this statement to the configuration.

output

Syntax output {
 service-set *service-set-name* <service-filter *filter-name*>;
}

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family inet service],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*
family inet service]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define one or more output service sets and filters to be applied to traffic.

Options The remaining statements are explained separately.

Usage Guidelines See the *JUNOS Services Interfaces Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

output-list

Syntax output-list [*filter-names*];

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family *family* filter],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*
family *family* filter]

Release Information Statement introduced in JUNOS Release 7.6.

Description Apply a group of filters to evaluate when packets are transmitted on an interface.

Options [*filter-names*]—Name of a filter to evaluate when packets are transmitted on the
interface. Up to 16 filters can be included in a filter input list.

Usage Guidelines See “Applying a Filter to an Interface” on page 203 and the *JUNOS Services Interfaces Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics input-list, *JUNOS Policy Framework Configuration Guide*, *JUNOS System Basics Configuration Guide*

output-policer

Syntax	<code>output-policer <i>policer-name</i>;</code>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer]</code>
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Associate a Layer 2 policer with a logical interface. The <code>output-policer</code> and <code>output-three-color</code> statements are mutually exclusive.
Options	<i>policer-name</i> —Name of the policer that you define at the <code>[edit firewall]</code> hierarchy level.
Usage Guidelines	See “Applying a Policer” on page 759.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Topics	<code>input-policer</code>

output-priority-map

Syntax	<pre>output-priority-map { classifier { premium { forwarding-class <i>class-name</i> { loss-priority (high low); } } } }</pre>
Hierarchy Level	<code>[edit interfaces <i>interface-name</i> gige-ether-options ethernet-switch-profile <i>ethernet-policer-profile</i>]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet interfaces only, define the output policer priority map to be applied to outgoing frames on this interface. The statements are explained separately.
Usage Guidelines	See “Specifying an Output Priority Map” on page 759.
Required Privilege Level	<code>interface</code> —To view this statement in the configuration. <code>interface-control</code> —To add this statement to the configuration.
Related Topics	<code>input-priority-map</code>

output-three-color

Syntax	<code>output-three-color <i>policer-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> layer2-policer]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Associate a Layer 2, three-color policer with a logical interface. The <code>output-three-color</code> and <code>output-policer</code> statements are mutually exclusive.
Options	<i>policer-name</i> —Name of the three-color policer that you define at the [edit firewall] hierarchy level.
Usage Guidelines	See “Applying a Policer” on page 759.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	input-three-color

output-vlan-map

See the following sections:

- output-vlan-map (Gigabit Ethernet IQ) on page 1168
- output-vlan-map (Aggregated Ethernet) on page 1169

output-vlan-map (Gigabit Ethernet IQ)

Syntax output-vlan-map {
 (pop | pop-pop | pop-swap | push | push-push | swap | swap-push | swap-swap);
 inner-tag-protocol-id *tpid*;
 inner-vlan-id *number*;
 tag-protocol-id *tpid*;
 vlan-id *number*;
 }

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before JUNOS Release 7.4.
 pop-pop, pop-swap, push-push, swap-push, and swap-swap statements added in JUNOS Release 8.1.

Description For Gigabit Ethernet IQ interfaces only, define the rewrite operation to be applied to outgoing frames on this logical interface.

The statements are explained separately.

Usage Guidelines See “Stacking and Rewriting Gigabit Ethernet VLAN Tags” on page 641.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics input-vlan-map

output-vlan-map (Aggregated Ethernet)

Syntax	output-vlan-map { (pop push swap); tag-protocol-id <i>tpid</i> ; vlan-id <i>number</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	For aggregated Ethernet interfaces using Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces only, define the rewrite profile to be applied to outgoing frames on this logical interface. The statements are explained separately.
Usage Guidelines	See “Stacking and Rewriting Gigabit Ethernet VLAN Tags” on page 641.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	input-vlan-map

overflow

See the following sections:

- overflow (Receive Bucket) on page 1170
- overflow (Transmit Bucket) on page 1170

overflow (Receive Bucket)

Syntax	overflow (discard tag);
Hierarchy Level	[edit interfaces <i>interface-name</i> receive-bucket]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify how to handle packets that exceed the threshold for the receive leaky bucket.
Options	<p>tag—Tag, count, and process received packets that exceed the threshold.</p> <p>discard—Discard received packets that exceed the threshold. No counting is done.</p>
Usage Guidelines	See “Configuring Receive and Transmit Leaky Bucket Properties” on page 129 and “Configuring Receive and Transmit Leaky Bucket Properties on SONET/SDH Interfaces” on page 876.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

overflow (Transmit Bucket)

Syntax	overflow discard;
Hierarchy Level	[edit interfaces <i>interface-name</i> transmit-bucket]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Discard packets that exceed the threshold for the transmit leaky bucket.
Usage Guidelines	See “Configuring Receive and Transmit Leaky Bucket Properties” on page 129 and “Configuring Receive and Transmit Leaky Bucket Properties on SONET/SDH Interfaces” on page 876.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

paired-group

Syntax	<code>paired-group <i>group-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure load sharing between two working-protect circuit pairs.
Options	<i>group-name</i> —Circuit's group name, as configured with the <code>protect-circuit</code> or <code>working-circuit</code> statement.
Usage Guidelines	See “Configuring APS Load Sharing Between Circuit Pairs” on page 870.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<code>working-circuit</code>

pap

Syntax pap {
 access-profile *name*;
 default-pap-password *password*;
 local-name *name*;
 local-password *password*;
 passive;
 }

Hierarchy Level [edit interfaces *interface-name* ppp-options],
 [edit interfaces *interface-name* unit *logical-unit-number* ppp-options],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*
 ppp-options]

Release Information Statement introduced in JUNOS Release 8.3.

Description Configure the Password Authentication Protocol (PAP). Use PAP authentication as a means to provide a simple method for the peer to establish its identity using a two-way handshake. This is done only upon initial link establishment.

After the link is established, an ID and password pair is repeatedly sent by the peer to the authenticator until authentication is acknowledged or the connection is terminated.

The statements are explained separately.

Usage Guidelines See “Configuring the PPP Challenge Handshake Authentication Protocol” on page 112, “Configuring PPP PAP Authentication” on page 164, and “Tracing Operations of the pppd Process” on page 119.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics traceoptions, *JUNOS System Basics Configuration Guide*

pap-password

Syntax	<code>pap-password password;</code>
Hierarchy Level	[edit access profile <i>profile-name</i> client <i>client-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the Password Authentication Protocol (PAP) password.
Options	<i>password</i> —PAP password.
Usage Guidelines	See “Configuring PPP PAP Authentication” on page 164.
Required Privilege Level	admin—To view this statement in the configuration. admin-control—To add this statement to the configuration.
Related Topics	chap-secret and <i>JUNOS System Basics Configuration Guide</i>

partition

Syntax	<code>partition <i>partition-number</i> oc-slice <i>oc-slice-range</i> interface-type <i>type</i> timeslots <i>time-slot-range</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For IQ interfaces and J Series interfaces on the Dual-Port Channelized E1 and T1 PIM, configure the channelized interface partition. The partition number is correlated with the channel number. Partition and channel numbering on IQ interfaces begins with :1, not :0.
Default	If you omit this statement, the channelized PIC or PIM is not partitioned, and no data channels are configured.
Options	<p><i>partition-number</i>—Sublevel interface partition index.</p> <p>Ranges:</p> <ul style="list-style-type: none"> ■ 1 through 4 for an OC3 interface on a channelized OC12 IQ interface. ■ 1 through 12 for a T3 interface on a channelized OC12 IQ interface. ■ 1 through 4 for a T3 interface on a channelized T3 IQ interface. ■ 1 through 28 for a T1 IQ interface on a channelized OC12 IQ or channelized T3 IQ interface. ■ 1 through 10 for an E1 interface on a channelized E1 IQ interface. ■ 1 through 30 on a channelized E1 interface. ■ 1 through 23 on a channelized T1 interface. ■ 1 through 24 for NxDS0 interfaces on either channelized OC12 IQ or channelized DS3 IQ interfaces. ■ 0 through 31 (with 0 reserved for framing) for NxDS0 interfaces on channelized E1 IQ interfaces. <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring Channelized E1 Interfaces” on page 501, “Configuring Channelized OC12/STM4 Interfaces” on page 423, and “Configuring Channelized T3 Interfaces” on page 479.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	no-partition

passive

See the following sections:

- passive (CHAP) on page 1175
- passive (PAP) on page 1176

passive (CHAP)

Syntax passive;

Hierarchy Level [edit interfaces *interface-name* ppp-options chap],
[edit interfaces *interface-name* unit *logical-unit-number* ppp-options chap],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* ppp-options chap]

Release Information Statement introduced before JUNOS Release 7.4.

Description Do not challenge the peer, but respond if challenged. If you omit this statement from the configuration, the interface always challenges its peer.

For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:

- atm-ppp-llc—PPP over AAL5 LLC encapsulation.
- atm-ppp-vc-mux—PPP over AAL5 multiplex encapsulation.

Usage Guidelines See “Configuring Passive Mode” on page 114.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics *JUNOS System Basics Configuration Guide*

passive (PAP)

Syntax	passive;
Hierarchy Level	[edit interfaces <i>interface-name</i> ppp-options pap], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options pap], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options pap]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Initiate an authentication request when the PAP option is received from a peer. If you omit this statement from the configuration, the interface requires the peer to initiate an authentication request.
Usage Guidelines	See “Configuring Passive Mode” on page 117.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS System Basics Configuration Guide</i>

passive-monitor-mode

Syntax	passive-monitor-mode;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM, Ethernet, and SONET/SDH interfaces only, monitor packet flows from another router. If you include this statement in the configuration, the interface does not send keepalives or alarms, and does not participate actively on the network. For ATM and Ethernet interfaces, you can include this statement on the physical interface only. For SONET/SDH interfaces, you can include this statement on the logical interface only.
Usage Guidelines	See “Enabling Passive Monitoring on ATM Interfaces” on page 293, “Enabling Passive Monitoring on Ethernet Interfaces” on page 677, and “Enabling Passive Monitoring on SONET/SDH Interfaces” on page 874.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	multiservice-options, <i>JUNOS Services Interfaces Configuration Guide</i>

path-database-size

Syntax	path-database-size <i>path-database-size</i> ;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management linktrace]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Number of linktrace reply entries to be stored per linktrace request.
Options	path-database-size—Database size. Range: 1 through 500 Default: 100
Usage Guidelines	See “Configuring the Linktrace Database Size” on page 690.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

path-trace

Syntax	<code>path-trace <i>trace-string</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For SONET/SDH interfaces and 10-Gigabit Ethernet interfaces in WAN PHY mode, configure a path trace identifier, which is a text string that identifies the circuit.</p> <p>On SONET/SDH OC48 interfaces that are configured for channelized (multiplexed) mode (by including the <code>no-concatenate</code> statement at the [edit chassis fpc <i>slot-number</i> pic <i>pic-number</i>] hierarchy level), the <code>bytes e1-quiet</code> and <code>bytes f1</code> options have no effect. The <code>bytes f2</code>, <code>bytes z3</code>, <code>bytes z4</code>, and <code>path-trace</code> options work correctly on channel 0 and work in the transmit direction only on channels 1, 2, and 3.</p> <p>For DS3 channels on a channelized OC12 interface, you can configure a unique path trace for each of the 12 channels. Each path trace can be up to 16 bytes. For channels on a channelized OC12 IQ interface, each path trace can be up to 64 bytes.</p>
Options	<p><i>trace-string</i>—Text string that identifies the circuit. If the string contains spaces, enclose it in quotation marks. A common convention is to use the circuit identifier as the path trace identifier. If you do not configure an identifier, the JUNOS Software uses the system and interface names to construct the default <i>trace-string</i>. For all nonchannelized SONET/SDH interfaces, the default <i>trace-string</i> is <i>system-name interface-name</i>. For channelized SONET/SDH interfaces and 10-Gigabit Ethernet WAN-PHY interfaces, the default <i>trace-string</i> is <i>interface-name</i>.</p>
Usage Guidelines	See “Configuring the SONET/SDH Path Trace Identifier” on page 853.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	sonet-options

payload-scrambler

Syntax	(payload-scrambler no-payload-scrambler);
Hierarchy Level	[edit interfaces <i>interface-name</i> e3-options], [edit interfaces <i>interface-name</i> sonet-options], [edit interfaces <i>interface-name</i> t3-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Enable or disable HDLC scrambling on an E3, a SONET/SDH, or a T3 interface. This type of scrambling provides better link stability. Both sides of a connection must either use or not use scrambling.</p> <p>If you commit a T3 interface configuration that has HDLC payload scrambling enabled, the interface must also be configured to be compatible with the channel service unit (CSU) at the remote end of the line.</p> <p>Disable payload scrambling on an E3 interface if Digital Link compatibility mode is used.</p> <p>On a channelized OC12 interface, the sonet payload-scrambler statement is ignored. To configure scrambling on the DS3 channels on the interface, you can include the t3-options payload-scrambler statement in the configuration for each DS3 channel.</p>
Default	Payload scrambling is disabled on all E3 and T3 interfaces; it is enabled by default on E3/T3 over ATM interfaces and on SONET/SDH interfaces.
Usage Guidelines	See “Configuring E3 and T3 Parameters on ATM Interfaces” on page 337, “Configuring E3 HDLC Payload Scrambling” on page 557, “Configuring SONET/SDH HDLC Payload Scrambling” on page 854, “Configuring T3 HDLC Payload Scrambling” on page 578, and “Examples: Configuring T3 Interfaces” on page 579.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	compatibility-mode

payload-size

Syntax	payload-size <i>bytes</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> satop-options],
Release Information	Option introduced in JUNOS Release 9.3.
Description	Specify the satop-options payload-size in integer number of bytes.
Usage Guidelines	See “Circuit Emulation PICs Overview” on page 519.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	satop-options

p-bit-timeout

Syntax	p-bit-timeout <i>time</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series routers only. On Ethernet interfaces configured for DLSw, configure the length of time the router waits for a response to a poll bit.
Options	<i>time</i> —Number of milliseconds. Range: 1 through 60000 Default: 3000 milliseconds
Usage Guidelines	See “Configuring LLC2 Options” on page 180.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

pdu-interval

Syntax	<code>pdu-interval <i>interval</i>;</code>
Hierarchy Level	[edit protocols oam ethernet link-fault-management interface <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	For Ethernet interfaces on M320, M120, MX Series, and T Series routers, specify the periodic OAM PDU sending interval for fault detection. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.
Options	<i>interval</i> —Periodic OAM PDU sending interval. Range: 100 through 1000 milliseconds Default: 1000 milliseconds
Usage Guidelines	See “Configuring the OAM PDU Interval” on page 747.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

pdu-threshold

Syntax	<code>pdu-threshold <i>threshold-value</i>;</code>
Hierarchy Level	[edit protocols oam ethernet link-fault-management interface <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	For Ethernet interfaces on M320, M120, MX Series, and T Series routers, specify the number of OAM PDUs to miss before an error is logged. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.
Options	<i>threshold-value</i> —The number of PDUs missed before declaring the peer lost. Range: 3 through 10 PDUs Default: 3 PDUs
Usage Guidelines	See “Configuring the OAM PDU Threshold” on page 747.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

peer

Syntax	<code>peer ip-address priority-cost priority;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track dlsw], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i> track dlsw]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	For J Series Services Routers only, Ethernet interfaces configured for DLSw, enable tracking options for a remote peer.
Options	<i>ip-address</i> —IP address of the remote peer. <i>priority-cost priority</i> —Cost value that is subtracted from the priority value when remote peer connectivity is lost. Specify a value from 1 through 254.
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 181.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Services Interfaces Configuration Guide</i>

peer-unit

Syntax	<code>peer-unit unit-number;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure a peer relationship between two logical systems.
Options	<i>unit-number</i> —Peering logical system unit number.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

performance-monitoring

Syntax	performance-monitoring { hardware-assisted-timestamping; }
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	For Ethernet interfaces on Dense Port Concentrators (DPCs) in MX Series routers only, specify performance monitoring support for Ethernet frame delay measurement.
Usage Guidelines	See “Ethernet Frame Delay Measurements Overview” on page 711, “Guidelines for Configuring Routers to Support an ETH-DM Session” on page 717, and “Enabling the Hardware-Assisted Timestamping Option” on page 726.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

periodic

Syntax	periodic <i>interval</i> ;
Hierarchy Level	[edit interfaces aex aggregated-ether-options lacp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For aggregated Ethernet interfaces only, configure the interval for periodic transmission of LACP packets.
Options	<i>interval</i> —Interval for periodic transmission of LACP packets. <ul style="list-style-type: none"> ■ fast—Transmit packets every second. ■ slow—Transmit packets every 30 seconds. <p>Default: fast</p>
Usage Guidelines	See “Configuring Aggregated Ethernet LACP” on page 627.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

per-unit-scheduler

Syntax per-unit-scheduler;

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description For channelized OC3 IQ, channelized OC12 IQ, channelized STM1 IQ, channelized T3 IQ, channelized E1 IQ, E3 IQ, link services IQ interfaces (lsq-), link services (ls-) on J Series routers, Gigabit Ethernet IQ, Gigabit Ethernet IQ2 and IQ2-E, and 10-Gigabit Ethernet interfaces only, enable association of scheduler map names with logical interfaces.



NOTE: Per-unit scheduling is not supported on T1 interfaces configured on the Channelized OC12 IQ PIC.



NOTE: On Gigabit Ethernet IQ2 and IQ2-E PICs without the **per-unit-scheduler** statement, the entire PIC supports 4071 VLANs and the user can configure all the VLANs on the same port.

On Gigabit Ethernet IQ2 and IQ2-E PICs with the **per-unit-scheduler** statement, the entire PIC supports $1024 - 2 * \text{number of ports}$ (1024 minus two times the number of ports), because each port is allocated two default schedulers.

Usage Guidelines When configuring the **per-unit-scheduler** statement on interfaces on the IQ2 and IQ2-E PIC, you must also include the **vlan-tagging** statement. See the *JUNOS Class of Service Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

pfc

Syntax	pfc;
Hierarchy Level	[edit interfaces <i>interface-name</i> ppp-options compression], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-options compression], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> ppp-option compression]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For interfaces with PPP encapsulation, configure the router to compress the protocol field to one byte.
Usage Guidelines	See “Configuring the PPP Protocol Field Compression” on page 121.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

pic-type

Syntax	pic-type (atm1 atm2);
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> atm-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM interfaces, configure the type of ATM PIC installed in your router.
Options	atm1—ATM1 PIC atm2—ATM2 IQ PIC
Usage Guidelines	See “Configuring the ATM PIC Type” on page 295.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

plp1

Syntax	<code>plp1 cells;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, define the EPD threshold on a VC. The EPD threshold is a limit on the number of transmit packets that can be queued. Packets that exceed the limit are discarded. This threshold applies to packets that have a PLP of 1.
Default	EPD threshold is unregulated.
Options	<i>cells</i> —Maximum number of cells. Range: For 1-port and 2-port OC12 interfaces, 1 through 425,984 cells. For 1-port OC48 interfaces, 1 through 425,984 cells. For 2-port OC3, DS3, and E3 interfaces, 1 through 212,992 cells. For 4-port DS3 and E3 interfaces, 1 through 106,496 cells.
Usage Guidelines	See “Configuring Two EPD Thresholds per Queue” on page 328 and “Configuring an ATM Scheduler Map” on page 341.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	epd-threshold, linear-red-profile

plp-to-clp

Syntax	plp-to-clp;
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> atm-options], [edit interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, enable the PLP setting to be copied to the cell-loss priority (CLP) bit.
Default	If you omit this statement, the JUNOS Software does not copy the PLP setting to the CLP bit.
Usage Guidelines	See “Enabling the PLP Setting to Be Copied to the CLP Bit” on page 348.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

point-to-point

Syntax	point-to-point;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For all interfaces except aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet, configure the interface unit as a point-to-point connection. This is the default connection type.
Default	If you omit this statement, the interface unit is configured as a point-to-point connection.
Usage Guidelines	See “Configuring a Point-to-Point Connection” on page 157.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	multipoint

policer

See the following sections:

- policer (CoS) on page 1188
- policer (Interface) on page 1189
- policer (MAC) on page 1190

policer (CoS)

Syntax `policer cos-policer-name {
 aggregate {
 bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
 burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
 }
 premium {
 bandwidth-limit (Policer for Gigabit Ethernet Interfaces) bps;
 burst-size-limit (Policer for Gigabit Ethernet Interfaces) bytes;
 }
 }`

Hierarchy Level [edit interfaces *interface-name* *gigether-options* ethernet-switch-profile
 ethernet-policer-profile]

Release Information Statement introduced before JUNOS Release 7.4.

Description For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), define a CoS policer template to specify the premium bandwidth and burst-size limits, and the aggregate bandwidth and burst-size limits. For Gigabit Ethernet interfaces with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), the premium policer is not supported.

Options *cos-policer-name*—Name of one policer to specify the premium bandwidth and burst-size limits, and the aggregate bandwidth and burst-size limits.

The remaining statements are explained separately.

Usage Guidelines See “Configuring Gigabit Ethernet Policers” on page 757.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

policer (Interface)

Syntax	<pre>policer { arp <i>policer-template-name</i>; input <i>policer-template-name</i>; output <i>policer-template-name</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Apply a policer to an interface.
Options	<p>arp <i>policer-template-name</i>—For inet family only, name of one policer to evaluate when ARP packets are received on the interface.</p> <p>input <i>policer-template-name</i>—Name of one policer to evaluate when packets are received on the interface.</p> <p>output <i>policer-template-name</i>—Name of one policer to evaluate when packets are transmitted on the interface.</p>
Usage Guidelines	See “Applying Policers” on page 194 and the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Policy Framework Configuration Guide</i>

policer (MAC)

Syntax `policer {
 input cos-policer-name;
 output cos-policer-name;
 }`

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* accept-source-mac
 mac-address *mac-address*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*
 accept-source-mac mac-address *mac-address*]

Release Information Statement introduced before JUNOS Release 7.4.

Description For Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), configure MAC policing.

**NOTE:**

On MX Series routers with Gigabit Ethernet or Fast Ethernet PICs, the following considerations apply:

- Interface counters do not count the 7-byte preamble and 1-byte frame delimiter in Ethernet frames.
 - In MAC statistics, the frame size includes MAC header and CRC before any VLAN rewrite/imposition rules are applied.
 - In traffic statistics, the frame size encompasses the L2 header without CRC after any VLAN rewrite/imposition rule.
-

Options `input cos-policer-name`—Name of one policer to specify the premium bandwidth and aggregate bandwidth.

`output cos-policer-name`—Name of one policer to specify the premium bandwidth and aggregate bandwidth.

Usage Guidelines See “Configuring MAC Address Filtering” on page 761.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

pool

Syntax	<code>pool <i>pool-name</i> <priority <i>priority</i>>;</code>
Hierarchy Level	[edit interfaces <i>br-pim</i> / <i>0</i> / <i>port</i> dialer-options], [edit interfaces <i>umd</i> <i>0</i> dialer-options], [edit interfaces <i>dl</i> <i>n</i> unit <i>logical-unit-number</i> dialer-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>dl</i> <i>n</i> unit <i>logical-unit-number</i> dialer-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	On J Series Services Routers, for logical and physical ISDN interfaces, specify the dial pool. The dial pool allows logical (dialer) and physical (<i>br-pim</i>/<i>0</i>/<i>port</i>) interfaces to be bound together dynamically on a per-call basis. On a dialer interface, pool directs the dialer interface which dial pool to use. On <i>br-pim</i>/<i>0</i>/<i>port</i> interface, pool defines the pool to which the interface belongs.
Options	<i>pool-name</i> —Pool identifier. <i>priority priority</i> —(Physical <i>br-pim</i>/<i>0</i>/<i>port</i> interfaces only) Specify a priority value of 0 (lowest) to 255 (highest) for the interface within the pool.
Usage Guidelines	See “Configuring ISDN Physical Interface Properties” on page 821 and “Specifying a USB Modem Interface on J Series Routers” on page 93.
Required Privilege Level	<i>interface</i> —To view this statement in the configuration. <i>interface-control</i> —To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

pop

Syntax	pop;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, specify the VLAN rewrite operation to remove a VLAN tag from the top of the VLAN tag stack. The outer VLAN tag of the frame is removed.
Usage Guidelines	See “Removing a VLAN Tag” on page 649.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

pop-all-labels

Syntax	pop-all-labels { required-depth <i>number</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options mpls], [edit interfaces <i>interface-name</i> sonet-options mpls], [edit interfaces <i>interface-name</i> fastether-options mpls], [edit interfaces <i>interface-name</i> gigeether-options mpls]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For passive monitoring on ATM and SONET/SDH interfaces only, removes up to two MPLS labels from incoming IP packets.</p> <p>This statement has no effect on IP packets with more than two MPLS labels. Packets with MPLS labels cannot be processed by the Monitoring Services PIC; if packets with MPLS labels are forwarded to the Monitoring Services PIC, they are discarded.</p> <p>The remaining statement is explained separately.</p>
Default	If you omit this statement, the MPLS labels are not removed, and the packet is not processed by the Monitoring Services PIC.
Usage Guidelines	See “Removing MPLS Labels from Incoming Packets” on page 294 and “Removing MPLS Labels from Incoming Packets” on page 874.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Services Interfaces Configuration Guide</i>

pop-pop

Syntax	pop-pop;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, specify the VLAN rewrite operation to remove both the outer and inner VLAN tags of the frame.
Usage Guidelines	See “Removing the Outer and Inner VLAN Tags” on page 649.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

pop-swap

Syntax	pop-swap;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, specify the VLAN rewrite operation to remove the outer VLAN tag of the frame, and replace the inner VLAN tag of the frame with a user-specified VLAN tag value. The inner tag becomes the outer tag in the final frame.
Usage Guidelines	See “Removing the Outer VLAN Tag and Rewriting the Inner VLAN Tag” on page 650.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

port

Syntax	port { minimum <i>port-number</i> ; maximum <i>port-number</i> ; }
Hierarchy Level	[edit interfaces <i>vsp-fpc/pic/port</i> unit <i>logical-unit-number</i> compression rtp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For voice services interfaces only, assign User Datagram Protocol (UDP) destination port numbers reserved for Real-Time Transport Protocol (RTP) traffic.
Options	<p>minimum <i>port-number</i>—Specify minimum port number. Range: 0 through 65,535</p> <p>maximum <i>port-number</i>—Specify maximum port number. Range: 0 through 65,535</p>
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

port-priority

Syntax	port-priority <i>priority</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> together-options 802.3ad lacp]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Define LACP port priority at the interface level.
Options	<p><i>priority</i>—Priority for being elected to be the active port and both collect and distribute traffic. A smaller value indicates a higher priority for being elected. Range: 1 through 255 Default: 127</p>
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

port-status-tlv

Syntax	port-status-tlv <blocked>;
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management action-profile (Defining for CFM) <i>tlv-action</i> event]
Release Information	Statement introduced in JUNOS Release 9.6.
Description	Define an action-profile consisting of various events and the action. Based on values of port-status-tlv in the received CCM packets, specific action such as <i>interface-down</i> can be taken using action-profile (Defining for CFM) options.
Options	blocked —When the incoming CCM packet contains port status TLV with value blocked, the action will be triggered for this action-profile.
Usage Guidelines	See “Configuring Remote MEP Action Profile Support” on page 707.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

post-service-filter

Syntax	post-service-filter <i>filter-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service input], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service input]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define the filter to be applied to traffic after service processing. The filter is applied only if a service set is configured and selected.
Options	filter-name —Identifier for postservice filter.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

pppoe-options

Syntax	<pre>pppoe-options { access-concentrator <i>name</i>; auto-reconnect <i>seconds</i>; (client server); service-name <i>name</i>; underlying-interface <i>interface-name</i>; }</pre>
Hierarchy Level	[edit interfaces pp0 unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces pp0 unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4. client Statement introduced in JUNOS Release 8.5. server Statement introduced in JUNOS Release 8.5.
Description	<p>For J Series Services Routers and M120 Multiservice Edge Routers with PPP over Ethernet interfaces, configure PPP over Ethernet-specific interface properties.</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring a PPPoE Interface” on page 790.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

ppp-options

Syntax	<pre> ppp-options { chap { access-profile <i>name</i>; default-chap-secret <i>name</i>; local-name <i>name</i>; passive; } compression { acfc; pfc; } dynamic-profile <i>profile-name</i>; lcp-max-conf-req <i>number</i> lcp-restart-timer <i>milliseconds</i>; loopback-clear-timer <i>seconds</i>; ncp-max-conf-req <i>number</i> ncp-restart-timer <i>milliseconds</i>; pap { access-profile <i>name</i>; default-pap-password <i>password</i>; local-name <i>name</i>; local-password <i>password</i>; passive; } } </pre>
Hierarchy Level	<pre> [edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] </pre>
Release Information	<p>Statement introduced before JUNOS Release 7.4.</p> <p><code>lcp-restart-timer</code> statement introduced in JUNOS Release 8.1.</p> <p><code>ncp-restart-timer</code> statement introduced in JUNOS Release 8.1.</p> <p><code>loopback-clear-timer</code> statement introduced in JUNOS Release 8.5.</p> <p><code>dynamic-profile</code> statement introduced in JUNOS Release 9.5.</p>
Description	<p>On interfaces with PPP encapsulation, configure PPP-specific interface properties.</p> <p>For ATM2 IQ interfaces only, you can configure CHAP on the logical interface unit if the logical interface is configured with one of the following PPP over ATM encapsulation types:</p> <ul style="list-style-type: none"> ■ <code>atm-ppp-llc</code>—PPP over AAL5 LLC encapsulation. ■ <code>atm-ppp-vc-mux</code>—PPP over AAL5 multiplex encapsulation. <p>The remaining statements are explained separately.</p>
Usage Guidelines	<p>See “Configuring the PPP Challenge Handshake Authentication Protocol” on page 112.</p>

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

preempt

Syntax (preempt | no-preempt) {
 hold-time *seconds*;
 }

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* family llc2 redundancy-group *group-number*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number* family llc2 redundancy-group *group-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure whether a DLSw backup router can preempt a master router:

- preempt—Allow the master router to be preempted.
- no-preempt—Prohibit the preemption of the master router.

The remaining statement is explained separately.

Default If you omit this statement, the backup router cannot preempt a master router.

Usage Guidelines See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 181.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics For information about the `preempt` statement at the [edit interfaces *interface-name* unit *unit-number* family inet address *address* (vrrp-group | vrrp-inet6-group) *group-number*] or [edit logical-systems *logical-system-name* interfaces *interface-name* unit *unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-number*] hierarchy level, see the *JUNOS High Availability Configuration Guide*.

preferred

Syntax	preferred;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure this address to be the preferred address on the interface. If you configure more than one address on the same subnet, the preferred source address is chosen by default as the source address when you originate packets to destinations on the subnet.
Default	The lowest numbered address on the subnet is the preferred address.
Usage Guidelines	See “Configuring the Interface Address” on page 174.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

preferred-source-address

Syntax	<code>preferred-source-address address;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> unnumbered-address <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> unnumbered-address <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	<p>For unnumbered Ethernet interfaces configured with a loopback interface as the donor interface, specify one of the loopback interface's secondary addresses as the preferred source address for the unnumbered Ethernet interface. Configuring the preferred source address enables you to use an IP address other than the primary IP address on some of the unnumbered Ethernet interfaces in your network.</p> <p>Configuration of a preferred source address for unnumbered Ethernet interfaces is supported for the IPv4 and IPv6 address families.</p>
Options	<i>address</i> —Secondary IP address of the donor loopback interface.
Usage Guidelines	See “Configuring a Preferred Source Address for Unnumbered Ethernet or Demux Interfaces” on page 187.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>address</i> , <i>JUNOS System Basics Configuration Guide</i>

premium

See the following sections:

- premium (Hierarchical Policer) on page 1202
- premium (Output Priority Map) on page 1203
- premium (Policer) on page 1203

premium (Hierarchical Policer)

Syntax premium {
 if-exceeding {
 bandwidth-limit *bandwidth*;
 burst-size-limit *burst*;
 }
 then {
 discard;
 }
 }

Hierarchy Level [edit firewall hierarchical-policer]

Release Information Statement introduced in JUNOS Release 9.5.

Description On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, to specify a premium level for a hierarchical policer, use the **premium** statement at the [edit firewall hierarchical-policer] hierarchy level.

Options Options are described separately.

Usage Guidelines See “Applying Policers” on page 194 and the *JUNOS Class of Service Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

premium (Output Priority Map)

Syntax	<pre>premium { forwarding-class <i>class-name</i> { loss-priority (high low); } }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> gigether-options ethernet-switch-profile ethernet-policer-profile output-priority-map classifier]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For Gigabit Ethernet IQ interfaces only, define the classifier for egress premium traffic.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Specifying an Output Priority Map” on page 759.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	input-priority-map

premium (Policer)

Syntax	<pre>premium { bandwidth-limit (Policer for Gigabit Ethernet Interfaces) <i>bps</i>; burst-size-limit (Policer for Gigabit Ethernet Interfaces) <i>bytes</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> gigether-options ethernet-switch-profile ethernet-policer-profile policer <i>cos-policer-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Define a policer to apply to nonpremium traffic.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Configuring Gigabit Ethernet Policers” on page 757.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	aggregate, ieee802.1p

preserve-interface

Syntax	preserve-interface;
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced in JUNOS Release 7.6.
Description	<p>Provide link PIC replication, providing MLPPP link redundancy at the port level. This feature is supported with SONET APS and the following link PICs:</p> <ul style="list-style-type: none"> ■ Channelized OC3 IQ PIC ■ Channelized OC12 IQ PIC ■ Channelized STM1 IQ PIC <p>Link PIC replication provides the ability to add two sets of links, one from the active SONET PIC and the other from the standby SONET PIC, to the same bundle. If the active SONET PIC fails, links from the standby PIC are used without triggering link renegotiation. All the negotiated state is replicated from the active links to the standby links to prevent link renegotiation.</p>
Usage Guidelines	See “Configuring Link PIC Redundancy” on page 869.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>JUNOS Services Interfaces Configuration Guide</i>

primary

See the following sections:

- primary (Address on Interface) on page 1205
- primary (AS PIC or MultiServices PIC Interfaces) on page 1205

primary (Address on Interface)

Syntax	primary;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure this address to be the primary address of the protocol on the interface. If the logical unit has more than one address, the primary address is used by default as the source address when packets originate from the interface and the destination does not indicate the subnet.
Default	For unicast traffic, the primary address is the lowest non-127 preferred address on the unit.
Usage Guidelines	See “Configuring the Interface Address” on page 174.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

primary (AS PIC or MultiServices PIC Interfaces)

Syntax	primary <i>interface-name</i> ;
Hierarchy Level	[edit interfaces (rsp0 rsp1) redundancy-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the primary AS PIC or MultiServices PIC interface.
Options	<i>interface-name</i> —The identifier for the AS PIC interface or MultiServices PIC interface, which must be of the form <i>sp-fpc/pic/port</i> .
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

priority

See the following sections:

- [priority \(DLSw\) on page 1206](#)
- [priority \(OAM Connectivity-Fault Management\) on page 1207](#)
- [priority \(Schedulers\) on page 1207](#)



NOTE: For information about the `priority` statement at the [edit interfaces *interface-name* unit *unit-number* family inet address *address* (vrrp-group | vrrp-inet6-group) *group-number*] or [edit logical-systems *logical-system-name* interfaces *interface-name* unit *unit-number* family (inet | inet6) address *address* (vrrp-group | vrrp-inet6-group) *group-number*] hierarchy level, see the *JUNOS High Availability Configuration Guide*.

priority (DLSw)

Syntax	<code>priority priority;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> redundancy-group <i>group-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet address <i>address</i> redundancy-group <i>group-number</i>]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	When configuring DLSw Ethernet redundancy on Fast Ethernet and Gigabit Ethernet interfaces, configure a DLSw router's priority for becoming the master default router. The router with the highest priority within the group becomes the master.
Options	<i>priority</i> —Router's priority for being elected to be the master router in the VRRP group. A larger value indicates a higher priority for being elected. Range: 1 through 255 Default: 100 (for backup routers)
Usage Guidelines	See "Configuring DLSw Ethernet Redundancy Using LLC2 Properties" on page 181.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

priority (OAM Connectivity-Fault Management)

Syntax	<code>priority number;</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> maintenance-association <i>ma-name</i> mep <i>mep-id</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	IEEE 802.1p priority bits used by the continuity check messages.
Options	<i>number</i> —Configure the IEEE 802.1p priority bits to be used in the VLAN header of the CFM packets. Range: 0 through 7
Usage Guidelines	See “Configuring the Maintenance End Point Priority” on page 687
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

priority (Schedulers)

Syntax	<code>priority (high low);</code>
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> atm-options scheduler-maps <i>map-name</i> forwarding-class <i>class-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, assign queuing priority to a forwarding class.
Options	<i>low</i> —Forwarding class has low priority. <i>high</i> —Forwarding class has high priority.
Usage Guidelines	See “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

priority-cost

Syntax	<code>priority-cost <i>priority</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>unit-number</i> family llc2 redundancy-group <i>group-number</i> priority <i>priority</i> track interface <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	When configuring DLSw Ethernet redundancy, configure a DLSw router's priority cost for becoming the master default router. The router with the highest priority within the group becomes the master.
Options	<p><code>priority-cost <i>priority</i></code>—The value subtracted from the configured DLSw priority when the tracked interface is down, forcing a new master router election. The sum of all the costs for all interfaces or routes that are tracked must be less than or equal to the configured priority of the DLSw group.</p> <p>Range: 1 through 254</p>
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 181.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

promiscuous-mode

Syntax	<pre>promiscuous-mode { vpi <i>vpi-identifier</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM interfaces with <code>atm-ccc-cell-relay</code> encapsulation, map all incoming cells from either an interface port or a VP to a single label-switched path (LSP) without restricting the VCI number. Promiscuous mode allows you to map traffic from all 65,535 VCIs to a single LSP, or from all 256 VPIs to a single LSP.
Options	<p><code>vpi-identifier</code>—Open this VPI in promiscuous mode.</p> <p>Range: 0 through 255</p>
Usage Guidelines	See “Configuring ATM Cell-Relay Promiscuous Mode” on page 296.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	vpi

protect-circuit

Syntax	protect-circuit <i>group-name</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the protect router in an APS circuit pair. When the working interface fails, APS brings up the protection circuit and the traffic is moved to the protection circuit.
Options	<i>group-name</i> —Circuit's group name.
Usage Guidelines	See “Configuring Basic APS Support” on page 861.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	working-circuit

protection-group

Syntax

```

protection-group {
  ethernet-ring ring-name (
    east-interface {
      control-channel channel-name {
        vlan number;
      }
    }
    guard-interval number;
    node-id mac-address;
    restore-interval number;
    ring-protection-link-owner;
    west-interface {
      control-channel channel-name {
        vlan number;
      }
    }
  }
}

```

Hierarchy Level [edit protocols]

Release Information Statement introduced in JUNOS Release 9.4.

Description Use this statement and its options to configure Ethernet ring protection switching.

Options The statement options are described separately.

Usage Guidelines See “Configuring Ethernet Ring Protection Switching” on page 799.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

protocol-down

Syntax protocol-down;

Hierarchy Level [edit protocols oam ethernet link-fault-management action-profile event]

Release Information Statement introduced in JUNOS Release 8.5.

Description Upper layer indication of protocol down event. When the **protocol-down** statement is included, the protocol down event triggers the action specified under the **action** statement.

Usage Guidelines See “Configuring an OAM Action Profile” on page 748.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

protocols

Syntax	<code>protocols [inet iso mpls];</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit logical-unit-number family tcc]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	For Layer 2.5 VPNs on T Series, M120, and M320 routers support, configure IS-IS (ISO traffic) or MPLS traffic to traverse a TCC interface. By default, IPv4 (inet) traffic runs on T Series, M120, and M320 routers and over TCC interfaces. You must configure the same traffic type on both ends of the Layer 2.5 VPN.
Usage Guidelines	See “Configuring IS-IS or MPLS Traffic for TCC Interfaces” on page 229.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

proxy

Syntax	<code>proxy inet-address <i>address</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family tcc], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family tcc]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Layer 2.5 VPNs using an Ethernet interface as the TCC router, configure the IP address for which the TCC router is proxying. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only. Ethernet TCC is not supported on the T640 router.
Options	inet-address—Configure the IP address of the neighbor to the TCC router.
Usage Guidelines	See “Configuring Ethernet TCC” on page 666 and “Example: Configuring an Ethernet TCC or Extended VLAN TCC” on page 667.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	remote, <i>JUNOS VPNs Configuration Guide</i>

proxy-arp

Syntax	proxy-arp;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Ethernet interfaces only, configure the router to respond to any ARP request, as long as the router has an active route to the ARP request's target address.
Usage Guidelines	See “Configuring Unrestricted Proxy ARP” on page 671.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

push

Syntax	push;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and aggregated Ethernet interfaces using Gigabit Ethernet IQ interfaces, specify the VLAN rewrite operation to add a new VLAN tag to the top of the VLAN stack. An outer VLAN tag is pushed in front of the existing VLAN tag. If you include the push statement in the configuration, you must also include the pop statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map] hierarchy level.
Usage Guidelines	See “Stacking a VLAN Tag” on page 648.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

push-push

Syntax	push-push;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, specifies the VLAN rewrite operation to push two VLAN tags in front of the frame.
Usage Guidelines	See “Stacking Two VLAN Tags” on page 651.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

queue-depth

Syntax	queue-depth <i>cells</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options linear-red-profiles <i>profile-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, define maximum queue depth in the CoS VC drop profile. Packets are always dropped beyond the defined maximum. This statement is mandatory; there is no default configuration.
Default	Buffer usage is unregulated.
Options	<i>cells</i> —Maximum number of cells the queue can contain. Range: 1 through 64,000 cells
Usage Guidelines	See “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	high-plp-threshold, low-plp-threshold

queue-length

Syntax	queue-length <i>number</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM1 interfaces only, define the maximum queue length in the traffic-shaping profile. For ATM1 PICs, each VC has its own independent shaping parameters.
Default	Buffer usage is unregulated.
Options	<i>number</i> —Maximum number of packets the queue can contain. Range: 1 through 16383 packets Default: 16383 packets
Usage Guidelines	See “Configuring the ATM1 Queue Length” on page 325.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

queues

Syntax	queues [<i>queue-numbers</i>];
Hierarchy Level	[edit interfaces <i>vsp-fpc/pic/port</i> unit <i>logical-unit-number</i> compression rtp]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For voice services interfaces only, assign queue numbers for RTP traffic.
Options	queues <i>queue-numbers</i> —Assign one or more of the following queues: q0, q1, q2, q3.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> . For VRRP services, specify the q3 option instead of q0.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

quiet-period

Syntax	quiet-period <i>seconds</i> ;
Hierarchy Level	[edit protocols dot1x authenticator interface <i>interface-id</i>]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specify the number of seconds the port remains in the wait state following a failed authentication exchange with the client, before reattempting authentication.
Options	<i>seconds</i> —Specify the number of seconds the port remains in the wait state following a failed authentication exchange with the client, before reattempting authentication. Range: 0 through 65,535 seconds Default: 60 seconds
Usage Guidelines	See “Configuring IEEE 802.1x Port-Based Network Access Control” on page 741.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	dot1x, authenticator, interface (IEEE 802.1x)

ranges

See the following sections:

- ranges (Dynamic Stacked VLAN) on page 1216
- ranges (Dynamic VLAN) on page 1216

ranges (Dynamic Stacked VLAN)

Syntax	ranges (any <i>low-tag</i> - <i>high-tag</i>) , (any <i>low-tag</i> - <i>high-tag</i>);
Hierarchy Level	[edit interfaces <i>interface-name</i> auto-configure stacked-vlan-ranges dynamic-profile <i>profile-name</i>]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	Configure VLAN ranges for dynamic, auto-sensed stacked VLANs.
Options	<p>any—The entire VLAN range.</p> <p><i>low-tag</i>—The lower limit of the VLAN range.</p> <p><i>high-tag</i>—The upper limit of the VLAN range.</p> <p>Range: 1 through 4094</p>
Usage Guidelines	See Configuring Stacked VLAN Ranges for Use with Stacked VLAN Dynamic Profiles.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

ranges (Dynamic VLAN)

Syntax	ranges (any <i>low-tag</i>) - (any <i>high-tag</i>);
Hierarchy Level	[edit interfaces <i>interface-name</i> auto-configure vlan-ranges dynamic-profile (VLAN) <i>profile-name</i>]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	Configure VLAN ranges for dynamic, auto-sensed VLANs.
Options	<p>any—The entire VLAN range.</p> <p><i>low-tag</i>—The lower limit of the VLAN range.</p> <p><i>high-tag</i>—The upper limit of the VLAN range.</p> <p>Range: 1 through 4094</p>
Usage Guidelines	See Configuring Single-Level VLAN Ranges for Use with VLAN Dynamic Profiles.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

rate

Syntax	<code>rate percentage;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> receive-bucket], [edit interfaces <i>interface-name</i> transmit-bucket]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify percentage of the interface line rate that is available to receive or transmit packets.
Options	<i>percentage</i> —Percentage of the interface line rate that is available to receive or transmit packets. Range: 0 through 100
Usage Guidelines	See “Configuring Receive and Transmit Leaky Bucket Properties” on page 129 and “Configuring Receive and Transmit Leaky Bucket Properties on SONET/SDH Interfaces” on page 876.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

reassemble-packets

Syntax	<code>reassemble-packets;</code>
Hierarchy Level	[edit interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>gr-fpc/pic/port</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Enable reassembly of fragmented tunnel packets on generic routing encapsulation (GRE) tunnel interfaces.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

reauthentication

Syntax	reauthentication (disable interval <i>seconds</i>);
Hierarchy Level	[edit protocols dot1x authenticator interface <i>interface-id</i>]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Set or disable the periodic reauthentication of the client.
Options	<ul style="list-style-type: none"> ■ disable—Disable the periodic reauthentication of the client. ■ interval <i>seconds</i>—Specify the periodic reauthentication time interval. <p>Range: 1 through 65,535 seconds Default: 3600 seconds</p>
Usage Guidelines	See “Configuring IEEE 802.1x Port-Based Network Access Control” on page 741
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	dot1x, interface (IEEE 802.1x), quiet-period

receive-bucket

Syntax	<pre>receive-bucket { overflow (discard tag); rate <i>percentage</i>; threshold <i>bytes</i>; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Set parameters for the receive leaky bucket, which specifies what percentage of the interface’s total capacity can be used to receive packets.</p> <p>For each DS3 channel on a channelized OC12 interface, you can configure a unique receive bucket.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Configuring Receive and Transmit Leaky Bucket Properties on SONET/SDH Interfaces” on page 876.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	transmit-bucket

receive-options-packets

Syntax	receive-options-packets;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For a Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.
Usage Guidelines	See “Enabling Passive Monitoring on ATM Interfaces” on page 293 and “Enabling Passive Monitoring on SONET/SDH Interfaces” on page 874.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

receive-ttl-exceeded

Syntax	receive-ttl-exceeded;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Monitoring Services PIC and an ATM or SONET/SDH PIC installed in an M160, M40e, or T Series router, guarantee conformity with cflowd records structure. This statement is required when you enable passive monitoring.
Usage Guidelines	See “Enabling Passive Monitoring on ATM Interfaces” on page 293 and “Enabling Passive Monitoring on SONET/SDH Interfaces” on page 874.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

red-differential-delay

Syntax	<code>red-differential-delay <i>milliseconds</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services and voice services interfaces only, configure the red differential delay among bundle links to give warning when a link has a differential delay that exceeds the configured threshold.
Options	<i>milliseconds</i> —Red differential delay threshold. Range: 1 through 2000 milliseconds Default: 10 milliseconds
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	action-red-differential-delay, yellow-differential-delay

redial-delay

Syntax	<code>redial-delay <i>time</i>;</code>
Hierarchy Level	[edit interfaces <i>dln</i> unit <i>logical-unit-number</i> dialer-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>dln</i> unit <i>logical-unit-number</i> dialer-options]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	<p>On J Series Services Routers with interfaces configured for ISDN with dialout, specify the delay (in seconds) between two successive calls made by the dialer. To configure callback mode, include the <code>callback</code> statement at the [edit interfaces <i>dln</i> unit <i>logical-unit-number</i> dialer-options] hierarchy level.</p> <p>If the <code>callback</code> statement is configured, you cannot use the <code>caller <i>caller-id</i></code> statement at the [edit interfaces <i>dln</i> unit <i>logical-unit-number</i> dialer-options] hierarchy level.</p>
Options	<i>time</i> —Delay (in seconds) between two successive calls. Range: 2 through 255 seconds Default: 3 seconds
Usage Guidelines	See “Configuring ISDN Interfaces” on page 819.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

redundancy-group

Syntax	<pre> redundancy-group <i>group-number</i> { advertise-interval <i>seconds</i>; map { local-mac <i>mac-address</i> request <i>mac-address</i>; } preempt hold-time <i>seconds</i>; no-preempt; priority <i>priority</i>; track { dls { peer <i>ip-address</i> priority-cost <i>priority</i>; destination <i>mac-address</i> priority-cost <i>priority</i>; } interface <i>interface-name</i> priority-cost <i>priority</i>; } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2]
Release Information	Statement introduced in JUNOS Release 7.5.
Description	<p>For J Series Services Routers only. On Ethernet interfaces configured for DLSw, configure the router for DLSw redundancy.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 181.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

redundancy-options

Syntax	<pre> redundancy-options { primary <i>interface-name</i>; secondary <i>interface-name</i>; hot-standby; } </pre>
Hierarchy Level	[edit interfaces (rsp0 rsp1)], [edit interfaces rlsqnumber]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the primary and secondary (backup) AS PIC interfaces or MultiServices PIC interfaces.
Options	<p>primary <i>interface-name</i>—The identifier for the primary LSQ AS or MultiServices PIC interface.</p> <p>secondary <i>interface-name</i>—The identifier for the secondary (backup) LSQ AS or MultiServices PIC interface.</p> <p>hot-standby—For one-to-one AS or MultiServices PIC redundancy configurations, specify that the failure detection and recovery must take place in less than 5 seconds.</p>
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.


remote

Syntax	remote { (inet-address <i>address</i> mac-address <i>address</i>); }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family tcc], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family tcc]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Layer 2.5 VPNs using an Ethernet interface as the TCC router, configure the location of the remote router. Ethernet TCC is supported on interfaces that carry IPv4 traffic only. Ethernet TCC encapsulation is supported on 1-port Gigabit Ethernet, 2-port Gigabit Ethernet, 4-port Gigabit Ethernet, and 4-port Fast Ethernet PICs only.
Options	mac-address—Configure the MAC address of the remote site. inet-address—Configure the IP address of the remote site.
Usage Guidelines	See “Configuring Ethernet TCC” on page 666 and “Example: Configuring an Ethernet TCC or Extended VLAN TCC” on page 667.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	proxy, <i>JUNOS VPNs Configuration Guide</i>

remote-loopback

Syntax	remote-loopback;
Hierarchy Level	[edit protocols oam link-fault-management interface <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	For Ethernet interfaces on M320, M120, MX Series, and T Series routers, set the remote DTE into loopback mode. Remove the statement from the configuration to take the remote DTE out of loopback mode. Used for IEEE 802.3ah Operation, Administration, and Management (OAM) support.
Usage Guidelines	See “Setting a Remote Interface into Loopback Mode” on page 751.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

remote-loopback-respond

Syntax	remote-loopback-respond;
Hierarchy Level	[edit interfaces <i>ct1-fpc/pic/port</i>], [edit interfaces <i>interface-name</i> t1-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For T1 interfaces only, configure the router to respond to remote loopback requests. Remote loopback requests can be from the facilities data link or inband.
<hr/>	
	NOTE: When configuring CT1 interfaces on the 10-port Channelized E1/T1 IQE PIC, the <code>remote-loopback-respond</code> statement must be included at the [edit interfaces <i>ct1-fpc/pic/port</i>] hierarchy level.
<hr/>	
Default	The router does not respond to remote loop requests.
Usage Guidelines	See “Configuring the T1 Remote Loopback Response” on page 564.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	feac-loop-respond, loopback

remote-mep

Syntax	<code>remote-mep mep-id { action-profile profile-name; }</code>
Hierarchy Level	<code>[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance-association ma-name mep mep-id]</code>
Release Information	Statement introduced in JUNOS Release 8.4.
Description	The numeric identifier of the remote maintenance association end point (MEP) within the maintenance association.
Options	<p>mep-id—Specify the numeric identifier of the MEP. Range: 1 through 8191</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Configuring a Remote Maintenance End Point” on page 688.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

request

Syntax	<code>request (protect working);</code>
Hierarchy Level	<code>[edit interfaces interface-name sonet-options aps]</code>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Perform a manual switch between the protect and working circuits. This statement is honored only if there are no higher-priority reasons to switch.
Options	<p>protect—Request that the circuit become the protect circuit.</p> <p>working—Request that the circuit become the working circuit.</p>
Usage Guidelines	See “Configuring Switching Between the Working and Protect Circuits” on page 866.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	force

required-depth

Syntax	<code>required-depth <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options mpls pop-all-labels], [edit interfaces <i>interface-name</i> sonet-options mpls pop-all-labels], [edit interfaces <i>interface-name</i> fastether-options mpls pop-all-labels], [edit interfaces <i>interface-name</i> gigether-options mpls pop-all-labels]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For passive monitoring on ATM and SONET/SDH interfaces only, specify the number of MPLS labels an incoming packet must have for the pop-all-labels statement to take effect. If you include the required-depth 1 statement, the pop-all-labels statement takes effect for incoming packets with one label only. If you include the required-depth 2 statement, the pop-all-labels statement takes effect for incoming packets with two labels only.
Options	<i>number</i> —Number of MPLS labels on incoming IP packets. Range: 1 through 2 labels. Default: If you omit this statement, the pop-all-labels statement takes effect for incoming packets with one or two labels. The default is equivalent to including the required-depth [1 2] statement.
Usage Guidelines	See “Removing MPLS Labels from Incoming Packets” on page 294 and “Removing MPLS Labels from Incoming Packets” on page 874.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Services Interfaces Configuration Guide</i>


restore-interval

Syntax	<code>restore-interval <i>number</i>;</code>
Hierarchy Level	[edit protocols protection-group ethernet-ring <i>ring-name</i>]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	Specify the wait time to restore the interval, in minutes.
Options	Range: 5 through 12 minutes
Usage Guidelines	See “Configuring Ethernet Ring Protection Switching” on page 799.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

retries

Syntax	<code>retries integer;</code>
Hierarchy Level	[edit protocols dot1x authenticator interface <i>interface-id</i>]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Set the number of tries after which the port remains in the wait state for <i>quiet-period</i> seconds before reattempting authentication.
Options	<i>integer</i> —Specify the number of retries. Range: 1 through 10 Default: 3 retries
Usage Guidelines	See “Configuring IEEE 802.1x Port-Based Network Access Control” on page 741
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	dot1x, interface (IEEE 802.1x), quiet-period

revertive

Syntax	<code>revertive;</code>
Hierarchy Level	[edit interfaces aeX aggregated-ether-options lacp link-protection]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Enable the ability to switch to a better priority link (if one is available).
<div style="display: flex; align-items: flex-start;"> <div style="flex: 1; text-align: center; padding-right: 10px;">  </div> <div> <p>NOTE: By default, LACP link protection is revertive. However, you can use this statement to define a specific aggregated Ethernet interface as revertive to override a global non-revertive statement specified at the [edit chassis] hierarchy.</p> </div> </div>	
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

revert-time

Syntax	<code>revert-time seconds;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure APS revertive mode.
Default	APS operates in nonrevertive mode.
Options	<p><i>seconds</i>—Amount of time to wait after the working circuit has again become functional before making the working circuit active again.</p> <p>Range: 1 through 65,535 seconds</p> <p>Default: none (APS operates in nonrevertive mode)</p>
Usage Guidelines	See “Configuring Revertive Mode” on page 867.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

rfc-2615

Syntax	<code>rfc-2615;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Include this statement to enable features described in RFC 2615, <i>PPP over SONET/SDH</i> .
Default	Settings required by RFC 1619, <i>PPP over SONET/SDH</i> .
Usage Guidelines	See “Configuring SONET/SDH RFC 2615 Support” on page 855.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

ring-protection-link-end

Syntax	ring-protection-link-end;
Hierarchy Level	[edit protocols protection-group ethernet-ring <i>ring-name</i> (east-interface west-interface)]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	<p>If this port is one side of ring protection link (RPL), the RPL end flag should be set.</p> <p>To set the RPL end, use the set ring-protection-link-end statement at the [edit protocols protection-group ethernet-ring <i>ring-name</i> (east-interface west-interface)] hierarchy level.</p> <p>To delete the RPL end, use the delete ring-protection-link-end statement at the [edit protocols protection-group ethernet-ring <i>ring-name</i> (east-interface west-interface)] hierarchy level.</p>
Usage Guidelines	See “Configuring Ethernet Ring Protection Switching” on page 799.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ring-protection-link-owner

Syntax	ring-protection-link-owner;
Hierarchy Level	[edit protocols protection-group ethernet-ring <i>ring-name</i>]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	<p>Use this statement to set the ring protection link (RPL) owner flag in the Ethernet protection ring. For each ring, only one node should be configured as a ring-protection-link-owner; so only one node in each ring should have the flag set.</p> <p>To set the RPL owner, use the set ring-protection-link-owner statement at the [edit protocols protection-group ethernet-ring <i>ring-name</i>] hierarchy level.</p> <p>To delete the RPL owner, use the delete ring-protection-link-owner statement at the [edit protocols protection-group ethernet-ring <i>ring-name</i>] hierarchy level.</p>
Usage Guidelines	See “Configuring Ethernet Ring Protection Switching” on page 799.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

routing-instance

Syntax	routing-instance { destination <i>routing-instance-name</i> ; bridge-domain <i>name</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel], [edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	To configure interfaces and logical-systems , specify the destination routing instance that points to the routing table containing the tunnel destination address. To configure protocols for the oam ethernet connectivity-fault-management maintenance-domain, specify the routing-instance <i>name</i> .
Default	The default Internet routing table is inet.0 .
Usage Guidelines	See <i>JUNOS Services Interfaces Configuration Guide</i> . See “Configuring Maintenance Intermediate Points” on page 683.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

rpf-check

Syntax	<pre>rpf-check { fail-filter <i>filter-name</i>; mode loose; }</pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Check whether traffic is arriving on an expected path. You can include this statement with the inet or inet6 protocol family only.</p> <p>The mode statement is explained separately.</p>
Options	fail-filter —A filter to evaluate when packets are received on the interface. If the RPF check fails, this optional filter is evaluated. If the fail filter is not configured, the default action is to silently discard the packet.
Usage Guidelines	See “Configuring Unicast RPF Strict Mode” on page 209 and “Configuring Unicast RPF Loose Mode” on page 210.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

rtp

Syntax rtp {
 f-max-period *number*;
 queues [*queue-numbers*];
 port {
 minimum *port-number*;
 maximum *port-number*;
 }
 }

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number* compression]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the real-time transport protocol (RTP) properties for voice services traffic.
 The remaining statements are described separately.

Usage Guidelines See the *JUNOS Services Interfaces Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

rts

Syntax rts (assert | de-assert | normal);

Hierarchy Level [edit interfaces *interface-name* serial-options dce-options],
 [edit interfaces *interface-name* serial-options dte-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description For EIA-530 and V.35 interfaces only, configure the to-DCE signal, request to send (RTS).

Options assert—The to-DCE signal must be asserted.
 de-assert—The to-DCE signal must be deasserted.
 normal—Normal RTS signal handling, as defined by the TIA/EIA Standard 530.
 Default: normal

Usage Guidelines See “Configuring the Serial Signal Handling” on page 271.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

rts-polarity

Syntax	rts-polarity (negative positive);
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure RTS signal polarity.
Options	negative—Negative signal polarity. positive—Positive signal polarity. Default: positive
Usage Guidelines	See “Configuring Serial Signal Polarities” on page 274.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

rtvbr

Syntax	<code>rtvbr peak <i>rate</i> sustained <i>rate</i> burst <i>length</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i> shaping],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping]</p>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For ATM2 IQ PICs only, define the real-time variable bandwidth utilization in the traffic-shaping profile.</p> <p>When you configure the real-time bandwidth utilization, you must specify all three options (burst, peak, and sustained). You can specify rate in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify rate in cells per second by entering a decimal number followed by the abbreviation c; values expressed in cells per second are converted to bits per second using the formula</p> $1 \text{ cps} = 384 \text{ bps}.$
Default	If the rtvbr statement is not included, bandwidth utilization is unlimited.
Options	<p>burst <i>length</i>—Burst length, in cells. If you set the length to 1, the peak traffic rate is used.</p> <p>Range: 1 through 4000 cells</p> <p>peak <i>rate</i>—Peak rate, in bits per second or cells per second.</p> <p>Range: For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure. For more information, see Table 27 on page 320.</p> <p>sustained <i>rate</i>—Sustained rate, in bps or cps.</p> <p>Range: For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, from 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure. For more information, see Table 27 on page 320.</p>
Usage Guidelines	See “Configuring ATM2 IQ Real-Time VBR” on page 321.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	cbr, vbr

sampling

Syntax	<code>sampling <i>direction</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the direction of traffic to be sampled.
Options	<p><i>direction</i> can be one of the following:</p> <p><code>input</code>—Configure at least one expected ingress point.</p> <p><code>output</code>—Configure at least one expected egress point.</p> <p><code>input output</code>—On a single interface, configure at least one expected ingress point and one expect egress point.</p>
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

satop-options

Syntax

```
satop-options {
    excessive-packet-loss-rate {
        apply-groups group-name
        apply-groups-except group-name
        groups group-name
        sample-period milliseconds
        threshold percentile
    }
    idle-pattern pattern
    jitter-buffer-auto-adjust
    jitter-buffer-latency milliseconds
    jitter-buffer-packets packets
    payload-size bytes;
}
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced in JUNOS Release 9.3.

Description Set Structure-Agnostic TDM over Packet (SATO-P) protocol options.

Options excessive-packet-loss-rate <options>—Packet loss options.

- apply-groups *group-name*—Groups from which to inherit configuration data.
- apply-groups-except *group-name*—Don't inherit configuration data from these groups.
- groups *group-name*—Specify groups.
- sample-period *milliseconds*—Number of milliseconds over which excessive packet loss rate is calculated.
- threshold *percentile*—Percentile designating the threshold of excessive packet loss rate (from 1 to 100).

idle-pattern *pattern*—An 8-bit hexadecimal pattern to replace TDM data in a lost packet (from 0 to 255).

jitter-buffer-auto-adjust—Automatically adjust the jitter buffer.

jitter-buffer-latency *milliseconds*—Number of milliseconds delay in jitter buffer (from 1 to 1000 milliseconds).

jitter-buffer-packets *packets*—Number of packets in jitter buffer (from 1 to 64).

payload-size *bytes*—Payload size in integer number of bytes.

Usage Guidelines See “Circuit Emulation PICs Overview” on page 519.

Required Privilege Level interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

scheduler-maps

Syntax	<code>scheduler-maps <i>map-name</i> { forwarding-class (<i>class-name</i> assured-forwarding best-effort expedited-forwarding network-control); vc-cos-mode (alternate strict); }</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, define CoS parameters assigned to forwarding classes.
Options	<i>map-name</i> —Name of the scheduler map. The remaining statements are explained separately.
Usage Guidelines	See “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	atm-scheduler-map, <i>JUNOS Class of Service Configuration Guide</i>

schedulers

Syntax	<code>schedulers <i>number</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 8.2.
Description	Specify the number of schedulers for Ethernet IQ2 and IQ2-E PIC port interfaces.
Default	If you omit this statement, the 1024 schedulers are distributed equally over all ports in multiples of 4.
Options	<i>number</i> —Number of schedulers to configure on the port. Range: 1 through 1024
Usage Guidelines	See the <i>JUNOS Class of Service Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

secondary

Syntax	secondary <i>interface-name</i> ;
Hierarchy Level	[edit interfaces (rsp0 rsp1) redundancy-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the secondary (backup) AS PIC interface or MultiServices PIC interface.
Options	<i>interface-name</i> —The identifier for the AS PIC interface or MultiServices PIC interface, which must be of the form <i>sp-fpc/pic/port</i> .
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

send-critical-event

Syntax	send-critical-event;
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile action]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Send OAM PDUs with the critical event bit set.
Usage Guidelines	See “Specifying the Actions to Be Taken for Link-Fault Management Events” on page 749.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

serial-options

Syntax

```
serial-options {
  clock-rate rate;
  clocking-mode (dce | loop);
  control-polarity (negative | positive);
  cts-polarity (negative | positive);
  dcd-polarity (negative | positive);
  dce-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
  }
  dsr-polarity (negative | positive);
  dte-options {
    control-signal (assert | de-assert | normal);
    cts (ignore | normal | require);
    dcd (ignore | normal | require);
    dsr (ignore | normal | require);
    dtr signal-handling-option;
    ignore-all;
    indication (ignore | normal | require);
    rts (assert | de-assert | normal);
    tm (ignore | normal | require);
  }
  dtr-circuit (balanced | unbalanced);
  dtr-polarity (negative | positive);
  encoding (nrz | nrzi);
  indication-polarity (negative | positive);
  line-protocol protocol;
  loopback (dce-local | dce-remote | local | remote);
  rts-polarity (negative | positive);
  tm-polarity (negative | positive);
  transmit-clock invert;
}
```

Hierarchy Level [edit interfaces *se-pim/0/port*]

Release Information Statement introduced prior to JUNOS Release 7.4.

Description Configure serial-specific interface properties.

The statements are explained separately.

Usage Guidelines See “Configuring Serial Interfaces” on page 263.

Required Privilege Level interface—To view this statement in the configuration.

interface-control—To add this statement to the configuration.

Related Topics no-concatenate in the *JUNOS System Basics Configuration Guide*

server

Syntax server;

Hierarchy Level [edit interfaces pp0 unit *logical-unit-number* pppoe-options],
[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number*
pppoe-options]

Release Information Statement introduced in JUNOS Release 8.5.

Description Configure the router to operate in the PPPoE server mode. Supported on M120 Multiservice Edge Routers operating as access concentrators.

Usage Guidelines See “Configuring the PPPoE Server Mode” on page 793.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

server-timeout

Syntax server-timeout *seconds*;

Hierarchy Level [edit protocols dot1x authenticator interface *interface-id*]

Release Information Statement introduced in JUNOS Release 9.3.

Description Specify the number of seconds the port waits for a response when relaying a request from the authentication server to the client before resending the request.

Options *seconds*—The number of seconds the port waits for a response when relaying a request from the authentication server to the client before resending the request.
Range: 1 through 60 seconds
Default: 30 seconds

Usage Guidelines See “Configuring IEEE 802.1x Port-Based Network Access Control” on page 741.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics dot1x, authenticator, interface (IEEE 802.1x)

service

Syntax	<pre> service { input { service-set service-set-name <service-filter filter-name>; post-service-filter filter-name; } output { service-set service-set-name <service-filter filter-name>; } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define one or more service sets and filters, and one postservice filter to be applied to an interface.
Options	The remaining statements are explained separately.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

service-domain

Syntax	service-domain (inside outside);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For adaptive services interfaces, specify a service interface domain. If you specify this interface using the <code>next-hop-service</code> statement at the [edit services service-set <i>service-set-name</i>] hierarchy level, the interface domain must match that used with the <code>inside-service-interface</code> and <code>outside-service-interface</code> statements.
Options	inside—Interface used within the network. outside—Interface used outside the network.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

service-filter

Syntax	<code>service-filter <i>filter-name</i>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service (input output)], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service (input output)]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define the filter to be applied to traffic before it is accepted for service processing. Configuration of a service filter is optional; if you include the service-set statement without a service-filter definition, the JUNOS Software assumes the match condition is true and selects the service set for processing automatically.
Options	<i>filter-name</i> —Identifies the filter to be applied in service processing.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

service-name

Syntax	<code>service-name <i>name</i>;</code>
Hierarchy Level	[edit interfaces pp0 unit <i>logical-unit-number</i> pppoe-options], [edit logical-systems <i>logical-system-name</i> interfaces pp0 unit <i>logical-unit-number</i> pppoe-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For J Series Services Routers with PPP over Ethernet interfaces, configure the service to be requested from the PPP over Ethernet server; that is, the access concentrator. For example, you can use this statement to indicate an Internet service provider (ISP) name or a class of service.
Options	<i>name</i> —Service to be requested from the PPP over Ethernet server.
Usage Guidelines	See “Configuring the PPPoE Service Name” on page 792.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

service-set

Syntax	<code>service-set service-set-name;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service (input output)], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet service (input output)]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Define one or more service sets to be applied to an interface. If you define multiple service sets, the JUNOS Software evaluates the filters in the order in which they appear in the configuration.
Options	<i>service-set-name</i> —Identifies the service set.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

services

Syntax	<code>services priority-level;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> services-options syslog host <i>hostname</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify system logging priority level.
Options	<i>priority-level</i> —Assigns a priority level to the facility. Valid entries are as follows: <ul style="list-style-type: none"> ■ alert—Conditions that should be corrected immediately ■ any—Matches any level. ■ emergency—Panic conditions. ■ critical—Critical conditions. ■ error—Error conditions. ■ info—Informational messages. ■ notice—Conditions that require special handling. ■ warning—Warning messages.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

services-options

Syntax

```

services-options {
  inactivity-timeout seconds;
  open-timeout seconds;
  syslog {
    host hostname {
      facility-override facility-name;
      log-prefix prefix-number;
      services priority-level;
    }
  }
}

```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define the service options to be applied on an interface.

Options The remaining statements are explained separately.

Usage Guidelines See the *JUNOS Services Interfaces Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

shaping

Syntax	<pre>shaping { (cbr rate rtvbr peak rate sustained rate burst length vbr peak rate sustained rate burst length); queue-length number; }</pre>
Hierarchy Level	<pre>[edit interfaces interface-name atm-options vpi vpi-identifier], [edit interfaces interface-name unit logical-unit-number], [edit interfaces interface-name unit logical-unit-number address address family family multipoint-destination address], [edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number], [edit logical-systems logical-system-name interfaces interface-name unit logical-unit-number address address family family multipoint-destination address]</pre>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For ATM encapsulation only, define the traffic-shaping profile.</p> <p>For ATM2 IQ interfaces, changing or deleting VP tunnel traffic shaping causes all logical interfaces on a VP to be deleted and then re-added.</p> <p>VP tunnels are not supported on multipoint interfaces.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Defining Virtual Path Tunnels” on page 316 and “Defining the ATM Traffic-Shaping Profile” on page 319.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

shdsl-options

Syntax	shdsl-options { annex (annex-a annex-b); line-rate <i>line-rate</i> ; loopback (local remote payload); snr-margin { current <i>margin</i> ; snext <i>margin</i> ; } }
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series Services Routers only, configure symmetric DSL (SHDSL) options. The statements are explained separately.
Usage Guidelines	See “Configuring ATM-over-SHDSL Interfaces” on page 361.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

short-name-format

Syntax	short-name-format (character-string vlan 2octet rfc-2685-vpn-id);
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> maintenance-association <i>ma-name</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Specify the name format of the maintenance association name.
Options	character-string—The name is an ASCII character string. vlan—The primary VLAN identifier. 2octet—A number in the range 0 through 65535. rfc-2685-vpn-id—A VPN identifier that complies with RFC 2685. Default: character-string
Usage Guidelines	See “Configuring the Maintenance Association Short Name Format” on page 685.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

short-sequence

Syntax	short-sequence;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For multilink interfaces only, set the length of the packet sequence identification number to 12 bits.
Default	If you omit this statement from the configuration, the length is set to 24 bits.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

snext

Syntax	snext <i>margin</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> shdsl-options snr-margin], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> shdsl-options snr-margin]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series Services Routers only, configure self-near-end crosstalk (SNEXT) signal-to-noise ratio (SNR) margin for a SHDSL line. When configured, the line trains at higher than SNEXT threshold. The SNR margin is the difference between the desired SNR and the actual SNR.
Options	<i>margin</i> —Desired SNEXT margin. Possible values are disabled or a margin between -10dB and 10 dB. Default: disabled
Usage Guidelines	See “Configuring ATM-over-SHDSL Interfaces” on page 361.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

snr-margin

Syntax	snr-margin { current <i>margin</i> ; snext <i>margin</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> shdsl-options], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> shdsl-options]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	<p>For J Series Services Routers only, configure the SHDSL signal-to-noise ratio (SNR) margin. The SNR margin is the difference between the desired SNR and the actual SNR. Configuring the SNR creates a more stable SHDSL connection by making the line train at a SNR margin higher than the threshold. If any external noise below the threshold is applied to the line, the line remains stable.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Configuring ATM-over-SHDSL Interfaces” on page 361.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

sonet-options

Syntax

```
sonet-options {
  aps {
    advertise-interval milliseconds;
    annex-b
    authentication-key key;
    force;
    hold-time milliseconds;
    lockout;
    neighbor address;
    paired-group group-name;
    protect-circuit group-name;
    request;
    revert-time seconds;
    switching-mode (bidirectional | unidirectional);
    working-circuit group-name;
  }
  bytes {
    c2 value;
    e1-quiet value;
    f1 value;
    f2 value;
    s1 value;
    z3 value;
    z4 value;
  }
  fcs (16 | 32);
  loopback (local | remote);
  mpls {
    pop-all-labels {
      required-depth number;
    }
  }
  path-trace trace-string;
  (payload-scrambler | no-payload-scrambler);
  rfc-2615;
  trigger {
    defect ignore;
    defect hold-time up milliseconds down milliseconds;
  }
}
vtmapping (itu-t | klm);
(z0-increment | no-z0-increment);
```

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure SONET/SDH-specific interface properties.

On SONET/SDH OC48 interfaces that you configure for channelized (multiplexed) mode (by including the `no-concatenate` statement at the `[edit chassis fpc slot-number pic pic-number]` hierarchy level), the `bytes e1-quiet` and `bytes f1` options have no effect. The `bytes f2`, `bytes z3`, `bytes z4`, and `path-trace` options work correctly on channel 0 and work in the transmit direction only on channels 1, 2, and 3.

On a channelized OC12 interface, the `bytes e1-quiet`, `bytes f1`, `bytes f2`, `bytes z3`, and `bytes z4` options are not supported. The `fcs` and `payload-scrambler` statements are also not supported; you must configure these for each DS3 channel using the `t3-options fcs` and `t3-options payload-scrambler` statements. The `aps` and `loopback` statements are supported only on channel 0 and are ignored if included in the configurations for channels 1 through 11. You can configure loopbacks for each DS3 channel with the `t3-options loopback` statement. The `path-trace` statement can be included in the configuration for each DS3 channel, thereby configuring a unique path trace for each channel.

To configure loopback on channelized IQ and IQE PICs, SONET/SDH level, use the `loopback` statement `local` and `remote` options at the controller interface (`coc48`, `cstm16`, `coc12`, `cstm4`, `coc3`, and `cstm1`). It is ignored for path-level interfaces `so-fpc/pic/port` or `so-fpc/pic/port:channel`.

If you are running Intermediate System-to-Intermediate System (IS-IS) over SONET/SDH interfaces, use PPP if you are running Cisco IOS Release 12.0 or later. If you need to run HDLC, configure an ISO family MTU of 4469 on the router.

The statements are explained separately.

Usage Guidelines	See “Configuring SONET/SDH Parameters on ATM Interfaces” on page 338, “Configuring Channelized OC12/STM4 Interfaces” on page 423, “Configuring Channelized STM1 Interfaces” on page 465, and “Configuring SONET/SDH Physical Interface Properties” on page 844.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<code>no-concatenate</code> in the <i>JUNOS System Basics Configuration Guide</i>

source

Syntax	source <i>source-address</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> tunnel <i>address</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the source address of the tunnel.
Default	If you do not specify a source address, the tunnel uses the unit's primary address as the source address of the tunnel.
Options	<i>source-address</i> —Address of the local side of the tunnel. This is the address that is placed in the outer IP header's source field.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	multicast-only, primary

source-address-filter

Syntax	source-address-filter { <code>mac-address</code> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> aggregated-ether-options], [edit interfaces <i>interface-name</i> fastether-options], [edit interfaces <i>interface-name</i> ggether-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), specify the MAC addresses from which the interface can receive packets. For this statement to have any effect, you must include the source-filtering statement in the configuration to enable source address filtering. This statement is not supported on the J Series Services Routers.
Options	<p>mac-address—MAC address filter. You can specify the MAC address as <code>nn:nn:nn:nn:nn:nn</code> or <code>nnnn.nnnn.nnnn</code>, where <i>n</i> is a decimal digit. To specify more than one address, include multiple mac-address options in the source-address-filter statement.</p> <p>If you enable the VRRP on a Fast Ethernet or Gigabit Ethernet interface, as described in “Configuring VRRP and VRRP for IPv6” on page 753, and if you enable MAC source address filtering on the interface, you must include the virtual MAC address in the list of source MAC addresses that you specify in the source-address-filter statement. MAC addresses ranging from <code>00:00:5e:00:01:00</code> through <code>00:00:5e:00:01:ff</code> are reserved for VRRP, as defined in RFC 3768, <i>Virtual Router Redundancy Protocol</i>. When you configure the VRRP group, the group number must be the decimal equivalent of the last hexadecimal byte of the virtual MAC address.</p> <p>On untagged Gigabit Ethernet interfaces you should not configure the source-address-filter statement and the accept-source-mac statement simultaneously. On tagged Gigabit Ethernet interfaces you should not configure the source-address-filter statement and the accept-source-mac statement with an identical MAC address specified in both filters.</p>
Usage Guidelines	See “Enabling Ethernet MAC Address Filtering” on page 591.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	source-filtering

source-class-usage

Syntax	source-class-usage { <i>direction</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet accounting]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable packet counters on an interface that count packets that arrive from specific prefixes on the provider core router and are destined for specific prefixes on the customer edge router.
Options	<i>direction</i> can be one of the following: input—Configure at least one expected ingress point. output—Configure at least one expected egress point. input output—On a single interface, configure at least one expected ingress point and one expect egress point.
Usage Guidelines	See “Enabling Source Class and Destination Class Usage” on page 214 or the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	accounting, destination-class-usage

source-filtering

Syntax	(source-filtering no-source-filtering);
Hierarchy Level	[edit interfaces <i>interface-name</i> aggregated-ether-options], [edit interfaces <i>interface-name</i> fastether-options], [edit interfaces <i>interface-name</i> gigether-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For aggregated Ethernet, Fast Ethernet, Gigabit Ethernet, and Gigabit Ethernet IQ interfaces only, enable the filtering of MAC source addresses, which blocks all incoming packets to that interface. To allow the interface to receive packets from specific MAC addresses, include the source-address-filter statement.</p> <p>If the remote Ethernet card is changed, the interface will no longer be able to receive packets from the new card because it will have a different MAC address.</p>
Default	Source address filtering is disabled.
Usage Guidelines	See “Enabling Ethernet MAC Address Filtering” on page 591.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	accept-source-mac, source-address-filter

speed

See the following sections:

- speed (Ethernet) on page 1255
- speed (MX Series DPC) on page 1256
- speed (SONET/SDH) on page 1257

speed (Ethernet)

Syntax	speed (10m 100m 1g auto);
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>ge-pim</i> /0/0 switch-options switch-port <i>port-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the interface speed. This statement applies to the management Ethernet interface (fxp0 or em0), Fast Ethernet 12-port and 48-port PICs, the built-in Fast Ethernet port on the FIC (M7i router), the built-in Ethernet interfaces on J Series Services Routers, Combo Line Rate DPCs and Tri-Rate Ethernet Copper interfaces on MX Series routers, and on the Gigabit Ethernet ports on J Series Services Routers with uPIMs installed and configured for access switching mode. When you configure the Tri-Rate Ethernet copper interface to operate at 1 Gbps, autonegotiation must be enabled. When you configure 100BASE-FX SFP, you must set the port speed at 100 Mbps.
Options	You can specify the speed as either 10m (10 Mbps), 100m (100 Mbps), or on J Series routers with uPIMs installed and on MX Series routers, 1g (1 Gbps). You can specify the auto option only on MX Series routers.
Usage Guidelines	See “Configuring the Interface Speed” on page 122, “Configuring the Interface Speed on Ethernet Interfaces” on page 597, “Configuring Gigabit Ethernet Autonegotiation” on page 767, and “Configuring J Series Services Router Switching Interfaces” on page 589.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

speed (MX Series DPC)

Syntax speed (auto | 1Gbps | 100Mbps | 10Mbps);

Hierarchy Level [edit interfaces *ge-fpc/pic/port*]

Release Information Statement introduced in JUNOS Release 9.5.

Description On MX Series routers with Combo Line Rate DPCs and Tri-Rate Copper SFPs you can set auto negotiation of speed. To specify the auto negotiation speed, use the **speed** (<auto | 1Gbps | 100Mbps | 10Mbps>) statement under the [edit interface *ge-/fpc/pic/port*] hierarchy level. The <auto> option will attempt to automatically match the rate of the connected interface. To set port speed negotiation to a specific rate, set the port speed to 1Gbps, 100Mbps, or 10Mbps.



NOTE: If the negotiated speed and the interface speed do not match, the link will not be brought up. Half duplex mode is not supported.

You can disable auto MDI/MDIX using the **no-auto-mdix** statement option at the [edit interface *ge-fpc/pic/port* *gether-options*] hierarchy level.

Options You can specify the speed as either **auto** (autonegotiate), **10Mbps** (10 Mbps), **100Mbps** (100 Mbps), or **1Gbps** (1 Gbps).

Usage Guidelines See “Configuring Gigabit Ethernet Autonegotiation” on page 767.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics no-auto-mdix

speed (SONET/SDH)

Syntax	<code>speed (oc3 oc12 oc48);</code>
Hierarchy Level	[edit interfaces <i>so-fpc/pic/port</i>], [edit interfaces <i>so-fpc/pic/port:channel</i>]
Release Information	Statement introduced in JUNOS Release 8.3.
Description	Configure the interface speed. This statement applies to SONET/SDH interfaces on next-generation SONET/SDH Type 1 and Type 2 PICs with SFP. Available speeds depend on whether the PIC is in concatenated mode or nonconcatenated mode. Include the channel in the interface name when configuring nonconcatenated interfaces.
Options	<code>oc3 oc12 oc48</code> —Speed when the PIC is in concatenated mode. For example, you can configure each port of a 4-port OC12 PIC to have a speed of <code>oc3</code> . You can configure port 0 of a 4-port OC12 PIC to have a speed of <code>oc12</code> . <code>oc3 oc12</code> —Speed when the PIC is in nonconcatenated mode.
Usage Guidelines	See “Configuring SONET/SDH Interface Speed” on page 847.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

spid1

Syntax	<code>spid1 <i>spid1-string</i>;</code>
Hierarchy Level	[edit interfaces <i>br-pim/0/port</i> isdn-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the Service Profile Identifier (SPID).
Options	<i>spid1-string</i> —Numeric SPID.
Usage Guidelines	See “Configuring ISDN Physical Interface Properties” on page 821.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

spid2

Syntax	<code>spid2 spid2-string;</code>
Hierarchy Level	[edit interfaces <i>br-pim</i> /0/ <i>port</i> isdn-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure an additional SPID.
Options	<i>spid2-string</i> —Numeric SPID.
Usage Guidelines	See “Configuring ISDN Physical Interface Properties” on page 821.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J Series Services Router Configuration Guide</i>


stacked-vlan-ranges

Syntax	<pre> stacked-vlan-ranges { dynamic-profile <i>profile-name</i> { accept (inet); ranges (Dynamic Stacked VLAN) (any <i>low-tag - high-tag</i>) , (any <i>low-tag - high-tag</i>); } } </pre>
Hierarchy Level	[edit interfaces <i>interface-name</i> auto-configure]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.
Options	<p><i>any</i>—Any valid VLAN ID number.</p> <p><i>vlan-id-low</i>—Specify the first VLAN ID number for the group of VLANs.</p> <p><i>vlan-id-high</i>—Specify the last VLAN ID number for the group of VLANs.</p> <p>Range: 1 through 4094</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See Configuring Stacked VLAN Ranges for Use with Stacked VLAN Dynamic Profiles and Configuring Dynamic Mixed VLAN Ranges.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

stacked-vlan-tagging

Syntax	stacked-vlan-tagging;
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Gigabit Ethernet IQ interfaces, enable stacked VLAN tagging for all logical interfaces on the physical interface.
Usage Guidelines	See “Configuring the Management Ethernet Interface” on page 775.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	vlan-tags (Stacked VLAN Tags)

start-end-flag

Syntax	start-end-flag (filler shared);
Hierarchy Level	[edit interfaces <i>e1-fpc/pic/port</i>], [edit interfaces <i>t1-fpc/pic/port</i>], [edit interfaces <i>interface-name</i> ds0-options], [edit interfaces <i>interface-name</i> e1-options], [edit interfaces <i>interface-name</i> e3-options], [edit interfaces <i>interface-name</i> t1-options], [edit interfaces <i>interface-name</i> t3-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For DS0, E1, E3, T1, and T3 interfaces, configure the interface to share the transmission of start and end flags.
	NOTE: When configuring E1 or T1 interfaces on the 10-port Channelized E1/T1 IQE PIC, the start-end-flag statement must be included at the [edit interfaces <i>e1-fpc/pic/port</i>] or [edit interfaces <i>t1-fpc/pic/port</i>] hierarchy level as appropriate.
Options	<p>filler—Wait two idle cycles between the start and end flags.</p> <p>shared—Share the transmission of the start and end flags. This is the default.</p>
Usage Guidelines	See “Configuring E1 Start and End Flags” on page 548, “Configuring the E3 Start and End Flags” on page 557, “Configuring T1 Start and End Flags” on page 567, and “Configuring T3 Start and End Flags” on page 579.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

static-tei-val

Syntax	<code>static-tei-val value;</code>
Hierarchy Level	[edit interfaces <i>br-pim</i> /0/ <i>port</i> isdn-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For J Series Services Routers only. Statically configure the Terminal Endpoint Identifier (TEI) value. The TEI value represents any ISDN-capable device attached to an ISDN network that is the terminal endpoint. TEIs are used to distinguish between several different devices using the same ISDN links.
Options	<i>value</i> —Value between 0 through 63.
Usage Guidelines	See “Configuring ISDN Physical Interface Properties” on page 821.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

supplicant

Syntax	<code>supplicant single;</code>
Hierarchy Level	[edit protocols dot1x authenticator interface <i>interface-id</i>]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specify the supplicant mode. Only single mode is supported. This option will authenticate only the first client that connects to a port. All other clients that connect later (802.1x compliant or non-compliant) will be allowed free access on that port without any further authentication. If the first authenticated client logs out, all other users are locked out until a client authenticates again.
Options	<i>single</i> —Sets single mode.
Usage Guidelines	See “Configuring IEEE 802.1x Port-Based Network Access Control” on page 741.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	dot1x, authenticator, interface (IEEE 802.1x)

supplicant-timeout

Syntax	supplicant-timeout <i>seconds</i> ;
Hierarchy Level	[edit protocols dot1x authenticator interface <i>interface-id</i>]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Specify the number of seconds the port waits for a response when relaying a request from the authentication server to the client before resending the request.
Options	<i>seconds</i> —Specify the number of seconds the port waits for the supplicant timeout. Range: 1 through 60 seconds Default: 30 seconds
Usage Guidelines	See “Configuring IEEE 802.1x Port-Based Network Access Control” on page 741.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	dot1x, authenticator, interface (IEEE 802.1x)

swap

Syntax	swap;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and aggregated Ethernet using Gigabit Ethernet IQ interfaces, specify the VLAN rewrite operation to replace a VLAN tag. The outer VLAN tag of the frame is overwritten with the user-specified VLAN tag information.
Usage Guidelines	See “Rewriting the VLAN Tag on Tagged Frames” on page 651.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

swap-push

Syntax	swap-push;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, specify the VLAN rewrite operation to replace the outer VLAN tag of the frame with a user-specified VLAN tag value. A user-specified outer VLAN tag is pushed in front. The outer tag becomes an inner tag in the final frame.
Usage Guidelines	See “Rewriting a VLAN Tag and Adding a New Tag” on page 655.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

swap-swap

Syntax	swap-swap;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced in JUNOS Release 8.1.
Description	For Gigabit Ethernet IQ, IQ2 and IQ2-E interfaces, and for aggregated Ethernet interfaces using Gigabit Ethernet IQ2 and IQ2-E or 10-Gigabit Ethernet PICs on MX Series routers, specify the VLAN rewrite operation to replace both the inner and the outer VLAN tags of the frame with a user-specified VLAN tag value.
Usage Guidelines	See “Rewriting the Inner and Outer VLAN Tags” on page 655.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

switching-mode

Syntax	switching-mode (bidirectional unidirectional);
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options aps]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For unchannelized OC3, OC12, and OC48 SONET/SDH interfaces on T Series routers only, configure the interface to interoperate with SONET/SDH line-terminating equipment (LTE) that is provisioned for unidirectional linear APS in 1 + 1 architecture.
Default	If the switching-mode statement is not configured, the mode is bidirectional, and the interface does not interoperate with a unidirectional SONET/SDH LTE.
Options	<p>bidirectional—Support bidirectional mode only.</p> <p>unidirectional—Interoperate with a SONET/SDH LTE provisioned for unidirectional mode.</p>
Usage Guidelines	See “Configuring Unidirectional Switching Mode Support” on page 867.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

switch-options

Syntax	<pre>switch-options { switch-port <i>port-number</i> { (auto-negotiation no-auto-negotiation); speed (10m 100m 1g); link-mode (full-duplex half-duplex); } }</pre>
Hierarchy Level	[edit interfaces <i>ge-pim/0/0</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	On a J Series Services Router with multiport Gigabit Ethernet uPIMs installed and operating in access switching mode, only one physical interface is configured for the entire multiport Gigabit Ethernet uPIM. Configuration of the physical port characteristics is done under the single physical interface.
Usage Guidelines	See “Configuring J Series Services Router Switching Interfaces” on page 589.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

switch-port

Syntax switch-port *port-number* {
 (auto-negotiation | no-auto-negotiation);
 speed (10m | 100m | 1g);
 link-mode (full-duplex | half-duplex);
 }

Hierarchy Level [edit interfaces *ge-pim*/0/0 switch-options]

Release Information Statement introduced in JUNOS Release 8.4.

Description On a J Series Services Router with Ethernet uPIMs installed and operating in access switching mode, configuration of the physical port characteristics, done under the single physical interface.

Default Autonegotiation is enabled by default. If the link speed and duplex are also configured, the interfaces use the values configured as the desired values in the negotiation.

Options *port-number*—Ports are numbered 0 through 5 on the 6-port Gigabit Ethernet uPIM, 0 through 7 on the 8-port Gigabit Ethernet uPIM, and 0 through 15 on the 16-port Gigabit Ethernet uPIM.

The remaining statements are explained separately.

Usage Guidelines See “Configuring J Series Services Router Switching Interfaces” on page 589.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

switch-type

Syntax	switch-type (att5e etsi ni1 ntdms-100)
Hierarchy Level	[edit interfaces br-pim/0/port isdn-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For J Series Services Routers only. Configure the ISDN variant supported.
Options	<p>att5e—AT&T switch variant.</p> <p>etsi—European Telecommunications Standards Institute switch variant.</p> <p>ni1—National ISDN 1 switch variant.</p> <p>ntdms-100—Northern Telecom DMS-100.</p> <p>ntt—NTT Group switch for Japan.</p>
Usage Guidelines	See “Configuring ISDN Interfaces” on page 819.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

symbol-period

Syntax	symbol-period <i>count</i> ;
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile event, link-event-rate], [edit protocols oam link-fault-management interface <i>interface-name</i> event-thresholds]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	<p>Configure the threshold for sending symbol period events or taking the action specified in the action profile.</p> <p>A symbol error is any symbol code error on the underlying physical layer. The symbol period threshold is reached when the number of symbol errors reaches the configured value within the period window. The default period window is the number of symbols that can be transmitted on the underlying physical layer in 1 second. The window is not configurable.</p>
Options	<i>count</i> —Threshold count for symbol period events. Range: 1 through 100
Usage Guidelines	See “Configuring Threshold Values for Local Fault Events on an Interface” on page 747 and “Configuring Threshold Values for Fault Events in an Action Profile” on page 750.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

syslog

See the following sections:

- syslog (Interfaces) on page 1268
- syslog (Monitoring) on page 1269
- syslog (OAM Action) on page 1269

syslog (Interfaces)

Syntax

```
syslog {
  host hostname {
    facility-override facility-name;
    log-prefix prefix-number;
    services priority-level;
  }
}
```

Hierarchy Level [edit interfaces *interface-name* services-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description For adaptive services interfaces, configure generation of system log messages for the service set. System log information is passed to the kernel for logging in the /var/log directory. Any values configured in the service set definition override these values.

Options The remaining statements are explained separately.

Usage Guidelines See the *JUNOS Services Interfaces Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

syslog (Monitoring)

Syntax	(syslog no-syslog);
Hierarchy Level	[edit interfaces mo-fpc/pic/port multiservice-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>System logging is enabled by default. The system log information of the Monitoring Services PIC is passed to the kernel for logging in the <code>/var/log</code> directory.</p> <ul style="list-style-type: none"> ■ syslog—Enable PIC system logging. ■ no-syslog—Disable PIC system logging.
Usage Guidelines	See “Configuring Multiservice Physical Interface Properties” on page 138 or the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

syslog (OAM Action)

Syntax	syslog;
Hierarchy Level	[edit protocols oam ethernet link-fault-management action-profile action]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Generate a syslog message for the Ethernet Operation, Administration, and Management (OAM) event.
Usage Guidelines	See “Specifying the Actions to Be Taken for Link-Fault Management Events” on page 749.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

system-priority

Syntax	system-priority <i>priority</i> ;
Hierarchy Level	[edit interfaces aeX aggregated-ether-options lacp]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Define LACP system priority at the aggregated Ethernet interface level. This system priority value takes precedence over a system priority value configured at the global ([edit chassis]) hierarchy level.
Options	<i>priority</i> —Priority for the aggregated Ethernet system. A smaller value indicates a higher priority. Range: 0 through 65535 Default: 127
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

t1-options

Syntax t1-options {
 bert-algorithm *algorithm*;
 bert-error-rate *rate*;
 bert-period *seconds*;
 buildout *value*;
 byte-encoding (nx56 | nx64);
 crc-major-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5);
 crc-minor-alarm-threshold (1e-3 | 5e-4 | 1e-4 | 5e-5 | 1e-5 | 5e-6 | 1e-6);
 fcs (16 | 32);
 framing (esf | sf);
 idle-cycle-flag (flags | ones);
 invert-data;
 line-encoding (ami | b8zs);
 loopback (local | payload | remote);
 remote-loopback-respond;
 start-end-flag (filler | shared);
 timeslots *time-slot-range*;
 }

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure T1-specific physical interface properties.

The statements are explained separately.

Usage Guidelines See “Configuring T1 Interfaces” on page 559.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

t1-time

Syntax	t1-time <i>time</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series Services Routers only. On Ethernet interfaces configured for DLSw, configure the length of time the router waits for an acknowledgment of transmitted frames.
Options	<i>time</i> —Number of milliseconds. Range: 1 through 60000 Default: 1000 milliseconds
Usage Guidelines	See “Configuring LLC2 Options” on page 180.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

t2-time

Syntax	t2-time <i>time</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series Services Routers only. On Ethernet interfaces configured for DLSw, configure the length of time the router withholds the I-frame response.
Options	<i>time</i> —Number of milliseconds. Range: 1 through 60000 Default: 100 milliseconds
Usage Guidelines	See “Configuring LLC2 Options” on page 180.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

t310

Syntax	t310-value <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>br-pim/0/port</i> isdn-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ISDN interfaces, configure the Q.931-specific timer for T310, in seconds. The Q.931 protocol is involved in the setup and termination of connections.
Options	<i>seconds</i> —Timer value, in seconds. Range: 1 through 65536 Default: 10 seconds
Usage Guidelines	See “Configuring ISDN Physical Interface Properties” on page 821.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

t391

Syntax	t391 <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services and voices interfaces only, set Frame Relay link integrity polling interval.
Options	<i>seconds</i> —Link integrity polling interval. Range: 5 through 30 seconds Default: 10 seconds
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	n391, n392, n393, t392

t392

Syntax	t392 <i>seconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services and voices interfaces only, set Frame Relay polling verification interval.
Options	<i>seconds</i> —Polling verification interval. Range: 5 through 30 seconds Default: 15 seconds
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	n391, n392, n393, timeslots

t3-options

Syntax t3-options {
 atm-encapsulation (direct | plcp);
 bert-algorithm *algorithm*;
 bert-error-rate *rate*;
 bert-period *seconds*;
 (cbit-parity | no-cbit-parity);
 compatibility-mode (digital-link | kentrox | larscom) <subrate *value*>;
 fcs (16 | 32);
 (feac-loop-respond | no-feac-loop-respond);
 idle-cycle-flag *value*;
 (long-buildout | no-long-buildout);
 (loop-timing | no-loop-timing);
 loopback (local | payload | remote);
 start-end-flag *value*;
 }

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure T3-specific physical interface properties, including the properties of DS3 channels on a channelized OC12 interface. The **long-buildout** statement is not supported for DS3 channels on a channelized OC12 interface.

On T3 interfaces, the default encapsulation is PPP.

For ATM1 interfaces, you can configure a subset of E3 options statements.

The statements are explained separately.

Usage Guidelines See “Configuring T3 Interfaces” on page 569.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

tag-protocol-id

See the following sections:

- tag-protocol-id (TPIDs Expected to Be Sent or Received) on page 1276
- tag-protocol-id (TPID to Rewrite) on page 1277

tag-protocol-id (TPIDs Expected to Be Sent or Received)

Syntax	tag-protocol-id [<i>tpids</i>];
Hierarchy Level	[edit interfaces <i>interface-name</i> ggether-options ethernet-switch-profile], [edit interfaces <i>interface-name</i> aggregated-ether-options ethernet-switch-profile]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces, aggregated Ethernet with Gigabit Ethernet IQ interfaces, and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC, and the built-in Gigabit Ethernet port on the M7i router), define the TPIDs expected to be sent or received on a particular VLAN. For each Gigabit Ethernet port, you can configure up to eight TPIDs using the tag-protocol-id statement; but only the first four TPIDs are supported on IQ2 and IQ2-E interfaces.
Options	<i>tpids</i> —TPIDs to be accepted on the VLAN. Specify TPIDs in hexadecimal.
Usage Guidelines	See “Configuring Frames with Particular TPIDs to Be Processed as Tagged Frames” on page 644.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

tag-protocol-id (TPID to Rewrite)

Syntax	tag-protocol-id <i>tpid</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces only, configure the outer TPID value. All TPIDs you include in input and output VLAN maps must be among those you specify at the [edit interfaces <i>interface-name</i> <i>gigether-options</i> ethernet-switch-profile tag-protocol-id [<i>tpids</i>]] hierarchy level.
Default	If the tag-protocol-id statement is not configured, the TPID value is 0x8100.
Usage Guidelines	See “Configuring Inner and Outer TPIDs and VLAN IDs” on page 645.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

tei-option

Syntax	tei-option (first-call power-up);
Hierarchy Level	[edit interfaces <i>br-pim</i> /0/ <i>port</i> isdn-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ISDN interfaces, configure when the Terminal Endpoint Identifier (TEI) negotiates with the ISDN provider.
Options	first-call—Activation does not occur until the call setup is sent. power-up—Activation occurs when the Services Router is powered on. Default: power-up
Usage Guidelines	See “Configuring ISDN Physical Interface Properties” on page 821.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

then

Syntax	then { discard; }
Hierarchy Level	[edit firewall hierarchical-policer aggregate], [edit firewall hierarchical-policer premium]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	On M40e, M120, and M320 (with FFPC and SFPC) edge routers and T320, T640, and T1600 core routers with Enhanced Intelligent Queuing (IQE) PICs, discard packets when a specified bandwidth or burst limits for an aggregate level of a hierarchical policer is reached.
Options	discard—Discard packets if condition is met.
Usage Guidelines	See “Applying Policers” on page 194 and the <i>JUNOS Class of Service Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

threshold

Syntax	threshold bytes;
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Specify the bucket threshold, which controls the burstiness of the leaky bucket mechanism. The larger the value, the more bursty the traffic, which means that over a very short amount of time, the interface can receive or transmit close to line rate, but the average over a longer time is at the configured bucket rate.
Options	bytes—Maximum size, in bytes, for traffic bursts. Range: 0 through 65,535 bytes.
Usage Guidelines	See “Configuring Receive and Transmit Leaky Bucket Properties” on page 129 and “Configuring Receive and Transmit Leaky Bucket Properties on SONET/SDH Interfaces” on page 876. For ease of entry, you can enter <i>number</i> either as a complete decimal number or as a decimal number followed by the abbreviation k (1000). For example, the entry threshold 2k corresponds to a threshold of 2000 bytes.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

timeslots

Syntax `timeslots time-slot-range;`

Hierarchy Level [edit interfaces *e1-fpc/pic/port*],
[edit interfaces *t1-fpc/pic/port*],
[edit interfaces *interface-name* *e1-options*],
[edit interfaces *interface-name* *partition partition-number*],
[edit interfaces *interface-name* *t1-options*]

Release Information Statement introduced before JUNOS Release 7.4.

Description For E1 and T1 interfaces, allocate the specific time slots by number.



NOTE: When configuring E1 or T1 interfaces on the 10-port Channelized E1/T1 IQE PIC, the `timeslots` statement must be included at the [edit interfaces *e1-fpc/pic/port*] or [edit interfaces *t1-fpc/pic/port*] hierarchy level as appropriate.

Options *time-slot-range*—Actual time slot numbers allocated:

Range:

Ranges vary by interface type and configuration option as follows:

- 1 through 24 for T1 interfaces (0 is reserved)
- 1 through 31 for 4-port E1 PICs (0 is reserved)
- 1 through 31 for NxDS0 interfaces (0 is reserved)
- 2 through 32 for 10-port Channelized E1 and 10-port Channelized E1 IQ PICs (1 is reserved)
- 2 through 32 for the setting under **e1-options** with IQE PICs (1 is reserved) (when creating fractional E1)
- 1 through 31 for the setting under **partition** with IQE PICs (0 is reserved) (when creating NxDS0)



NOTE: When creating fractional E1 interfaces only, if you connect a 4-port E1 PIC interface to a device that uses time slot numbering from 2 through 32, you must subtract 1 from the configured number of time slots.

Usage Guidelines See “Configuring Fractional E1 IQ and IQE Interfaces” on page 502, “Configuring Fractional T1 IQ and IQE Interfaces” on page 480, “Configuring Fractional E1 Time Slots” on page 548, “Configuring Fractional T1 Time Slots” on page 567, and configuring “Configuring a Channelized T1/E1 Interface to Drop and Insert Time Slots” on page 510.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

tm

Syntax tm (ignore | normal | require);

Hierarchy Level [edit interfaces *interface-name* serial-options dce-options],
 [edit interfaces *interface-name* serial-options dte-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description For EIA-530 interfaces only, configure the from-DCE signal, test-mode (TM).

Options ignore—The from-DCE signal is ignored.

normal—Normal TM signal handling as defined by the TIA/EIA Standard 530.

require—The from-DCE signal must be asserted.
Default: normal

Usage Guidelines See “Configuring the Serial Signal Handling” on page 271.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

tm-polarity

Syntax tm-polarity (negative | positive);

Hierarchy Level [edit interfaces *interface-name* serial-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure TM signal polarity.

Options negative—Negative signal polarity.

positive—Positive signal polarity.
Default: positive

Usage Guidelines See “Configuring Serial Signal Polarities” on page 274.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

traceoptions

See the following sections:

- traceoptions (Individual Interfaces) on page 1282
- traceoptions (Interface Process) on page 1283
- traceoptions (LACP) on page 1285
- traceoptions (PPP Process) on page 1287



NOTE: For information about the `traceoptions` statement at the `[edit protocols vrrp]` hierarchy level, see the *JUNOS High Availability Configuration Guide*.

traceoptions (Individual Interfaces)

Syntax `traceoptions {
 flag flag;
 }`

Hierarchy Level [edit interfaces *interface-name*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define tracing operations for individual interfaces.

To specify more than one tracing operation, include multiple **flag** statements.

The interfaces **traceoptions** statement does not support a trace file. The logging is done by the kernel, so the tracing information is placed in the system **syslog** file in the directory `/var/log`.

Default If you do not include this statement, no interface-specific tracing operations are performed.

Options *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the interface-specific tracing options.

- **all**—All interface tracing operations
- **event**—Interface events
- **ipc**—Interface interprocess communication (IPC) messages
- **media**—Interface media changes
- **q921**—Trace ISDN Q.921 frames
- **q931**—Trace ISDN Q.931 frames

Usage Guidelines See “Tracing Operations of an Individual Router Interface” on page 241.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

traceoptions (Interface Process)

Syntax traceoptions {
 file <filename> <files number> <match regular-expression> <size size> <world-readable |
 no-world-readable>;
 flag flag <disable>;
 no-remote-trace;
 }

Hierarchy Level [edit interfaces]

Release Information Statement introduced before JUNOS Release 7.4.

Description Define tracing operations for the interface process (dcd).

Default If you do not include this statement, no interface-specific tracing operations are performed.

Options disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.

filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory */var/log*. By default, interface process tracing output is placed in the file *dcd*.

files number—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file.0*, then *trace-file.1*, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the *size* option.

Range: 2 through 1000

Default: 3 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple *flag* statements. You can include the following flags:

- *change-events*—Log changes that produce configuration events
- *config-states*—Log the configuration state machine changes
- *kernel*—Log configuration IPC messages to kernel
- *kernel-detail*—Log details of configuration messages to kernel

no-world-readable—(Optional) Disallow any user to read the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file.0*. When the *trace-file* again reaches its maximum size, *trace-file.0* is renamed *trace-file.1* and *trace-file* is renamed *trace-file.0*. This

renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option.

Syntax: **xk** to specify kilobytes, **xm** to specify megabytes, or **xg** to specify gigabytes

Range: 10 KB through the maximum file size supported on your router

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

match regex—(Optional) Refine the output to include only those lines that match the given regular expression.

Usage Guidelines See “Tracing Operations of the Interface Process” on page 241.

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

traceoptions (LACP)

Syntax traceoptions {
 file <filename> <files number> <size size> <world-readable | no-world-readable>;
 flag flag;
 no-remote-trace;
 }

Hierarchy Level [edit protocols lacp]

Release Information Statement introduced in JUNOS Release 7.6.

Description Define tracing operations for the LACP protocol.

Default If you do not include this statement, no LACP protocol tracing operations are performed.

Options disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

filename—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, interface process tracing output is placed in the file **lacpd**.

files number—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed **trace-file.0**, then **trace-file.1**, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option.

Range: 2 through 1000

Default: 3 files

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. You can include the following flags:

- **all**—All LACP tracing operations
- **configuration**—Configuration code
- **packet**—Packets sent and received
- **process**—LACP process events
- **protocol**—LACP protocol state machine
- **routing-socket**—Routing socket events
- **startup**—Process startup events

no-world-readable—(Optional) Prevent any user from reading the log file.

size size—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file.0*. When the *trace-file* again reaches its maximum size, *trace-file.0* is renamed *trace-file.1* and *trace-file* is renamed *trace-file.0*. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option:

Syntax: **xk** to specify kilobytes, **xm** to specify megabytes, or **xg** to specify gigabytes

Range: 10 KB through the maximum file size supported on your router

Default: 1 MB

world-readable—(Optional) Allow any user to read the log file.

Usage Guidelines See “Tracing LACP Operations” on page 631.

Required Privilege Level **interface**—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

traceoptions (PPP Process)

Syntax traceoptions {
 file *filename* <files *number*> <match *regular-expression*> <size *size*> <world-readable |
 no-world-readable>;
 flag *flag*;
 level *severity-level*;
 no-remote-trace;
 }

Hierarchy Level [edit protocols ppp]

Release Information Statement introduced in JUNOS Release 7.5.

Description Define tracing operations for the PPP process.

To specify more than one tracing operation, include multiple **flag** statements.

You cannot specify a separate trace tile. Tracing information is placed in the system syslog file in the directory `/var/log/pppd`.

Default If you do not include this statement, no PPPD-specific tracing operations are performed.

Options *filename*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory `/var/log`. By default, commit script process tracing output is placed in the file `pppd`. If you include the **file** statement, you must specify a filename. To retain the default, you can specify **eventd** as the filename.

files number—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file.0*, then *trace-file.1*, and so on, until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum number of files, you also must specify a maximum file size with the **size** option and a filename.

Range: 2 through 1000

Default: 3 files

disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

flag—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements. The following are the PPPD-specific tracing options.

- **access**—Access code
- **address-pool**—Address pool code
- **all**—All areas of code

- **auth**—Authentication code
- **chap**—Challenge Handshake Authentication Protocol (CHAP) code
- **config**—Configuration code
- **ifdb**—Interface database code
- **lcp**—LCP state machine code
- **memory**—Memory management code
- **message**—Message processing code
- **mlppp**—Trace MLPPP code
- **ncp**—NCP state machine code
- **pap**—Password Authentication Protocol (PAP) code
- **ppp**—PPP protocol processing code
- **radius**—RADIUS processing code
- **rtsock**—Routing socket code
- **session**—Session management code
- **signal**—Signal handling code
- **timer**—Timer code
- **ui**—User interface code

match *regex*—(Optional) Refine the output to include only those lines that match the given regular expression.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file.0***. When the ***trace-file*** again reaches its maximum size, ***trace-file.0*** is renamed ***trace-file.1*** and ***trace-file*** is renamed ***trace-file.0***. This renaming scheme continues until the maximum number of trace files is reached. Then the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the **files** option and filename.

Syntax: **xk** to specify KB, **xm** to specify MB, or **xg** to specify GB

Range: 10 KB through 1 GB

Default: 128 KB

world-readable—(Optional) Enable unrestricted file access.

non-world-readable—(Optional) By default, log files can be accessed only by the user who configures the tracing operation. Specify **non-world-readable** to reset the default.

Usage Guidelines See “Tracing Operations of the pppd Process” on page 119.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

track

Syntax	<pre>track { dlsw { destination <i>mac-address</i> priority-cost <i>priority</i>; peer <i>ip-address</i> priority-cost <i>priority</i>; } interface <i>interface-name</i> priority-cost <i>priority</i>; }</pre>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i>],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2 redundancy-group <i>group-number</i>]</p>
Release Information	Statement introduced in JUNOS Release 7.5.
Description	<p>For J Series Services Routers only. On Ethernet interfaces configured for DLSw Ethernet redundancy, enable tracking options for an interface, remote peer, or destination MAC address.</p> <p>The statements are explained separately.</p>
Options	<p>destination <i>mac-address</i>—Local MAC address. Specify the MAC address as six hexadecimal bytes in one of the following formats: <i>nnnn.nnnn.nnnn</i> or <i>nn:nn:nn:nn:nn:nn</i>. For example, 0011.2233.4455 or 00:11:22:33:44:55.</p> <p>dlsw—DLSw protocol.</p> <p>interface <i>interface-name</i>—Interface name. Include the logical portion of the name, which corresponds to the logical unit number.</p> <p>peer <i>ip-address</i>—IP address of the remote peer.</p> <p>priority-cost <i>priority</i>—Cost value that is subtracted from the priority value when remote peer connectivity is lost. Specify a value from 1 through 254.</p>
Usage Guidelines	See “Configuring DLSw Ethernet Redundancy Using LLC2 Properties” on page 181.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	<p>For information about DLSw, see the <i>JUNOS Services Interfaces Configuration Guide</i>.</p> <p>For information about the track statement at the [edit interfaces <i>interface-name</i> unit <i>unit-number</i> family inet address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-number</i>] or [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>unit-number</i> family (inet inet6) address <i>address</i> (vrrp-group vrrp-inet6-group) <i>group-number</i>] hierarchy level, see the <i>JUNOS High Availability Configuration Guide</i>.</p>

translate-discard-eligible

Syntax	(translate-discard-eligible no-translate-discard-eligible);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ccc], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ccc]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For interfaces with encapsulation type Frame Relay CCC, enable or disable translation of Frame Relay discard eligible (DE) control bits.
Default	DE bit translation is disabled.
Usage Guidelines	See “Configuring Frame Relay Control Bit Translation” on page 375.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

translate-fecn-and-becn

Syntax	(translate-fecn-and-becn no-translate-fecn-and-becn);
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ccc], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family ccc]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For interfaces with encapsulation type Frame Relay CCC, enable or disable translation of Frame Relay forward explicit congestion notification (FECN) control bits and Frame Relay backward explicit congestion notification (BECN) control bits.
Default	FECN and BECN bit translation is disabled.
Usage Guidelines	See “Configuring Frame Relay Control Bit Translation” on page 375.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

transmit-bucket

Syntax	transmit-bucket { overflow discard; rate <i>percentage</i> ; threshold <i>bytes</i> ; }
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>Set parameters for the transmit leaky bucket, which specifies what percentage of the interface's total capacity can be used to transmit packets.</p> <p>For each DS3 channel in a channelized OC12 interface, you can configure a unique transmit bucket.</p> <p>The statements are explained separately.</p>
Usage Guidelines	See “Configuring Receive and Transmit Leaky Bucket Properties” on page 129 and “Configuring Receive and Transmit Leaky Bucket Properties on SONET/SDH Interfaces” on page 876.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	receive-bucket

transmit-clock

Syntax	transmit-clock invert;
Hierarchy Level	[edit interfaces <i>interface-name</i> serial-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure the transmit clock signal.
Options	invert—Shift the clock phase 180 degrees.
Usage Guidelines	See “Configuring the Serial Clocking Mode” on page 269.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

transmit-period

Syntax	transmit-period <i>seconds</i> ;
Hierarchy Level	[edit protocols dot1x authenticator interface <i>interface-id</i>]
Release Information	Statement introduced in JUNOS Release 9.3.
Description	Set the number of seconds the port waits before retransmitting the initial EAPOL PDUs to the client.
Options	<i>seconds</i> —The number of seconds the port waits before retransmitting the initial EAPOL PDUs to the client. Range: 1 through 65,535 seconds Default: 30 seconds
Usage Guidelines	See “Configuring IEEE 802.1x Port-Based Network Access Control” on page 741.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	dot1x, authenticator, interface (IEEE 802.1x)

transmit-weight

See the following sections:

- transmit-weight (ATM2 IQ CoS Forwarding Class) on page 1294
- transmit-weight (ATM2 IQ Virtual Circuit) on page 1295

transmit-weight (ATM2 IQ CoS Forwarding Class)

Syntax	transmit-weight (cells <i>number</i> percent <i>number</i>);
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options scheduler-maps <i>map-name</i> forwarding-class <i>class-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, assign a transmission weight to a forwarding class.
Default	95 percent for queue 0, 5 percent for queue 3.
Options	percent <i>percent</i> —Transmission weight of the forwarding class as a percentage of the total bandwidth. Range: 5 through 100 cells <i>number</i> —Transmission weight of the forwarding class as a number of cells. Range: 0 through 32,000
Usage Guidelines	See “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

transmit-weight (ATM2 IQ Virtual Circuit)

Syntax	<code>transmit-weight number;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ PICs only, configure the transmission weight. Each VC is serviced in weighted round robin (WRR) mode. When VCs have data to send, they send the number of cells equal to their weight before passing control to the next active VC. This allows proportional bandwidth sharing between multiple VCs within a rate-shaped VP tunnel. VP tunnels are not supported on multipoint interfaces.
Options	<i>number</i> —Number of cells a VC sends before passing control to the next active VC within a VP tunnel. Range: 1 through 32,767
Usage Guidelines	See “Configuring the ATM2 IQ Transmission Weight” on page 329.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

traps

Syntax	<code>(traps no-traps);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Enable or disable the sending of Simple Network Management Protocol (SNMP) notifications when the state of the connection changes.
Usage Guidelines	See “Enabling or Disabling SNMP Notifications on Physical Interfaces” on page 139 and “Enabling or Disabling SNMP Notifications on Logical Interfaces” on page 159.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

trej-time

Syntax	<code>trej-time time;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family llc2]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	For J Series Services Routers only. On Ethernet interfaces configured for DLSw, configure the length of time a router waits for a rejected frame to be re-sent before the router sends the reject command.
Options	<i>time</i> —Number of milliseconds. Range: 1 through 60000 Default: 3000 milliseconds
Usage Guidelines	See “Configuring LLC2 Options” on page 180.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>

trigger

Syntax trigger {
 defect ignore;
 defect hold-time up *milliseconds* down *milliseconds*;
 }

Hierarchy Level [edit interfaces *interface-name* sonet-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description For ATM over SONET/SDH, SONET/SDH interfaces, and 10-Gigabit Ethernet interfaces in WAN PHY mode, configure SONET/SDH defect triggers to be ignored.

Default If you do not include this statement, all SONET/SDH defect triggers are honored.

Options *defect*—Defect to ignore or hold. It can be one of the following:

- *ais-l*—Line alarm indication signal
- *ais-p*—Path alarm indication signal
- *ber-sd*—Bit error rate signal degrade
- *ber-sf*—Bit error rate signal fault
- *locd* (ATM only)—Loss of cell delineation
- *lof*—Loss of frame
- *lol*—PHY loss of light
- *lop-p*—Path loss of pointer
- *los*—Loss of signal
- *pll*—PHY phase-locked loop out of lock
- *plm-p*—Path payload label mismatch
- *rfl-l*—Line remote failure indication
- *rfl-p*—Path remote failure indication
- *uneq-p*—Path unequipped

The remaining statements are explained separately.

Usage Guidelines See “Configuring SONET/SDH Defect Triggers to Be Ignored” on page 855.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

trigger-link-failure

Syntax	[trigger-link-failure <i>interface-name</i>];
Hierarchy Level	[edit interfaces <i>lsq-fpc/pic/port</i> lsq-failure-options]
Release Information	Statement introduced in JUNOS Release 7.4.
Description	List of SONET interfaces connected to the LSQ interface that can implement Automatic Protection Switching (APS) if the LSQ PIC fails.
Options	<i>interface-name</i> —Name of SONET interface.
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

trunk-bandwidth

Syntax	trunk-bandwidth <i>rate</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces configured to use Layer 2 circuit trunk mode, configure a scheduler so that unused bandwidth from any inactive trunk is proportionally shared among the active trunks. During congestion, each trunk receives a proportional share of the leftover bandwidth, thus minimizing the latency on each trunk.
Options	<i>rate</i> —Peak rate, in bits per second (bps) or cells per second (cps). You can specify a value in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify a value in cells per second by entering a decimal number followed by the abbreviation c; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps. Range: 1,000,000 through 542,526,792 bps
Usage Guidelines	See “Configuring Layer 2 Circuit Trunk Mode Scheduling” on page 309.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

trunk-id

Syntax	<code>trunk-id number;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces with ATM CCC cell-relay encapsulation, configure the trunk identification number. When you associate a trunk ID number with a logical interface, you are in effect specifying the interfaces that are allowed to send ATM traffic over an LSP.
Options	<i>number</i> —A valid trunk identifier. Range: For UNI mode, 0 through 7. For NNI mode, 0 through 31.
Usage Guidelines	See “Configuring Layer 2 Circuit Transport Mode” on page 300.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

ttl

Syntax	<code>ttl value;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>number</i> tunnel]
Release Information	Statement introduced before JUNOS Release 7.4
Description	Set the time-to-live value bit in the header of the outer IP packet.
Options	<i>value</i> —Time-to-live value. Range: 0 through 255 Default: 64
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

tunnel

Syntax tunnel {
 backup-destination *address*;
 destination *address*;
 key *number*;
 routing-instance {
 destination *routing-instance-name*;
 }
 source *source-address*;
 ttl *number*;
 }

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
 [edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure a tunnel. You can use the tunnel for unicast and multicast traffic or just for multicast traffic. You can also use tunnels for encrypted traffic or VPNs.

The statements are explained separately.

Usage Guidelines See the *JUNOS Services Interfaces Configuration Guide* and *JUNOS VPNs Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

underlying-interface

Syntax `underlying-interface interface-name;`

Hierarchy Level [edit interfaces pp0 unit *logical-unit-number* pppoe-options],
[edit interfaces demux0 unit *logical-unit-number* demux-options],
[edit logical-systems *logical-system-name* interfaces demux0 unit *logical-unit-number* demux-options],
[edit logical-systems *logical-system-name* interfaces pp0 unit *logical-unit-number* pppoe-options],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* interfaces demux0 unit *logical-unit-number* demux-options]
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* interfaces pp0 unit *logical-unit-number* pppoe-options],

Release Information Statement introduced before JUNOS Release 7.4.
Support for aggregated Ethernet added in JUNOS Release 9.4.

Description For J Series Services Routers and M120 Internet routers with PPP over Ethernet interfaces, configure the interface on which PPP over Ethernet is running.

For demux interfaces, configure the underlying interface on which the demultiplexing (demux) interface is running.

Options *interface-name*—Name of the interface on which PPP over Ethernet or demux is running. For example, **at-0/0/1.0** (ATM VC), **fe-1/0/1.0** (Fast Ethernet interface), **ge-2/0/0.0** (Gigabit Ethernet interface), or **ae1.0** (aggregated Ethernet interface).



NOTE: Logical demux interfaces are currently supported only on Gigabit Ethernet, Fast Ethernet, 10-Gigabit Ethernet interfaces, or aggregated Ethernet.

Usage Guidelines See “Configuring an IP Demux Underlying Interface” on page 252, “Specifying the Demux Underlying Interface” on page 253, and “Configuring the PPPoE Underlying Interface” on page 791.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics *J-series Services Router Basic LAN and WAN Access Configuration Guide*

unframed

Syntax	(unframed no-unframed);
Hierarchy Level	[edit interfaces <i>interface-name</i> e3-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For E3 IQ interfaces only, enable or disable unframed mode. In unframed mode, the E3 IQ interface do not detect yellow (ylw) or loss-of-frame (lof) alarms.
Default	Unframed mode is disabled.
Usage Guidelines	See “Configuring E3 IQ and IQE Unframed Mode” on page 558.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

unidirectional

Syntax	unidirectional;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 8.5.
Description	Create two new, unidirectional (transmit-only and receive-only) physical interfaces subordinate to the original parent interface. Unidirectional links are currently supported only on 10-Gigabit Ethernet interfaces on the following hardware: <ul style="list-style-type: none"> ■ 4-port 10-Gigabit Ethernet DPC on the MX960 router ■ 10-Gigabit Ethernet IQ2 PIC and 10-Gigabit Ethernet IQ2E PIC on the T Series router
Default	Disabled.
Usage Guidelines	See “Enabling Unidirectional Traffic Flow on Physical Interfaces” on page 139.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

unit

Syntax `unit logical-unit-number {`
 `accept-source-mac {`
 `mac-address mac-address {`
 `policer {`
 `input cos-policer-name;`
 `output cos-policer-name;`
 `}`
 `}`
 `}`
 `accounting-profile name;`
 `allow-any-vci;`
 `atm-scheduler-map (map-name | default);`
 `backup-options {`
 `interface interface-name;`
 `}`
 `bandwidth rate;`
 `cell-bundle-size cells;`
 `clear-dont-fragment-bit;`
 `compression {`
 `rtp {`
 `maximum-contexts number <force>;`
 `f-max-period number;`
 `queues [queue-numbers];`
 `port {`
 `minimum port-number;`
 `maximum port-number;`
 `}`
 `}`
 `}`
 `compression-device interface-name;`
 `copy-tos-to-outer-ip-header;`
 `demux-destination family;`
 `demux-source family;`
 `demux-options {`
 `underlying-interface interface-name;`
 `}`
 `description text;`
 `dial-options {`
 `l2tp-interface-id name;`
 `(dedicated | shared);`
 `}`
 `dialer-options {`
 `activation-delay seconds;`
 `callback;`
 `callback-wait-period time;`
 `deactivation-delay seconds;`
 `dial-string [dial-string-numbers];`
 `idle-timeout seconds;`
 `incoming-map {`
 `caller caller-id) | accept-all;`
 `initial-route-check seconds;`
 `}`
 `}`

```

        load-interval seconds;
        load-threshold percent;
        pool pool-name;
        redial-delay time;
        watch-list {
            [ routes ];
        }
    }
}
disable;
disable-mlppp-inner-ppp-pfc;
dlci dlci-identifier;
drop-timeout milliseconds;
dynamic-call-admission-control {
    activation-priority priority;
    bearer-bandwidth-limit kilobits-per-second;
}
encapsulation type;
epd-threshold cells plp1 cells;
fragment-threshold bytes;
inner-vlan-id-range start start-id end end-id;
input-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap |
    swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}
interleave-fragments;
inverse-arp;
layer2-policer {
    input-policer policer-name;
    input-three-color policer-name;
    output-policer policer-name;
    output-three-color policer-name;
}
link-layer-overhead percent;
minimum-links number;
mrru bytes;
multicast-dlci dlci-identifier;
multicast-vci vpi-identifier.vci-identifier;
multilink-max-classes number;
multipoint;
oam-liveness {
    up-count cells;
    down-count cells;
}
oam-period (disable | seconds);
output-vlan-map {
    (pop | pop-pop | pop-swap | push | push-push | swap |
    swap-push | swap-swap);
    inner-tag-protocol-id tpid;
    inner-vlan-id number;
    tag-protocol-id tpid;
    vlan-id number;
}

```



```

}
passive-monitor-mode;
peer-unit unit-number;
plp-to-clp;
point-to-point;
ppp-options {
    chap {
        access-profile name;
        default-chap-secret name;
        local-name name;
        passive;
    }
    compression {
        acfc;
        pfc;
    }
    dynamic-profile profile-name;
    lcp-restart-timer milliseconds;
    loopback-clear-timer seconds;
    ncp-restart-timer milliseconds;
    pap {
        access-profile name;
        default-pap-password password;
        local-name name;
        local-password password;
        passive;
    }
}
pppoe-options {
    access-concentrator name;
    auto-reconnect seconds;
    (client | server);
    service-name name;
    underlying-interface interface-name;
}
proxy-arp;
service-domain (inside | outside);
shaping {
    (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
    burst length);
    queue-length number;
}
short-sequence;
transmit-weight number;
(traps | no-traps);
trunk-bandwidth rate;
trunk-id number;
tunnel {
    backup-destination address;
    destination address;
    key number;
    routing-instance {
        destination routing-instance-name;
    }
    source source-address;
    ttl number;

```

```

}
vci vpi-identifier.vci-identifier;
vci-range start start-vci end end-vci;
vpi vpi-identifier;
vlan-id number;
vlan-id-range number-number;
vlan-tags (Stacked VLAN Tags) inner tpid.vlan-id outer tpid.vlan-id;
family family {
    accounting {
        destination-class-usage;
        source-class-usage {
            direction;
        }
    }
}
bundle interface-name;
filter {
    group filter-group-number;
    input filter-name;
    input-list {
        [ filter-names ];
        output filter-name;
    }
    output-list {
        [ filter-names ];
    }
}
ipsec-sa sa-name;
interface-mode (access | trunk);
keep-address-and-control;
mac-validate (loose | strict);
mtu bytes;
multicast-only;
no-redirects;
policer {
    arp policer-template-name;
    input policer-template-name;
    output policer-template-name;
}
primary;
proxy inet-address address;
receive-options-packets;
receive-ttl-exceeded;
remote (inet-address address | mac-address address);
rpf-check <fail-filter filter-name> {
    <mode loose>;
}
sampling {
    direction;
}
service {
    input {
        service-set service-set-name <service-filter filter-name>;
        post-service-filter filter-name;
    }
    output {
        service-set service-set-name <service-filter filter-name>;
    }
}

```

```

    }
}
(translate-discard-eligible | no-translate-discard-eligible);
(translate-fecn-and-becn | no-translate-fecn-and-becn);
unnumbered-address interface-name destination address destination-profile
    profile-name;
address address {
    arp ip-address (mac | multicast-mac) mac-address <publish>;
    broadcast address;
    destination address;
    eui-64;
    master-only;
    multipoint-destination address (dlci dlci-identifier | vci vci-identifier);
    multipoint-destination address {
        epd-threshold cells plp1 cells;
        inverse-arp;
        oam-liveness {
            up-count cells;
            down-count cells;
        }
        oam-period (disable | seconds);
        shaping {
            (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained
                rate burst length);
            queue-length number;
        }
        vci vpi-identifier.vci-identifier;
    }
}
preferred;
primary;
(vrrp-group | vrrp-inet6-group) group-number {
    (accept-data | no-accept-data);
    advertise-interval seconds;
    authentication-type authentication;
    authentication-key key;
    fast-interval milliseconds;
    (preempt | no-preempt) {
        hold-time seconds;
    }
    priority-number number;
    track {
        priority-cost seconds;
        priority-hold-time interface-name {
            interface priority;
            bandwidth-threshold bits-per-second {
                priority;
            }
        }
    }
}
virtual-address [ addresses ];
}
}
}

```

Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>], [edit interfaces interface-set (Ethernet Interfaces) <i>interface-set-name</i> interface <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure a logical interface on the physical device. You must configure a logical interface to be able to use the physical device.
Options	<i>logical-unit-number</i> —Number of the logical unit. Range: 0 through 16,384 The remaining statements are explained separately.
Usage Guidelines	See “Configuring Logical Interface Properties” on page 143.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS Services Interfaces Configuration Guide</i>

unnumbered-address

See the following sections:

- unnumbered-address (Demux) on page 1309
- unnumbered-address (Ethernet) on page 1310
- unnumbered-address (PPP) on page 1310

unnumbered-address (Demux)

Syntax	<code>unnumbered-address interface-name <preferred-source-address address>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced in JUNOS Release 8.2. <code>preferred-source-address</code> option introduced in JUNOS Release 9.0. IP demultiplexing interfaces supported in JUNOS Release 9.2.
Description	For IP demultiplexing interfaces, enable the local address to be derived from the specified interface. Configuring an unnumbered interface enables IP processing on the interface without assigning an explicit IP address to the interface.
Options	<i>interface-name</i> —Name of the interface from which the local address is derived. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface.
	The <code>preferred-source-address</code> statement is explained separately.
Usage Guidelines	See “Configuring an Unnumbered Interface” on page 185.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	address, <i>JUNOS System Basics Configuration Guide</i>

unnumbered-address (Ethernet)

Syntax	<code>unnumbered-address interface-name <preferred-source-address address>;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i>]
Release Information	Statement introduced in JUNOS Release 8.2. <code>preferred-source-address</code> option introduced in JUNOS Release 9.0.
Description	For Ethernet interfaces, enable the local address to be derived from the specified interface. Configuring an unnumbered Ethernet interface enables IP processing on the interface without assigning an explicit IP address to the interface.
Options	<i>interface-name</i> —Name of the interface from which the local address is derived. The specified interface must have a logical unit number and a configured IP address, and must not be an unnumbered interface. The <code>preferred-source-address</code> statement is explained separately.
Usage Guidelines	See “Configuring an Unnumbered Interface” on page 185.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	address, <i>JUNOS System Basics Configuration Guide</i>

unnumbered-address (PPP)

Syntax	<code>unnumbered-address interface-name destination address destination-profile profile-name;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family inet]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For interfaces with PPP encapsulation, enable the local address to be derived from the specified interface.
Options	<i>interface-name</i> —Interface from which the local address is derived. The interface name must include a logical unit number and must have a configured address. The remaining statements are explained separately.
Usage Guidelines	See “Configuring IPCP Options” on page 177.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	address, negotiate-address, <i>JUNOS System Basics Configuration Guide</i>

up-count

Syntax	up-count <i>cells</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i> oam-liveness], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> oam-liveness], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i> oam-liveness], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> oam-liveness], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i> oam-liveness]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM encapsulation only, configure Operation, Administration, and Maintenance (OAM) F5 loopback cell count thresholds. Not supported on ATM-over-SHDSL interfaces. For ATM2 IQ PICs only, configure OAM F4 loopback cell count thresholds at the [edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i>] hierarchy level.
Options	<i>cells</i> —Minimum number of consecutive OAM F4 or F5 loopback cells received before a VC is declared up. Range: 1 through 255 Default: 5 cells
Usage Guidelines	See “Configuring the ATM OAM F5 Loopback Cell Threshold” on page 329.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

vbr

Syntax	<code>vbr peak <i>rate</i> sustained <i>rate</i> burst <i>length</i>;</code>
Hierarchy Level	<p>[edit interfaces <i>interface-name</i> atm-options vpi <i>vpi-identifier</i> shaping],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping],</p> <p>[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> address <i>address</i> family <i>family</i> multipoint-destination <i>address</i> shaping],</p> <p>[edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> shaping]</p>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For ATM encapsulation only, define the variable bandwidth utilization in the traffic-shaping profile.</p> <p>When you configure the variable bandwidth utilization, you must specify all three options (burst, peak, and sustained). You can specify rate in bits per second either as a complete decimal number or as a decimal number followed by the abbreviation k (1000), m (1,000,000), or g (1,000,000,000). You can also specify rate in cells per second by entering a decimal number followed by the abbreviation c; values expressed in cells per second are converted to bits per second by means of the formula 1 cps = 384 bps.</p>
Default	If the vbr statement is not specified, bandwidth utilization is unlimited.
Options	<p>burst <i>length</i>—Burst length, in cells. If you set the length to 1, the peak traffic rate is used.</p> <p>Range: 1 through 4000 cells</p> <p>peak <i>rate</i>—Peak rate, in bits per second or cells per second.</p> <p>Range: For ATM1 interfaces, 33 Kbps through 135.6 Mbps (ATM OC3); 33 Kbps through 276 Mbps (ATM OC12). For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps. For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps. For ATM2 IQ DS3 and E3 interfaces, from 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure.</p> <p>For more information, see Table 27 on page 320.</p> <p>sustained <i>rate</i>—Sustained rate, in bits per second or cells per second.</p> <p>Range: For ATM1 interfaces, 33 Kbps through 135.6 Mbps (ATM OC3); 33 Kbps through 276 Mbps (ATM OC12).</p> <p>For ATM2 IQ OC3 and OC12 interfaces, 33 Kbps through 542,526,792 bps.</p> <p>For ATM2 IQ OC48 interfaces, 33 Kbps through 2,170,107,168 bps.</p>

For ATM2 IQ DS3 and E3 interfaces, from 33 Kbps through the maximum rate, which depends on the ATM encapsulation and framing you configure. For more information, see Table 27 on page 320.

Usage Guidelines	See “Defining the ATM Traffic-Shaping Profile” on page 319.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	cbr, rtvbr, shaping

vc-cos-mode

Syntax	vc-cos-mode (alternate strict);
Hierarchy Level	[edit interfaces <i>interface-name</i> atm-options scheduler-maps <i>map-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For ATM2 IQ interfaces only, specify packet-scheduling priority value for ATM2 IQ VC tunnels.
Options	<p>alternate—VC CoS queue has high priority. The scheduling of the queues alternates between the high-priority queue and the remaining queues, so every other scheduled packet is from the high-priority queue.</p> <p>strict—VC CoS queue has strictly high priority. A queue with strict high priority is always scheduled before the remaining queues. The remaining queues are scheduled in round-robin fashion.</p> <p>Default: alternate</p>
Usage Guidelines	See “Configuring ATM2 IQ VC Tunnel CoS Components” on page 339.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

vci

Syntax	<code>vci vpi-identifier.vci-identifier;</code>
Hierarchy Level	<p>[edit interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i> family <i>family</i> address <i>address</i> multipoint-destination <i>address</i>]</p>
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For ATM point-to-point logical interfaces only, configure the virtual circuit identifier (VCI) and virtual path identifier (VPI).</p> <p>To configure a VPI for a point-to-multipoint interface, specify the VPI in the multipoint-destination statement.</p> <p>VCIs 0 through 31 are reserved for specific ATM values designated by the ATM Forum.</p>
Options	<p>vci-identifier—ATM virtual circuit identifier. Unless you configure the interface to use promiscuous mode, this value cannot exceed the largest numbered VC configured for the interface with the maximum-vcs option of the vpi statement.</p> <p>Range: 0 through 4089 or 0 through 65,535 with promiscuous mode, with VCIs 0 through 31 reserved.</p> <p>vpi-identifier—ATM virtual path identifier.</p> <p>Range: 0 through 255</p> <p>Default: 0</p>
Usage Guidelines	See “Configuring a Point-to-Point ATM1 or ATM2 IQ Connection” on page 316.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>
Related Topics	multipoint-destination, promiscuous-mode, vpi

vci-range

Syntax	<code>vci-range start <i>start-vci</i> end <i>end-vci</i>;</code>
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Range of VCI values used in ATM-to-Ethernet interworking cross-connects. VCI 0 through 31 are reserved. VCI 0 through 31 should not be used.
Options	<code>start-vci</code> —Lowest number VCI in the range. <code>end-vci</code> —Highest number VCI in the range. Range: 0 through 255
Usage Guidelines	See “Configuring ATM-to-Ethernet Interworking” on page 229.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

virtual-switch

Syntax	<code>virtual-switch <i>name</i> bridge-domain <i>name</i> vlan-id [<i>id1 id2 ... idn</i>]</code>
Hierarchy Level	[edit protocols oam ethernet connectivity-fault-management maintenance-domain <i>domain-name</i> default-x]
Description	Specify the routing-instance type as a virtual switch, under which bridge-domain MIPs must be enabled.
Usage Guidelines	See “Configuring MIP for Bridge Domains of a Virtual Switch” on page 683.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

vlan-id

See the following sections:

- [vlan-id \(VLAN ID to Be Bound to a Logical Interface\)](#) on page 1316
- [vlan-id \(Logical Port in Bridge Domain\)](#) on page 1317
- [vlan-id \(VLAN ID to Rewrite\)](#) on page 1317
- [vlan-id \(Outer VLAN ID\)](#) on page 1318

vlan-id (VLAN ID to Be Bound to a Logical Interface)

Syntax `vlan-id number;`

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description For Fast Ethernet, Gigabit Ethernet, and Aggregated Ethernet interfaces only, bind a 802.1Q VLAN tag ID to a logical interface.

Options *number*—A valid VLAN identifier.

Range: For aggregated Ethernet, 4-port, 8-port, and 12-port Fast Ethernet PICs, and for management and internal Ethernet interfaces, 1 through 1023.

For 48-port Fast Ethernet and Gigabit Ethernet PICs, 1 through 4094.

VLAN ID 0 is reserved for tagging the priority of frames.

Usage Guidelines See “Configuring Mixed Tagging” on page 602.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

vlan-id (Logical Port in Bridge Domain)

Syntax	<code>vlan-id number;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	The VLAN ID configured on the logical port. Received packets with no VLAN tags are forwarded within the bridge domain with the matching VLAN ID.
Options	number—The VLAN ID. Range: 1 through 4095
Usage Guidelines	See “Configuring a Logical Interface for Access Mode” on page 619.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

vlan-id (VLAN ID to Rewrite)

Syntax	<code>vlan-id number;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> input-vlan-map], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Gigabit Ethernet IQ and 10-Gigabit Ethernet IQ2 and IQ2-E interfaces and aggregated Ethernet using Gigabit Ethernet IQ interfaces, specify the line VLAN identifiers to be rewritten at the input or output interface. You cannot include the <code>vlan-id</code> statement with the <code>swap</code> statement, <code>swap-push</code> statement, <code>push-push</code> statement, or <code>push-swap</code> statement at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> output-vlan-map] hierarchy level. If you include any of those statements in the output VLAN map, the VLAN ID in the outgoing frame is rewritten to the <code>vlan-id</code> statement you include at the [edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>] hierarchy level.
Usage Guidelines	See “Rewriting the VLAN Tag on Tagged Frames” on page 651 and “Binding VLAN IDs to Logical Interfaces” on page 604.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

vlan-id (Outer VLAN ID)

Syntax	<code>vlan-id outer-vlan-id;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	The outer VLAN ID to be used in ATM-to-Ethernet interworking cross-connects. Outer VLAN IDs are converted to the ATM VPI. The outer VLAN ID must match the VPI value configured. The allowable VPI range is 0 to 255. Do not configure the outer VLAN ID to be greater than 255.
Options	outer-vlan-id—Outer VLAN ID number. Range: 0 through 4094
Usage Guidelines	See “Configuring ATM-to-Ethernet Interworking” on page 229.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

vlan-id-list

See the following sections:

- [vlan-id-list \(Interface in Bridge Domain\)](#) on page 1319
- [vlan-id-list \(Ethernet VLAN Circuit\)](#) on page 1320

vlan-id-list (Interface in Bridge Domain)

Syntax	<code>vlan-id-list [<i>number number-number</i>];</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i> family bridge]
Release Information	Statement introduced in JUNOS Release 9.2.
Description	Configure a logical interface to forward packets and learn MAC addresses within each bridge domain configured with a VLAN ID that matches a VLAN ID specified in the list. VLAN IDs can be entered individually using a space to separate each ID, entered as an inclusive list separating the starting VLAN ID and ending VLAN ID with a hyphen, or a combination of both.
Options	<p><i>number number</i>—Individual VLAN IDs separated by a space.</p> <p><i>number-number</i>—Starting VLAN ID and ending VLAN ID in an inclusive range. Range: 1 through 4095</p>
Usage Guidelines	See “Configuring a Logical Interface for Trunk Mode” on page 620 and “Configuring the VLAN ID List for a Trunk Interface” on page 620.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

vlan-id-list (Ethernet VLAN Circuit)

Syntax	<code>vlan-id-list [<i>vlan-id</i> <i>vlan-id-vlan-id</i>];</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	(MX Series routers only) Binds a single-tag logical interface to a list of VLAN IDs. Configures a logical interface to receive and forward any tag frame whose VLAN ID tag matches the list of VLAN IDs you specify.

**NOTE:**

When you create a circuit cross-connect (CCC) using VLAN-bundled single-tag logical interfaces on Layer 2 VPN routing instances, the circuit automatically uses **ethernet** encapsulation. For Layer 2 VPN, you need to include the **encapsulation-type** statement and specify the value **ethernet** at either of the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols l2vpn]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn]

For more information about the **encapsulation-type** configuration statement and the Layer 2 encapsulation types **ethernet** and **ethernet-vlan**, see the *JUNOS VPNs Configuration Guide*.

Options	[<i>vlan-id</i> <i>vlan-id-vlan-id</i>] —A list of valid VLAN ID numbers. Specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both. Range: 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.
----------------	--



NOTE: Configuring **vlan-id-list** with the entire *vlan-id* range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

Inefficient	[edit interfaces <i>interface-name</i>] vlan-tagging; unit <i>number</i> {
--------------------	---


```

        vlan-id-range 1-4094;
    }

```

Best Practice

```

[edit interfaces interface-name]

    unit 0;

```

Usage Guidelines See “Binding VLAN IDs to Logical Interfaces” on page 604.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

- Related Topics**
- encapsulation (Logical Interface)
 - encapsulation (Physical Interface)
 - encapsulation-type (Layer 2 VPN routing instance), see the *JUNOS VPNs Configuration Guide*.
 - flexible-vlan-tagging
 - vlan-tagging
 - vlan-tags (Dual-Tagged Logical Interface)

vlan-id-range

Syntax	<code>vlan-id-range <i>vlan-id-vlan-id</i></code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 8.4.
Description	Bind a range of VLAN IDs to a logical interface.
Options	number—The first number is the lowest VLAN ID in the range the second number is the highest VLAN ID in the range. Range: 1 through 4094



NOTE: Configuring `vlan-id-range` with the entire `vlan-id` range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

Inefficient	<pre>[edit interfaces <i>interface-name</i> vlan-tagging; unit <i>number</i> { vlan-id-range 1-4094; }</pre>
Best Practice	<pre>[edit interfaces <i>interface-name</i> unit 0;</pre>

VLAN ID 0 is reserved for tagging the priority of frames.

Usage Guidelines	See “Binding a Range of VLAN IDs to a Logical Interface” on page 606.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

vlan-ranges

Syntax `vlan-ranges {
 dynamic-profile (VLAN) profile-name {
 accept (inet);
 ranges (Dynamic VLAN) (any | low-tag) - (any | high-tag);
 }
}`

Hierarchy Level [edit interfaces *interface-name* dynamic-profile *profile-name*]

Release Information Statement introduced in JUNOS Release 9.5.

Description Configure multiple VLANs. Each VLAN is assigned a VLAN ID number from the range.

Options any—Any valid VLAN ID number.

vlan-id-low—Specify the first VLAN ID number for the group of VLANs.

vlan-id-high—Specify the last VLAN ID number for the group of VLANs.

Range: 1 through 4094

The remaining statements are explained separately.

Required Privilege Level routing—To view this statement in the configuration.
 routing-control—To add this statement to the configuration.

Related Topics

- Configuring Single-Level VLAN Ranges for Use with VLAN Dynamic Profiles
- Configuring Dynamic Mixed VLAN Ranges

vlan-rewrite

Syntax	vlan-rewrite translate (200 500 201 501)
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>number</i> family bridge interface-mode trunk]
Release Information	Statement introduced in JUNOS Release 9.4.
Description	Translates an incoming VLAN to a bridge-domain VLAN, corresponding counter translation at egress. Supports translation of VLAN 200 to VLAN 500 and VLAN 201 to VLAN 501. Other valid vlan pass through without translation.
Options	translate 200 500—Translates incoming packets with VLAN 200 to 500. translate 201 501—Translates incoming packets with VLAN 201 to 501. translate 202 502—Translates incoming packets with VLAN 202 to 502.
Usage Guidelines	See “Rewriting a VLAN Tag and Adding a New Tag” on page 655.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

vlan-tagging

Syntax	vlan-tagging;
Hierarchy Level	[edit interfaces <i>interface-name</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For Fast Ethernet and Gigabit Ethernet interfaces and aggregated Ethernet interfaces configured for VPLS, enables the reception and transmission of 802.1Q VLAN-tagged frames on the interface.
Usage Guidelines	See “Configuring 802.1Q VLANs” on page 599 and “Configuring Tagged Aggregated Ethernet Interfaces” on page 632.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

vlan-tags

See the following sections:

- [vlan-tags \(Dual-Tagged Logical Interface\)](#) on page 1326
- [vlan-tags \(Stacked VLAN Tags\)](#) on page 1328

vlan-tags (Dual-Tagged Logical Interface)

Syntax	<code>vlan-tags inner-list [vlan-id vlan-id-vlan-id] outer <tpid.>vlan-id;</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 9.5.
Description	(MX Series routers only) Binds a dual-tag logical interface to a list of VLAN IDs. Configures the logical interface to receive and forward any dual-tag frame whose inner VLAN ID tag matches the list of VLAN IDs you specify.

**NOTE:**

To create a circuit cross-connect (CCC) using VLAN-bundled dual-tag logical interfaces on Layer 2 VPN routing instances, you must include the **encapsulation-type** statement and specify the value **ethernet-vlan** at the one of the following hierarchy levels:

- [edit routing-instances *routing-instance-name* protocols l2vpn]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols l2vpn]

For more information about the **encapsulation-type** configuration statement and the Layer 2 encapsulation types **ethernet** and **ethernet-vlan**, see the *JUNOS VPNs Configuration Guide*.

Options	<p><code>inner-list [vlan-id vlan-id vlan-id-vlan-id]</code>—A list of valid VLAN ID numbers. Specify the VLAN IDs individually by using a space to separate each ID, as an inclusive list by separating the starting VLAN ID and ending VLAN ID with a hyphen, or as a combination of both.</p> <p>Range: 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.</p> <p><code>outer <tpid.>vlan-id</code>—An optional Tag Protocol ID (TPID) and a valid VLAN ID.</p> <p>Range: For TPID, specify a hexadecimal value in the format <code>0xnnnn</code>.</p> <p>Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.</p>
----------------	--



NOTE: Configuring `inner-list` with the entire `vlan-id` range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup tables entries, and so on.

The following examples illustrate this further:

Inefficient

```
[edit interfaces interface-name]

vlan-tagging;
unit number {
    vlan-tags outer vid inner-list 1-4094;
}
```

Best Practice

```
[edit interfaces interface-name]

vlan-tagging;
unit number {
    vlan-id vid;
}
```

Usage Guidelines See “Binding VLAN IDs to Logical Interfaces” on page 604.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics

- encapsulation (Logical Interface)
- encapsulation (Physical Interface)
- encapsulation-type (Layer 2 VPN routing instance), see the *JUNOS VPNs Configuration Guide*.
- flexible-vlan-tagging
- vlan-id-list (Ethernet VLAN Circuit)
- vlan-tagging

vlan-tags (Stacked VLAN Tags)

Syntax `vlan-tags inner tpid.vlan-id inner-range vid1-vid2 outer tpid.vlan-id;`

Hierarchy Level [edit interfaces *interface-name* unit *logical-unit-number*],
[edit logical-systems *logical-system-name* interfaces *interface-name* unit *logical-unit-number*]

Release Information Statement introduced before JUNOS Release 7.4.

Description For Gigabit Ethernet IQ and IQE interfaces only, binds TPIDs and 802.1Q VLAN tag IDs to a logical interface. You must include the **stacked-vlan-tagging** statement at the [edit interfaces *interface-name*] hierarchy level.



NOTE: The inner-range *vid1-vid2* option is supported on MX Series with IQE PICs only.

Options inner *tpid.vlan-id*—A TPID and a valid VLAN identifier.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

inner-range *vid1-vid2*—For MX Series routers with Enhanced IQ (IQE) PICs only; specify a range of VLAN IDs where vid1 is the start of the range and vid2 is the end of the range.

Range: For VLAN ID, 1 through 4094. VLAN ID 0 is reserved for tagging the priority of frames.

outer *tpid.vlan-id*—A TPID and a valid VLAN identifier.

Range: For VLAN ID, 1 through 511 for normal interfaces, and 512 through 4094 for VLAN CCC interfaces. VLAN ID 0 is reserved for tagging the priority of frames.



NOTE: Configuring inner-range with the entire vlan-id range is an unnecessary waste of system resources and is not best practice. It should be used only when a subset of VLAN IDs of inner tag (not the entire range) needs to be associated with a logical interface. If you specify the entire range (1-4094), it has the same result as not specifying a range; however, it consumes PFE resources such as VLAN lookup table entries, and so on.

The following examples illustrate this further:

Inefficient

```
[edit interfaces interface-name

stacked-vlan-tagging;
unit number {
    vlan-tags outer vid inner-range 1-4094;
}
```


Best Practice `[edit interfaces interface-name]`

```
vlan-tagging;
unit number {
    vlan-id vid;
}
```

Usage Guidelines See “Configuring Dual VLAN Tags” on page 645.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Topics stacked-vlan-tagging

vlan-tags-outer

Syntax `vlan-tags-outer vlan-tag;`

Hierarchy Level `[edit interfaces interface-set (Ethernet Interfaces) interface-set-name interface interface-name]`

Release Information Statement introduced in JUNOS Release 8.5.

Description The S-VLAN outer tag that belongs to a set of interfaces used to configure hierarchical CoS schedulers.

Usage Guidelines See the *JUNOS Class of Service Configuration Guide*.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

vlan-vci-tagging

Syntax	vlan-vci-tagging;
Hierarchy Level	[edit interfaces <i>interface-name</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>interface-name</i>]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	Enable the ATM-to-Ethernet interworking cross-connect function on a Gigabit Ethernet, 10-Gigabit Ethernet, or aggregated Ethernet interface.
Usage Guidelines	See “Configuring ATM-to-Ethernet Interworking” on page 229.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

vpi

See the following sections:

- vpi (ATM CCC Cell-Relay Promiscuous Mode) on page 1331
- vpi (Define Virtual Path) on page 1332
- vpi (Logical Interface and Interworking) on page 1333

vpi (ATM CCC Cell-Relay Promiscuous Mode)

Syntax	<code>vpi vpi-identifier;</code>
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> atm-options promiscuous-mode]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	<p>For ATM interfaces, allow all VCIs in this VPI to open in ATM CCC cell-relay mode.</p> <p>When you include <code>vpi</code> statements at the [edit interfaces <i>interface-name</i> atm-options promiscuous-mode] hierarchy level, the specified VPIs open in promiscuous mode.</p>
Options	<p>vpi-identifier—ATM virtual path identifier. This is one of the VPIs that you define in the <code>vci</code> statement. (For a list of hierarchy levels at which you can include the <code>vci</code> statement, see <code>vci</code>.)</p> <p>Range: 0 through 255</p>
Usage Guidelines	See “Configuring ATM Cell-Relay Promiscuous Mode” on page 296.
Required Privilege Level	<p>interface—To view this statement in the configuration.</p> <p>interface-control—To add this statement to the configuration.</p>

vpi (Define Virtual Path)

Syntax `vpi vpi-identifier {
 maximum-vcs maximum-vcs;
 oam-liveness {
 up-count cells;
 down-count cells;
 }
 oam-period (disable | seconds);
 shaping {
 (cbr rate | rtvbr peak rate sustained rate burst length | vbr peak rate sustained rate
 burst length);
 queue-length number;
 }
}`

Hierarchy Level [edit interfaces *at-fpc/pic/port* atm-options]

Release Information Statement introduced before JUNOS Release 7.4.

Description For ATM interfaces, configure the virtual path (VP).

Options *vpi-identifier*—ATM virtual path identifier. This is one of the VPIs that you define in the *vci* statement. (For a list of hierarchy levels at which you can include the *vci* statement, see *vci*.)

Range: 0 through 255

The remaining statements are explained separately.

Usage Guidelines See “Configuring the Maximum Number of ATM1 VCs on a VP” on page 300.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics multipoint-destination, promiscuous-mode, vci

vpi (Logical Interface and Interworking)

Syntax	<code>vpi virtual-path-identifier;</code>
Hierarchy Level	[edit interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i>], [edit logical-systems <i>logical-system-name</i> interfaces <i>at-fpc/pic/port</i> unit <i>logical-unit-number</i>]
Release Information	Statement introduced in JUNOS Release 9.0.
Description	VPI used in an ATM-to-Ethernet interworking cross-connect.
Options	virtual-path-identifier—VPI to be used. Range: 0 through 255
Usage Guidelines	See “Configuring ATM-to-Ethernet Interworking” on page 229 and “Configuring ATM Cell-Relay Promiscuous Mode” on page 296.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.

vtmapping

Syntax	<code>vtmapping (itu-t klm);</code>
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options]; [edit chassis <i>fpc number pic number</i>]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For the Channelized STM1 IQ PIC or Channelized STM1 PIC, configure virtual tributary mapping. For the Channelized STM1 PIC, you configure virtual tributary mapping at the [edit chassis <i>fpc number pic number</i>] hierarchy level.
Options	itu-t—International Telephony Union standard klm—KLM standard Default: klm
Usage Guidelines	See “Configuring Virtual Tributary Mapping of Channelized STM1 Interfaces” on page 472.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>JUNOS System Basics Configuration Guide</i>

watch-list

Syntax	<pre>watch-list { [routes]; }</pre>
Hierarchy Level	[edit interfaces <i>dl</i> n unit <i>logical-unit-number</i> dialer-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	On J Series Services Routers with ISDN interfaces, configure an ISDN list of routes to watch. Used only for dialer watch.
Options	<i>routes</i> —IP prefix of a route. Specify one or more. The primary interface is considered up if there is at least one valid route for any of the addresses in the watch list to an interface other than the backup interface.
Usage Guidelines	See “Configuring Dialer Watch” on page 835.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>

wavelength

Syntax	wavelength <i>nm</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> optics-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For 10-Gigabit Ethernet DWDM interfaces only, configure full C-band ITU-Grid tunable optics.
Options	<p><i>nm</i>—Wavelength value. It can be one of the following:</p> <ul style="list-style-type: none"> ■ 1528.77—1528.77 nanometers (nm), corresponds to 196.10 terahertz (THz) ■ 1529.55—1529.55 nm, corresponds to 196.00 THz ■ 1530.33—1530.33 nm, corresponds to 195.90 THz ■ 1531.12—1531.12 nm, corresponds to 195.80 THz ■ 1531.90—1531.90 nm, corresponds to 195.70 THz ■ 1532.68—1532.68 nm, corresponds to 195.60 THz ■ 1533.47—1533.47 nm, corresponds to 195.50 THz ■ 1534.25—1534.25 nm, corresponds to 195.40 THz ■ 1535.04—1535.04 nm, corresponds to 195.30 THz ■ 1535.82—1535.82 nm, corresponds to 195.20 THz ■ 1536.61—1536.61 nm, corresponds to 195.10 THz ■ 1537.40—1537.40 nm, corresponds to 195.00 THz ■ 1538.19—1538.19 nm, corresponds to 194.90 THz ■ 1538.98—1538.98 nm, corresponds to 194.80 THz ■ 1539.77—1539.77 nm, corresponds to 194.70 THz ■ 1540.56—1540.56 nm, corresponds to 194.60 THz ■ 1541.35—1541.35 nm, corresponds to 194.50 THz ■ 1542.14—1542.14 nm, corresponds to 194.40 THz ■ 1542.94—1542.94 nm, corresponds to 194.30 THz ■ 1543.73—1543.73 nm, corresponds to 194.20 THz ■ 1544.53—1544.53 nm, corresponds to 194.10 THz ■ 1545.32—1545.32 nm, corresponds to 194.00 THz ■ 1546.12—1546.12 nm, corresponds to 193.90 THz ■ 1546.92—1546.92 nm, corresponds to 193.80 THz ■ 1547.72—1547.72 nm, corresponds to 193.70 THz

- 1548.52—1548.52 nm, corresponds to 193.60 THz
- 1549.32—1549.32 nm, corresponds to 193.50 THz
- 1550.12—1550.12 nm, corresponds to 193.40 THz
- 1550.92—1550.92 nm, corresponds to 193.30 THz
- 1551.72—1551.72 nm, corresponds to 193.20 THz
- 1552.52—1552.52 nm, corresponds to 193.10 THz
- 1553.33—1553.33 nm, corresponds to 193.00 THz
- 1554.13—1554.13 nm, corresponds to 192.90 THz
- 1554.94—1554.94 nm, corresponds to 192.80 THz
- 1555.75—1555.75 nm, corresponds to 192.70 THz
- 1556.56—1556.56 nm, corresponds to 192.60 THz
- 1557.36—1557.36 nm, corresponds to 192.50 THz
- 1558.17—1558.17 nm, corresponds to 192.40 THz
- 1558.98—1558.98 nm, corresponds to 192.30 THz
- 1559.79—1559.79 nm, corresponds to 192.20 THz
- 1560.61—1560.61 nm, corresponds to 192.10 THz
- 1561.42—1561.42 nm, corresponds to 192.00 THz
- 1562.23—1562.23 nm, corresponds to 191.90 THz
- 1563.05—1563.05 nm, corresponds to 191.80 THz
- 1563.86—1563.86 nm, corresponds to 191.70 THz
- **Default:** 1550.12—1550.12 nm, corresponds to 193.40 THz

Usage Guidelines See “Configuring the 10-Gigabit Ethernet DWDM Interface Wavelength” on page 779.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

west-interface

Syntax west-interface {
 control-channel *channel-name* {
 vlan *number*;
 }
 }

Hierarchy Level [edit protocols protection-group ethernet-ring *ring-name*]

Release Information Statement introduced in JUNOS Release 9.5.

Description For Ethernet ring protection, each ring should have two interface ports; an east-interface and a west-interface.



NOTE: Always configure the **east-interface** first, before configuring the **west-interface**.

The interface must use the control channel's logical interface name. The control channel is a dedicated VLAN channel for the ring port.

Options ring-protection-link-end—If this port is one side of RPL, this flag should be set.

Usage Guidelines See “Configuring Ethernet Ring Protection Switching” on page 799.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics ethernet-ring

working-circuit

Syntax working-circuit *group-name*;

Hierarchy Level [edit interfaces *interface-name* sonet-options aps]

Release Information Statement introduced before JUNOS Release 7.4.

Description Configure the working router in an APS circuit pair.

Options *group-name*—Circuit's group name.

Usage Guidelines See “Configuring Basic APS Support” on page 861.

Required Privilege Level interface—To view this statement in the configuration.
 interface-control—To add this statement to the configuration.

Related Topics protect-circuit

yellow-differential-delay

Syntax	yellow-differential-delay <i>milliseconds</i> ;
Hierarchy Level	[edit interfaces <i>interface-name</i> mlfr-uni-nni-bundle-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	For link services and voices interfaces only, configure the yellow differential delay among bundle links to give warning when a link has a differential delay that exceeds the configured threshold.
Options	<i>milliseconds</i> —Yellow differential delay threshold. Range: 1 through 2000 milliseconds Default: 6 milliseconds
Usage Guidelines	See the <i>JUNOS Services Interfaces Configuration Guide</i> .
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	action-red-differential-delay, remote

z0-increment

Syntax	(z0-increment no-z0-increment);
Hierarchy Level	[edit interfaces <i>interface-name</i> sonet-options]
Release Information	Statement introduced before JUNOS Release 7.4.
Description	Configure an incrementing STM ID rather than a static one.
Usage Guidelines	See “Configuring an Incrementing STM ID” on page 850.
Required Privilege Level	interface—To view this statement in the configuration. interface-control—To add this statement to the configuration.
Related Topics	sonet-options

Part 14

Index

- Index on page 1341
- Index of Statements and Commands on page 1367

Index

Symbols

#, comments in configuration statements.....	lxvi
(), in syntax descriptions.....	lxvi
10-Gigabit Ethernet interfaces.....	591
802.3ah OAM.....	752
DWDM.....	779
framing.....	781
SONET.....	872
10-Gigabit Ethernet IQ PIC.....	781
128-bit IPv6 address.....	176
32-bit IPv4 address.....	176
802.1ag OAM	
configuring Ethernet interfaces.....	679
802.1Q VLANs	
dynamic.....	921
mixed VLAN tagging.....	603, 1027, 1144
VLAN IDs.....	1316, 1322
values, listed by Ethernet interface	
type.....	600
VLAN tagging.....	599, 665, 1324
802.1Q/Ethernet 802.3 encapsulation	
encapsulation overhead.....	104
802.1Q/Ethernet SNAP encapsulation	
encapsulation overhead.....	104
802.1x Port-Based Network Access Control.....	987
802.3ad statement.....	889
usage guidelines.....	625
802.3ah OAM	
configuring Ethernet interfaces.....	745
example configuration.....	752
< >, in syntax descriptions.....	lxv
[], in configuration statements.....	lxvi
{ }, in configuration statements.....	lxvi
(pipe), in syntax descriptions.....	lxvi

A

accept statement.....	890
accept-source-mac statement.....	891
usage guidelines.....	759
access interface	
interface-mode statement.....	619

access profile	
configuring.....	113
default CHAP secret.....	113
default PAP password.....	165
access-concentrator statement.....	892
usage guidelines.....	790
access-profile statement.....	893
usage guidelines.....	112, 114
accounting profiles	
logical interfaces.....	158
example configuration.....	158
physical interfaces.....	130
example configuration.....	131
accounting statement.....	894
usage guidelines.....	214
accounting-profile statement.....	894
usage guidelines.....	130, 158
acfc statement.....	895
usage guidelines.....	120
ack-delay-time statement.....	895
usage guidelines.....	180
ack-max statement.....	896
usage guidelines.....	180
acknowledge-retries statement.....	896
acknowledge-timer statement.....	897
usage guidelines.....	897
action statement.....	898
action-profile statement	
applying to remote MEP.....	899
CFM.....	899
LFM.....	900
action-red-differential-delay statement.....	901
usage guidelines.....	901
activation-delay statement.....	901
usage guidelines.....	835
activation-priority statement.....	902
active statement	
usage guidelines.....	627
address and control field compression.....	120
address entries.....	50
address statement.....	903
usage guidelines.....	174, 185
ADSL	
example configuration.....	795, 796

advertise-interval statement.....	905	load sharing.....	870
APS		load sharing between circuits	
usage guidelines.....	868	example configuration.....	872
advertisement intervals		overview.....	859
APS interfaces.....	868	revertive mode.....	867
ae		timers.....	868
physical part of interface name.....	52	unidirectional mode.....	867
age statement.....	906	aps statement.....	914
aggregate statement.....	907	annex-b.....	913
hierarchical policer.....	908	usage guidelines.....	861
usage guidelines.....	757, 882	ARP	
aggregate-ports statement.....	909	aging timer.....	596
aggregated Ethernet interfaces.....	625	arp option	
CCC.....	236, 237	policers.....	194
configuring.....	623	ARP proxy, unrestricted	
example configuration.....	814	Ethernet interfaces.....	671
LACP.....	627	arp statement.....	915
example configuration.....	631	usage guidelines.....	669
interval.....	628	ARP table, static	
traceoptions.....	631	Ethernet interfaces.....	669
Layer 2 VPNs.....	237	as	
link speed.....	634	physical part of interface name.....	52
minimum links.....	635	asynchronous-notification statement.....	916
VLAN IDs.....	600	at	
aggregated SONET interfaces.....	882	physical part of interface name.....	52
configuring.....	881	ATM cell-relay encapsulation	
example configuration.....	885	encapsulation overhead.....	104
firewall filters.....	884	ATM encapsulation.....	281
example configuration.....	884	ATM CCC VC multiplex.....	330
link speed.....	883	ATM cell-relay.....	330
minimum links.....	883	ATM NLPID.....	330
aggregated-ether-options statement.....	910	ATM PVC.....	330
usage guidelines.....	583	ATM PVC encapsulation.....	104, 106
aggregated-sonet-options statement.....	911	ATM SNAP.....	330
usage guidelines.....	881	ATM TCC SNAP.....	330
aging timer		ATM TCC SNAP multiplex.....	330
ARP.....	596	ATM VC multiplex.....	330
alarm triggers, SONET interfaces		cell-relay accumulation mode.....	330
hold timers.....	856	Cisco ATM NLPID.....	330
all (tracing flag)		Ethernet over ATM.....	330
interfaces.....	137	Ethernet VPLS over ATM.....	330
VRRP.....	596	keepalives.....	126
allow-any-vci statement.....	911	Layer 2 switching cross-connects.....	225
allow-fragmentation statement.....	912	Multilink PPP over ATM.....	330
allow-remote-loopback statement.....	912	PPP over ATM.....	330
annex statement.....	913	PPP over ATM multiplex.....	330
usage guidelines.....	364	<i>See also</i> ATM interfaces	
apply-action-profile statement.....	913	ATM interfaces	
APS		ATM overview.....	281
basic support.....	861	communication with ATM switches	
circuit pairs.....	861	example configuration.....	292
circuits, switching between.....	866	encapsulation.....	106
configuration example.....	862	ILMI with cell relay	
link PIC redundancy		example configuration.....	293
example configuration.....	870	MTU sizes.....	98, 99, 100, 101
link state replication.....	869		

PIC type
 example configuration.....296
 PPP over ATM encapsulation
 example configuration.....335
 virtual circuits
 example configuration.....297
 virtual paths
 example configuration.....297
 ATM PVC encapsulation.....104, 106
 atm-encapsulation statement.....916
 ATM-for-ADSL
 example configuration.....795
 atm-options statement.....917
 ATM-over-ADSL
 encapsulation types.....358
 operational mode.....357
 ATM-over-SHDLS.....361
 atm-scheduler-map statement918
 ATM-to-Ethernet interworking.....223, 229, 1330
 VCI range.....1315
 virtual path identifier.....1333
 VLAN tagging.....1062, 1318
 ATM1 and ATM2 IQ
 configuration differences.....287
 ATM1 interfaces
 cell-relay circuit
 example configuration.....333
 shaping values
 example calculation.....324
 ATM2 IQ interfaces
 early packet discard threshold
 example configuration.....328
 eight forwarding classes
 example configuration.....343
 Layer 2 circuit AAL5 mode
 example configuration.....303
 Layer 2 circuit cell-relay mode
 example configuration.....303
 Layer 2 circuit cell-relay promiscuous mode
 example configuration.....308
 Layer 2 circuit trunk mode
 example configuration.....303
 MTU sizes.....100, 101
 authentication-key statement.....918
 APS.....861
 usage guidelines.....861
 authentication-profile-name statement.....919
 authenticator statement.....920
 auto-configure statement.....921
 auto-discovery statement.....921
 auto-negotiation statement
 Gigabit Ethernet.....922
 usage guidelines.....767
 J Series uPIM.....589, 923
 auto-reconnect statement.....923
 usage guidelines.....790

auto-synchronize statement
 usage guidelines.....271
 Automatic Protection Switching *See* APS
 autonegotiation
 configuring manually.....767

B

backup routers
 VRRP.....753
 backup-destination statement.....924
 backup-interface statement.....924
 backup-options statement.....925
 bandwidth statement.....925
 usage guidelines.....159
 bandwidth-limit statement
 hierarchical policer.....926
 policer for Gigabit Ethernet interface.....927
 usage guidelines.....757
 bchannel-allocation statement.....927
 bcm0
 internal Ethernet interface.....32
 physical part of interface name.....52
 bcm0 interface.....246
 bearer-bandwidth-limit statement.....928
 BERT
 configuring interface diagnostics.....134
 bert-algorithm statement.....929
 usage guidelines.....134, 553, 560, 570
 bert-error-rate statement.....931
 usage guidelines.....134
 bert-period statement.....932
 usage guidelines.....134
 bit error rate test *See* BERT
 BOOTP
 accepting packets.....213
 Bootstrap Protocol *See* BOOTP
 borrower interface
 unnumbered Ethernet or demux.....186
 br
 physical part of interface name.....52
 braces, in configuration statements.....lxvi
 brackets
 angle, in syntax descriptions.....lxv
 square, in configuration statements.....lxvi
 Bridge Domain.....684
 bridge network
 trunk interface.....621
 bridge-domain.....684
 bridge-domain statement.....933
 broadcast statement.....933
 usage guidelines.....174
 buildout statement.....934
 bundle statement.....936

burst-size-limit statement	
hierarchical policer.....	937
policer for Gigabit Ethernet interface.....	937
usage guidelines.....	757
byte encoding.....	938
byte-encoding statement.....	938
bytes statement.....	939
usage guidelines.....	849
C	
C-bit parity mode.....	943
c2 SONET header byte.....	849
callback statement.....	940
usage guidelines.....	834
callback-wait-period statement.....	941
usage guidelines.....	834
caller statement.....	942
usage guidelines.....	833
calling-number statement.....	943
usage guidelines.....	821
cau4	
physical part of interface name.....	52
cbit-parity statement.....	943
cbr statement.....	944
CCC.....	223
aggregated Ethernet.....	236, 237
encapsulation	
VLAN-bundled dual-tag logical	
interfaces.....	1326
VLAN-bundled single-tag logical	
interfaces.....	1320
Layer 2 VPNs.....	237
CE PIC E1 interfaces	
configuration statements	
E1 options.....	521
CE PIC T1 interfaces	
configuration statements	
T1 options.....	521
ce1	
physical part of interface name.....	52
cell-bundle-size statement.....	945
Challenge Handshake Authentication Protocol <i>See</i>	
CHAP	
channel part of interface name.....	56
channelized AU-4 interfaces.....	416, 436, 442
example configuration.....	416, 437, 442
channelized E1 interfaces	
example configuration.....	418, 439, 471, 503, 506
interface naming.....	471, 503
channelized E1 IQ and IQE interfaces	
time slots.....	467
channelized E1 IQ interfaces	
example configuration.....	505
time slots.....	502
channelized E1 IQE interfaces	
example configuration.....	505
time slots.....	502
channelized E3 (COC48/STM16) interfaces	
example configuration.....	417
channelized interfaces	
clock sources.....	390
channelized IQ interfaces	
supported options.....	393
channelized NxDS0 IQ interfaces	
example configuration.....	443, 497, 503
channelized NxDS0 IQE interfaces	
example configuration.....	503
channelized OC12 (COC48/STM16 IQE)	
interfaces.....	407
example configuration.....	408
channelized OC12 interfaces	
example configuration.....	445, 446, 451
traffic-shaping rates.....	320, 322, 327
channelized OC12 IQ interfaces	
example configuration.....	448
channelized OC3 (COC48/STM16 IQE) interfaces.....	408
example configuration.....	409
channelized OC3 interfaces.....	429
example configuration.....	429
channelized OC48 interfaces	
SONET header bytes.....	849
channelized OC48 IQE interfaces	
example configuration.....	419
channelized STM1 interfaces	
example configuration.....	475, 785
interface naming.....	785
time slots.....	468, 502
virtual tributary mapping.....	785
channelized STM1 IQ interfaces	
example configuration.....	469
channelized T1 (COC48/STM16 IQE) interfaces	
VT mapping.....	411
channelized T1 interfaces	
example configuration.....	444, 514
VT mapping.....	431, 439
channelized T1 IQ interfaces	
VT mapping.....	444, 460
channelized T1 IQE interfaces	
VT mapping.....	419
channelized T3 (COC48/STM16 IQE) interfaces	
example configuration.....	410
channelized T3 interfaces	
example configuration.....	428
loop timing.....	852
channelized T3 IQ interfaces	
example configuration.....	458
CHAP.....	112, 114
configuring default CHAP secret.....	113
example configuration.....	114, 117

- chap statement.....946
 - usage guidelines.....112, 114
- chap-secret statement.....947
- checksums *See* frame checksums
- ci
 - physical part of interface name.....52
- circuit cross-connect (CCC)
 - encapsulation
 - VLAN-bundled dual-tag logical
 - interfaces.....1326
 - VLAN-bundled single-tag logical
 - interfaces.....1320
- circuit cross-connect CCC.....223
- Circuit Emulation PIC E1 interfaces
 - configuration statements
 - E1 options.....521
- Circuit Emulation PIC T1 interfaces
 - configuration statements
 - T1 options.....521
- Circuit Emulation PICs.....517
- circuit pairs, APS.....861
- Cisco HDLC encapsulation.....106
 - configuring on physical interfaces.....106
 - encapsulation overhead.....104
 - example configuration.....579
 - keepalives.....126
 - Layer 2 switching cross-connect.....225
- cisco-interoperability statement.....947
- class of service *See* CoS
- classifier statement.....948
 - usage guidelines.....759
- clear-dont-fragment-bit statement.....948
- client statement.....949
- clock rates.....270
- clock sources.....64, 128, 875, 950
 - channelized interfaces.....390
- clock-rate statement.....952
 - usage guidelines.....270
- clocking modes.....270
- clocking statement.....950
 - usage guidelines.....64, 128, 875
- clocking-mode statement.....951
 - usage guidelines.....269
- coc1
 - physical part of interface name.....52
- coc12
 - physical part of interface name.....52
- coc3
 - physical part of interface name.....52
- coc48
 - physical part of interface name.....52
- comments, in configuration statements.....lxvi
- communication with ATM switches
 - example configuration.....292
- compatibility-mode statement.....953
- compression.....162
 - compression statement.....954
 - usage guidelines.....120, 121
 - compression-device statement.....955
- Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface.....615
- Configuring a VLAN-Bundled Logical Interface.....614, 616
- Configuring Logical Link-Layer Encapsulation to Support CCCs.....612
- Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs.....612
- connections
 - configuration statements.....24, 27
- connections statement.....956
 - usage guidelines.....228
- connectivity-fault management.....899
- connectivity-fault-management statement.....957
- container interfaces.....37
 - configuration example.....863
- container-devices statement.....958
 - device-count.....958
- container-list statement.....958
- container-options statement.....959
- container-type statement.....959
- continuity-check statement.....960
- control leads
 - serial interfaces.....271
- control-leads statement
 - usage guidelines.....271
- control-polarity statement.....961
 - usage guidelines.....274
- control-signal statement.....962
 - usage guidelines.....271
- conventions
 - text and syntax.....lxv
- copy-tos-to-outer-ip-header statement.....962
- core-dump statement.....963
 - usage guidelines.....138
- CoS
 - eight forwarding classes
 - example configuration.....343
- cp
 - physical part of interface name.....52
- crc-major-alarm-threshold statement.....963
- crc-minor-alarm-threshold statement.....964
- cstm1
 - physical part of interface name.....52
- cstm16
 - physical part of interface name.....52
- cstm4
 - physical part of interface name.....52
- ct1
 - physical part of interface name.....52
- ct3
 - physical part of interface name.....52

cts statement.....	964
usage guidelines.....	271
cts-polarity statement.....	965
usage guidelines.....	274
curly braces, in configuration statements.....	lxvi
current statement.....	965
customer support.....	lxvii
contacting JTAC.....	lxvii

D

damping	
interface transitions.....	138, 877
data circuit-terminating equipment <i>See</i> DCE	
data terminal equipment <i>See</i> DTE	
data-input statement.....	966
database (tracing flag).....	596
dcd statement.....	967
usage guidelines.....	271
dcd-polarity statement.....	967
usage guidelines.....	274
DCE.....	380, 579, 968
dce statement.....	968
usage guidelines.....	380
dce-options statement.....	968
deactivation-delay statement.....	969
usage guidelines.....	824
default router addresses.....	192
default-action statement.....	969
default-chap-secret statement.....	970
default-pap-password statement.....	970
defect triggers, SONET interfaces.....	855
configuration example.....	858
hold timers.....	856
demux	
physical part of interface name.....	52
demux interfaces.....	251
configuring underlying interfaces for.....	252
unnumbered.....	186
demux-destination statement.....	972
demux interfaces.....	973
demux-options statement.....	973
demux-source statement.....	974
demux interfaces.....	975
demux0 statement.....	971
dense wavelength-division multiplexing <i>See</i> DWDM	
description statement.....	976
example configuration.....	96
usage guidelines.....	96, 156
descriptors, interface.....	50
destination statement.....	977
usage guidelines.....	174, 179
destination-class usage	
example configuration.....	217
destination-class-usage statement.....	979
usage guidelines.....	214

destination-profile	
usage guidelines.....	179
destination-profile statement.....	980
usage guidelines.....	174
dfc	
physical part of interface name.....	52
DHCP	
accepting.....	213
dial-options statement.....	982
dial-string statement.....	983
usage guidelines.....	835
dialer statement.....	980
dialer-options	
usage guidelines.....	824
dialer-options statement.....	981
dialin statement.....	982
direction statement.....	983
disable statement.....	984
link protection.....	984
logical interfaces	
usage guidelines.....	167
physical interfaces	
usage guidelines.....	140
disable-mlppp-inner-ppp-pfc statement.....	985
disabling	
keepalives.....	126
logical interfaces.....	167
physical interfaces.....	140
example configuration.....	141
discard interface.....	249
dlci statement.....	985
dlsr statement.....	986
usage guidelines.....	181
do-not-fragment statement.....	986
documentation set	
comments on.....	lxvi
donor interface	
unnumbered Ethernet or demux.....	186
dot1x.....	987
authenticator.....	987
authentication-profile-name.....	987
interface.....	987
<i>See also</i> maximum-requests	
<i>See also</i> quiet-period	
<i>See also</i> reauthentication	
<i>See also</i> retries	
<i>See also</i> supplicant	
<i>See also</i> supplicant-timeout	
<i>See also</i> transmit-period	
configuration statements.....	25
interface	
.....	987
<i>See also</i> server-timeout	
<i>See also</i> server-timeout	
down-count statement.....	988
down-timeout statement.....	989

ds
 physical part of interface name.....52
 ds0-options statement.....989
 usage guidelines.....471, 503
 DS3 IQ interfaces
 MTU sizes.....99, 100
 dsc
 physical part of interface name.....52
 dsl-options statement.....990
 usage guidelines.....356
 dsr statement.....990
 usage guidelines.....271
 dsr-polarity statement.....991
 usage guidelines.....274
 DTE.....579
 dte-options statement.....991
 DTR circuit
 serial interfaces.....274
 dtr statement.....992
 usage guidelines.....271
 dtr-circuit statement.....993
 usage guidelines.....274
 dtr-polarity statement.....993
 usage guidelines.....274
 dual-tag framing
 VLAN ID list.....1326
 dump-on-flow-control statement.....994
 DWDM.....819, 843
 Dynamic Host Configuration Protocol *See* DHCP
 dynamic VLANs.....921
 dynamic-call-admission-control statement.....994
 dynamic-profile statement
 stacked VLAN ranges.....995
 usage guidelines.....163
 VLAN ranges.....996

E

E-LMI.....690, 1104
 e1
 physical part of interface name.....52
 E1 interfaces
 configuration statements.....543
 example configuration.....471, 503
 ITU-T standards.....543
 MTU sizes.....98, 99, 100
 physical interface properties.....543
 time slots
 example configuration.....549
 E1 IQ interfaces
 example configuration.....466
 E1 IQE interfaces
 example configuration.....466
 e1-options statement.....997
 usage guidelines.....471, 503, 543, 785
 e3
 physical part of interface name.....52
 E3 interfaces
 configuration statements.....551
 encapsulation.....106
 ITU-T standards.....551
 MTU sizes.....98, 99, 100
 E3 IQ interfaces
 MTU sizes.....99, 100
 e3-options statement.....998
 usage guidelines.....551
 east-interface statement.....999
 em0
 configuring.....775
 management Ethernet interface.....32, 775
 physical part of interface name.....52
 em1
 internal Ethernet interface.....32
 physical part of interface name.....52
 em1 interface.....246
 em2
 physical part of interface name.....52
 encapsulation.....106
 CCC.....225
 Ethernet 802.3.....104
 example configuration.....111, 880
 extended VLAN CCC.....610, 665
 media MTU size and.....377, 873
 on logical interfaces.....160, 880
 on physical interfaces.....106, 878
 overheads.....104
 See also ATM encapsulation
 encapsulation statement.....1000
 Layer 2 switching cross-connect.....225
 logical interfaces
 usage guidelines.....160, 880
 physical interfaces
 usage guidelines.....106, 878
 encoding
 byte.....938
 line.....277
 encoding statement.....1007
 usage guidelines.....277
 Enhanced IQ (IQE) interfaces
 channelized OC3.....455
 epd-threshold statement.....1008
 usage guidelines.....341
 es
 physical part of interface name.....52
 es-options statement.....1009
 ETH-DM
 configuring routers to support.....717, 724, 732
 displaying statistics and frame
 counts.....720, 729, 732
 overview.....711
 starting an ETH-DM session.....718, 727, 732

Ethernet bridging.....	663, 664	ethernet-switch-profile statement.....	762, 1014
Ethernet CCC and TCC encapsulation		usage guidelines.....	644, 757
physical interfaces.....	106	eui-64 statement.....	1017
Ethernet CCC encapsulation		usage guidelines.....	174
encapsulation overhead.....	104	evcs statement.....	1015
Ethernet configurations, example.....	813	event (tracing flag).....	137
Ethernet frame delay measurement		event statement.....	1016
configuring routers to support.....	717, 724, 732	interface-status-tlv statement.....	1075
displaying statistics and frame		port-status-tlv statement.....	1196
counts.....	720, 729, 732	event-thresholds	
overview.....	711	frame-error statement.....	1032
starting an ETH-DM session.....	718, 727, 732	frame-period statement.....	1033
Ethernet interfaces.....	663	frame-period-summary statement.....	1034
802.1ag OAM.....	679	symbol-period statement.....	1267
802.3ah OAM.....	745	event-thresholds statement.....	1016
configuration statements.....	583	extended VLAN	
dynamic VLANs.....	921	CCC	
example configuration.....	813	applying.....	610
Fast Ethernet interfaces.....	583	encapsulation.....	106
Gigabit Ethernet interfaces.....	583	encapsulation overhead.....	104
gratuitous ARP.....	596	example configuration.....	611
internal Ethernet interface.....	32, 245	TCC	
management Ethernet interface.....	775	applying.....	665
mixed VLAN tagging.....	603, 1027, 1144	encapsulation.....	106, 665
passive monitoring.....	677	encapsulation overhead.....	104, 105
proxy ARP, unrestricted.....	671	external clock sources.....	128, 875, 950
static ARP table entries.....	669	external synchronization interface	
unnumbered.....	186	usage guidelines.....	65
preferred source address.....	1201		
VLAN IDs.....	1316	F	
VLAN tagging.....	599, 665, 1324	f-max-period statement.....	1028
VRRP.....	753	f1 SONET header byte.....	849
Ethernet link aggregation.....	625	f2 SONET header byte.....	849
Ethernet Local Management Interface <i>See</i> E-LMI		facility-override statement.....	1017
Ethernet over ATM encapsulation		fail-filter statement.....	1231
encapsulation overhead.....	104	usage guidelines.....	208
physical interfaces.....	106	failover-delay statement.....	754, 1018
Ethernet Ring Protection		family bridge	
configuration statements.....	27	VLAN ID list.....	1319
Ethernet Ring Protection Switching, Configuring.....	799	VLAN IDs.....	1317
Ethernet Service OAM	711	family descriptors.....	50
Ethernet SNAP encapsulation		family statement.....	1019
encapsulation overhead.....	104	usage guidelines.....	172
ethernet statement.....	1010	family type	
Ethernet switching.....	663, 664	specifying for underlying interfaces.....	253
Ethernet switching interfaces.....	589	Fast Ethernet interfaces	
Ethernet TCC		configuration statements.....	583
applying.....	666	dynamic VLANs.....	921
encapsulation.....	665	Ethernet link aggregation.....	625
encapsulation overhead.....	104	example configuration.....	813
example configuration.....	667	ignoring Layer 3 incomplete errors.....	594
Ethernet VLAN circuit		ingress rate-limit.....	597
VLAN ID list.....	1320	interface speed.....	123
ethernet-policer-profile statement.....	1012	link modes.....	97, 595
usage guidelines.....	757	link protection.....	626
ethernet-ring statement.....	1013		

- loopback mode.....593
 - MAC address filtering.....591
 - MTU sizes.....98, 99, 100
 - physical interface properties.....583
 - proxy ARP, unrestricted.....671
 - speed.....597
 - static ARP table entries.....669
 - usage guidelines.....665
 - VLAN IDs.....600, 1316
 - VLAN tagging.....599, 665, 1324
 - VRRP.....753
 - fastether-options statement.....1023
 - usage guidelines.....583
 - fcs statement.....1024
 - SONET interfaces
 - usage guidelines.....851
 - fe
 - physical part of interface name.....52
 - feac-loop-respond statement.....1025
 - filter statement.....1026
 - usage guidelines.....203
 - filters.....884
 - on aggregated links.....884
 - example configuration.....884
 - See also* firewall filters
 - firewall
 - hierarchical-policer.....1043
 - firewall filters
 - applying.....203
 - example configuration.....206
 - logical interfaces.....203
 - flexible-vlan-tagging statement.....1027
 - flow control.....594
 - flow-control statement.....1028
 - usage guidelines.....594
 - font conventions.....lxv
 - force statement.....1029
 - forwarding-class statement.....1030
 - usage guidelines.....341, 759
 - fractional E1 IQ and IQE interfaces
 - time slots.....467
 - fractional E1 IQ interfaces
 - example configuration.....468, 502
 - time slots.....502
 - fractional E1 IQE interfaces
 - example configuration.....502
 - fractional T1 (COC48/STM16 IQE) interfaces
 - example configuration.....412
 - time slots.....412
 - fractional T1 interfaces
 - example configuration.....434
 - time slots.....433
 - fractional T1 IQ interfaces
 - example configuration.....462, 481, 496
 - fragment-threshold statement
 - usage guidelines.....831
 - fragment-threshold statement (interfaces).....1031
 - frame checksums
 - SONET interfaces.....851
 - frame delay, Ethernet *See* Ethernet frame delay
 - measurement
 - Frame Relay encapsulation.....106
 - DCE.....380, 968
 - encapsulation overhead.....104
 - example configuration.....374, 579
 - Frame Relay protocol.....371
 - keepalives.....126
 - Layer 2 switching cross-connect.....226
 - logical interfaces.....880
 - physical interfaces.....106, 878
 - Frame Relay ether type encapsulation
 - physical interfaces.....106
 - Frame Relay Ether Type encapsulation.....374
 - physical interfaces.....878
 - frame-error statement.....1032
 - frame-period statement.....1033
 - frame-period-summary statement.....1034
 - framing statement.....1035
 - channelized OC12 IQ interfaces
 - usage guidelines.....440
 - channelized OC12 IQE interfaces
 - usage guidelines.....434
 - channelized OC48 IQE interfaces
 - usage guidelines.....415
 - usage guidelines.....781, 846
 - fxp
 - physical part of interface name.....52
 - fxp0
 - management Ethernet interface.....32
 - physical part of interface name.....52
 - fxp0 interface
 - configuring.....775
 - management Ethernet interface.....775
 - fxp1
 - internal Ethernet interface.....32
 - physical part of interface name.....52
 - fxp1 interface.....245
 - fxp2
 - internal Ethernet interface.....32
 - physical part of interface name.....52
- ## G
- ge
 - physical part of interface name.....52
 - ge interface
 - configuring.....773
 - general (tracing flag).....596
 - Gigabit Ethernet interfaces.....591, 663
 - 802.1ag679
 - 802.3ah745
 - autonegotiation.....767

configuration statements.....	583
dynamic VLANs.....	921
Ethernet link aggregation.....	625
example configuration.....	813
flow control.....	594
ignoring Layer 3 incomplete errors.....	594
link protection.....	626
loopback mode.....	593
MAC address filtering.....	591
MTU sizes.....	98, 99, 100
proxy ARP, unrestricted.....	671
static ARP table entries.....	669
usage guidelines.....	665
VLAN IDs.....	600, 1316, 1322
VLAN tagging.....	599, 665, 1062, 1318, 1324
VRRP.....	753
Gigabit Ethernet IQ interfaces	
configuring.....	755
MAC address accounting.....	766
MAC address filtering.....	761
policer	
example configuration.....	762
rate limiting.....	757
Gigabit Ethernet OTN.....	773
Gigabit Ethernet uPIM interfaces	
speed.....	597
Gigabit Ethernet, IQ, IQ2 and IQ2-E interfaces	
VLAN tag stacking and rewriting.....	641
gigether-options statement.....	1038
usage guidelines.....	583
gr	
physical part of interface name.....	52
gratuitous ARP.....	596
gratuitous-arp-reply statement.....	1039
usage guidelines.....	596
gre	
physical part of interface name.....	52
group option	
firewall filters.....	203
H	
hardware-assisted-timestamping statement.....	1040
HDLC encapsulation	
Cisco HDLC encapsulation.....	106
HDLC payload scrambling	
SONET interfaces.....	854
header byte values.....	849
hello-timer statement.....	1041
hierarchical policer	
premium.....	1202
hierarchical-policer.....	1043
high-plp-max-threshold statement.....	1041
high-plp-threshold statement.....	1042
hold-interval statement	
connectivity-fault management.....	1044
Ethernet ring protection switching.....	1044
hold-time statement.....	1045
APS	
usage guidelines.....	869
damping interface transitions	
usage guidelines.....	138, 877
SONET defect triggers	
usage guidelines.....	856
host statement.....	1049
I	
icons defined, notice.....	lxiv
idle-cycle-flag statement.....	1050
serial interfaces	
usage guidelines.....	271
idle-time statement.....	1051
usage guidelines.....	180
idle-timeout	
usage guidelines.....	824
idle-timeout statement.....	1051
IEEE 802.1p policer profile	
usage guidelines.....	758
ieee802.1p statement.....	1052
if-exceeding statement	
hierarchical policer.....	1052
ignore statement.....	1053
ignore-all statement.....	1053
usage guidelines.....	271
ignore-l3-incompletes statement.....	1054
ignoring Layer 3 incomplete errors.....	594
ilmi statement.....	1054
usage guidelines.....	292
ILMI with cell relay	
encapsulation types supported.....	292, 331
example configuration.....	293
inactivity-timeout statement.....	1055
incoming-called-number statement.....	1055
usage guidelines.....	823
incoming-map statement.....	1056
usage guidelines.....	833
incrementing STM ID.....	850
indication statement.....	1057
usage guidelines.....	271
indication-polarity statement.....	1057
usage guidelines.....	274
inet protocol family	
interface addresses.....	176
inet6 protocol family	
interface addresses.....	176
ingress-rate-limit statement.....	1058
usage guidelines.....	597
init-command-string statement.....	1059
usage guidelines.....	93

- initial-route-check statement.....1060
- inner-tag-protocol-id statement.....1060
 - usage guidelines.....645
- inner-vlan-id statement.....1061
 - usage guidelines.....645
- inner-vlan-id-range statement.....1062
- input option
 - firewall filters.....203
 - policers.....194
- input statement.....1062
- input-list statement.....1063
 - usage guidelines.....203
- input-policer statement.....1063
- input-priority-map statement.....1064
 - usage guidelines.....758
- input-three-color statement.....1064
- input-vlan-map statement.....1065
 - usage guidelines.....641
- instance.....684
- Instance.....684
- instance statement.....1066
- integrated routing and bridging interfaces *See* IRB
- interface.....1069
 - maximum-requests.....1069
 - quiet-period.....1069
 - reauthentication.....1069
 - retries.....1069
 - server-timeout.....1069
 - supplicant.....1069
 - supplicant-timeout.....1069
 - transmit-period.....1069
- interface addresses
 - logical interfaces.....174
 - preferred interface addresses.....192, 194
 - primary interface addresses.....192, 193
- interface groups.....205
- interface naming
 - chassis, routing matrix.....61
 - routing matrix based on a TX Matrix Plus
 - router.....59
 - routing matrix based on a TX Matrix router.....57
 - TX Matrix Plus router.....59
 - TX Matrix router.....57
- interface preservation.....869
- interface statement.....1067
 - CoS.....1068
 - DLSw.....1067
 - IEEE 802.1ag.....1070
 - IEEE 802.3ah.....1071
 - usage guidelines.....621
- interface transitions
 - damping.....138, 877
- interface-down statement.....969, 1072
- interface-mode statement.....1073
 - usage guidelines.....619, 620
- interface-set statement.....1074
- interface-status-tlv statement.....1075
- interface-switch statement.....1075
 - usage guidelines.....224
- interface-type statement.....1076
 - channelized (COC48/STM16 IQE) interfaces
 - usage guidelines.....407, 408
 - channelized E1 (COC48/STM16) IQE interfaces
 - usage guidelines.....418
 - channelized E1 IQE interfaces
 - usage guidelines.....438
 - channelized E3 (COC48/STM16) IQE interfaces
 - usage guidelines.....417
 - channelized E3 interfaces
 - usage guidelines.....437
 - channelized OC3 interfaces
 - usage guidelines.....436
 - channelized OC3 IQ interfaces
 - usage guidelines.....429, 442
 - channelized OC48 IQE interfaces
 - usage guidelines.....416
 - channelized T1 (COC48/STM16 IQE) interfaces
 - usage guidelines.....410
 - channelized T1 interfaces
 - usage guidelines.....429
 - channelized T1 IQ interfaces
 - usage guidelines.....443, 466
 - channelized T3 IQ interfaces
 - usage guidelines.....443
 - E1 (COC48/STM16) IQE interfaces
 - usage guidelines.....418
 - NxDS0 (COC48/STM16 IQE) interfaces
 - usage guidelines.....413
 - NxDS0 interfaces
 - usage guidelines.....431
- interfaces.....62
 - 10-Gigabit Ethernet DWDM.....779
 - 10-Gigabit Ethernet framing.....781
 - aggregated Ethernet.....623
 - aggregated SONET.....881
 - channelized E1 PRI.....509
 - channelized T1 PRI.....509
 - clock sources.....64, 128, 950
 - configuration statements.....4, 889
 - container interfaces.....37
 - demux underlying.....252
 - descriptive text.....976
 - descriptors.....50
 - disabling.....140, 167
 - display order in configurations.....64
 - encapsulation *See* encapsulation
 - firewall filters firewall filters.....203
 - Gigabit Ethernet
 - configuring.....767
 - Gigabit Ethernet IQ.....755
 - Gigabit Ethernet IQ policer
 - example configuration.....762

mixed VLAN tagging.....	603, 1027, 1144
names.....	51
overview.....	819, 843
permanent interfaces.....	32
physical.....	67
services interfaces.....	36
transient interfaces.....	32, 35
interfaces (tracing flag).....	596
interfaces statement.....	1072
interleave-fragments statement.....	1077
internal clock sources.....	128, 875, 950
internal Ethernet interface.....	32, 245
Internet Protocol Control Protocol	<i>See</i> IPCP
interval statement.....	1078
inverse-arp statement.....	1079
invert-data statement.....	1080
ip	
physical part of interface name.....	52
IP addresses	
128-bit.....	176
32-bit.....	176
IPCP.....	178
management Ethernet interface.....	775
mapping to MAC address.....	669, 671
unnumbered interfaces.....	185
ipc (tracing flag).....	137
IPCP.....	177
assigning PPP properties.....	179
configuring IP address.....	178
negotiating IP addresses.....	178
unnumbered interfaces.....	179
ipip	
physical part of interface name.....	52
ipsec-sa statement.....	1080
IPv4 Protocol family	
interface addresses.....	174
on logical interfaces.....	172
IPv6.....	174
standards documents.....	174
transition.....	174
IPv6 Protocol family	
on logical interfaces.....	172
IQ interfaces.....	385
channelized.....	385
channelized E1.....	501
channelized OC12.....	423
channelized OC3.....	455
channelized STM1.....	785
channelized STM4.....	423
channelized T1.....	495
channelized T3.....	479
Gigabit Ethernet.....	755
IQE interfaces.....	385
channelized.....	385
channelized E1.....	501
channelized OC12.....	423
channelized OC48.....	405
channelized OC48/STM16.....	405
channelized STM1.....	785
channelized STM16.....	405
channelized STM4.....	423
channelized T1.....	495
ISDN interfaces	
callback-wait-period.....	941
caller.....	942
calling-number.....	943
configuration (called router).....	839
configuration (calling router).....	837
redial-delay.....	1220
isdn-options statement.....	1081
usage guidelines.....	821
ISO Protocol family.....	172, 174
ITU-T Recommendation Y.1731.....	711
ITU-T standards	
E1 interfaces.....	543
E3 interfaces.....	551
Y.1731 ETH-DM.....	711
ixgbe0	
internal Ethernet interface.....	32
physical part of interface name.....	52
ixgbe0 interface.....	246
ixgbe1	
internal Ethernet interface.....	32
physical part of interface name.....	52
ixgbe1 interface.....	246
J	
J Series Routers.....	589
jitter, Ethernet frame	<i>See</i> Ethernet frame delay
measurement	
K	
keep-address-and-control statement.....	1082
usage guidelines.....	192
keepalives	
disabling.....	126
keepalives statement.....	1083
usage guidelines.....	126
key statement.....	1084
L	
l2tp-interface-id statement.....	1084
LACP	
Ethernet aggregation.....	627
example configuration.....	631
interval.....	628
traceoptions.....	631
tracing operations.....	631

- lACP statement.....1086
 - 802.3ad.....1085
 - usage guidelines.....627
- LAN PHY.....781, 872
- Layer 2 bridging
 - Ethernet.....664
- Layer 2 circuit cell-relay promiscuous mode
 - ATM2 IQ interfaces
 - example configuration.....308
- Layer 2 circuit transport mode
 - ATM2 IQ interfaces
 - example configuration.....303
- Layer 2 switching cross-connect
 - CCC connections.....228
 - CCC encapsulation.....225
 - example configuration.....233
 - MPLS.....229
 - router configuration.....224
- Layer 2.5 VPNs
 - Ethernet.....610, 665, 666
- Layer 3 incomplete errors.....1054
- layer2-policer statement.....1087
- lc
 - physical part of interface name.....52
- LCC.....57
 - in a routing matrix based on a TX Matrix Plus
 - router.....59
- LCP
 - address and control field compression.....120
 - protocol field compression.....121, 162
- lcp-max-conf-req.....1088
- lcp-restart-timer statement.....1088
 - usage guidelines.....162
- leaky bucket properties.....129, 876
- level statement.....1089
- line.....277
 - protocol.....265
- line-card chassis *See* LCC
- line-encoding statement.....1091
- line-protocol statement.....1091
 - usage guidelines.....265
- line-rate statement.....1092
 - usage guidelines.....364
- line-vlan-id statement
 - usage guidelines.....641
- linear-red-profile statement.....1089
 - usage guidelines.....341
- linear-red-profiles statement.....1090
- link aggregation.....625
 - SONET interfaces.....882
- link modes.....97, 595
- link PIC failover
 - channelized OC12 IQ interfaces.....447
 - channelized OC48 IQE interfaces.....419
 - channelized STM1 IQ and IQE PICs.....475
- link protection
 - aggregated Ethernet interfaces.....626
 - disable statement.....984
 - non-revertive statement.....1150
 - revertive statement.....1227
- link speed
 - Ethernet aggregation.....634
 - SONET aggregation.....883
- link state replication
 - APS.....869
- link-adjacency-loss statement.....1092
- link-discovery statement.....1093
- link-down statement.....1093
- link-event-rate statement.....1094
- link-fault management.....900
- link-fault-management statement.....1095
- link-layer-overhead statement.....1096
- link-mode statement.....1097
 - usage guidelines.....97, 123, 595
- link-protection statement.....1098
- link-speed statement.....1099
 - usage guidelines.....634, 883
- linktrace statement.....1100
- llc2 statement.....1101
 - usage guidelines.....180
- LMI packets *See* keepalives
- lmi statement
 - Ethernet OAM.....1104
 - Frame Relay keepalives.....1103
- lmi-type statement.....1105
- lo
 - physical part of interface name.....52
- load sharing
 - APS.....870
- load-interval statement.....1105
 - usage guidelines.....824
- load-threshold statement.....1106
 - usage guidelines.....824
- Local Management Interface packets *See* keepalives
- local name, configuring.....113, 116, 165
- local password, configuring.....116, 165
- local-mac statement.....1106
- local-name statement.....1107
 - usage guidelines.....112, 114
- local-password statement.....1108
- local-window statement.....1108
 - usage guidelines.....180
- lockout statement.....1109
 - usage guidelines.....866
- log-prefix statement.....1110
- logical interface properties, statements for.....147
- logical interfaces
 - accounting profiles.....158
 - example configuration.....158
 - CCC.....223, 233
 - clear loopback detected timer.....162

configuration statements.....	143
default router addresses.....	192
descriptive text.....	156
descriptors.....	50
disabling.....	167
encapsulation <i>See</i> encapsulation	
example configuration	
dynamic CAC.....	167
firewall filters.....	203
interface addresses.....	174, 192
interface bandwidth.....	159
Layer 2 switching cross-connect.....	224
logical interface properties.....	143
multipoint connections.....	157
point-to-point connections.....	157
policers.....	194
PPP over ATM encapsulation	
example configuration.....	335
PPP restart timers.....	162
preferred interface addresses.....	192, 194
primary interface addresses.....	192, 193
primary router addresses.....	192
primary router interfaces.....	193
protocol families.....	169, 172
protocol MTU.....	191
protocol redirect messages.....	192
SNMP notifications.....	159
static ARP table entries.....	669, 671
unit numbers.....	155
unnumbered interfaces.....	185
VLAN IDs.....	1316, 1322
VLAN-bundled	
dual-tag.....	1326
single-tag.....	1320
logical part of interface name.....	56
logical routers <i>See</i> logical systems	
logical systems	
configuration statements.....	19
interfaces.....	155
logical-interface-policer statement.....	1109
logical-systems statement.....	1110
usage guidelines.....	155
long-buildout statement.....	1111
loop timing	
channelized T3 interfaces.....	852
loop-timing statement.....	1116
usage guidelines.....	852
loopback capability.....	275, 852
E1 interfaces	
example configuration.....	548
E3 interfaces	
example configuration.....	556
serial interfaces	
example configuration.....	276
SONET interfaces	
example configuration.....	853
loopback mode.....	593
loopback statement.....	1112
10-Gigabit Ethernet interfaces	
usage guidelines.....	872
Fast Ethernet interfaces	
usage guidelines.....	593
Gigabit Ethernet interfaces	
usage guidelines.....	593
serial interfaces	
usage guidelines.....	275
SONET interfaces	
usage guidelines.....	852
loopback testing.....	131
loopback-clear-timer statement.....	1115
usage guidelines.....	162
loss-priority statement.....	1116
usage guidelines.....	759
loss-threshold statement.....	1117
low-plp-max-threshold statement.....	1117
low-plp-threshold statement.....	1118
ls	
physical part of interface name.....	52
lsi	
physical part of interface name.....	52
lsq-failure-options statement.....	1118
M	
MAC address accounting	
Gigabit Ethernet IQ interfaces.....	766
MAC address filtering	
Fast Ethernet interfaces.....	591
Gigabit Ethernet interfaces.....	591
Gigabit Ethernet IQ interfaces.....	761
MAC addresses	
management Ethernet interface.....	591, 777
mapping to IP addresses.....	669, 671
mac statement.....	1119
usage guidelines.....	777
mac-address statement.....	1119
mac-learn-enable statement.....	1120
usage guidelines.....	766
mac-validate statement.....	1121
Maintenance Intermediate Points.....	683
Bridge Domain.....	684
Instance.....	684
MIP.....	683
MIP Half Function.....	684
Routing Instance.....	684
maintenance-association statement.....	1122
maintenance-domain	
mip-half-function.....	1131
virtual-switch.....	1315
maintenance-domain statement.....	1123

- management Ethernet interface
 - configuring.....775
 - configuring for M Series and T Series
 - routers.....122
 - configuring J Series Gigabit Ethernet
 - interfaces.....123
 - interface speed.....122
 - IP address.....775
 - link modes.....97, 595
 - MAC address.....777
 - overview.....32
 - speed.....597
 - manuals
 - comments on.....lxvi
 - map statement.....1124
 - master routers
 - VRRP.....753
 - master-only statement.....1125
 - usage guidelines.....775
 - max-retry statement.....1127
 - usage guidelines.....180
 - maximum transmission unit *See* MTU
 - maximum-contexts statement.....1125
 - maximum-requests statement.....1126
 - maximum-vcs statement.....1126
 - media (tracing flag).....137
 - media MTUs.....98
 - SONET interfaces.....873
 - See also* MTU
 - member-interface-speed statement.....1127
 - member-interface-type statement.....1128
 - member-interface-speed.....1128
 - mep statement.....1129
 - minimum links for aggregation
 - Ethernet links.....635
 - SONET links.....883
 - minimum-links statement.....1130
 - usage guidelines.....635, 883
 - MIP Half Function.....684
 - mip-half-function.....684, 1131
 - mixed VLAN tagging.....603, 1027, 1144
 - ml
 - physical part of interface name.....52
 - mlfr-uni-nni-bundle-options statement.....1132
 - mo
 - physical part of interface name.....52
 - mode statement.....1133
 - usage guidelines.....210
 - modem-options statement.....1133
 - usage guidelines.....93
 - monitor-session statement.....1134
 - usage guidelines.....118
 - monitoring services interfaces
 - physical interface properties.....138
 - MPLS
 - Layer 2 switching cross-connect.....229
 - protocol family.....172, 174
 - mpls statement
 - 10-Gigabit Ethernet interfaces
 - usage guidelines.....872
 - ATM interfaces.....1134
 - SONET/SDH interfaces.....1134
 - usage guidelines.....874
 - mrru statement.....1135
 - usage guidelines.....830
 - ms
 - physical part of interface name.....52
 - MSP.....859
 - mt
 - physical part of interface name.....52
 - MTU
 - physical interfaces.....98
 - mtu statement.....1136
 - logical interfaces
 - usage guidelines.....191
 - SONET interfaces
 - usage guidelines.....873
 - usage guidelines.....98
 - mtun
 - physical part of interface name.....52
 - MTUs
 - logical interfaces.....191
 - media MTUs.....873
 - physical interfaces.....873
 - protocol MTUs.....191
 - SONET interfaces.....873
 - multicast-dlci statement.....1137
 - multicast-only statement.....1138
 - usage guidelines.....173
 - multicast-vci statement.....1138
 - multilink-max-classes statement.....1139
 - multiplex section protection *See* MSP
 - multipoint connections.....157, 317
 - See also* point-to-multipoint connections
 - multipoint statement.....1139
 - usage guidelines.....157
 - multipoint-destination statement.....1140
 - Multiprotocol Label Switching *See* MPLS
 - multiservice-options
 - dump-on-flow-control.....994
 - multiservice-options statement.....1141
 - usage guidelines.....138
 - MX Series Routers.....591
- N**
- n391 statement.....1141
 - n392 statement.....1142
 - n393 statement.....1142
 - name-format statement.....1143

names.....	62	channelized T3 interfaces	
of interfaces.....	51	usage guidelines.....	427
native-vlan-id statement.....	1144	channelized T3 IQ interfaces	
nbp-max-conf-req.....	1145	usage guidelines.....	465
nbp-restart-timer statement.....	1145	no-payload-scrambler statement.....	1179
usage guidelines.....	162	SONET interfaces	
negotiate-address statement.....	1146	usage guidelines.....	854
usage guidelines.....	178, 790	no-preempt statement.....	1199
negotiating IP addresses		no-redirects statement.....	1154
IPCP.....	178	usage guidelines.....	192
negotiation-options		no-source-filtering statement.....	1254
allow-remote-loopback statement.....	912	usage guidelines.....	591, 761
no-allow-link-events statement.....	1147	no-syslog statement.....	1269
negotiation-options statement.....	1146	usage guidelines.....	138
neighbor statement.....	1147	no-termination-request statement.....	1155
usage guidelines.....	861	no-translate-discard-eligible statement.....	1291
no unidirectional statement.....	1302	no-translate-fecn-and-beqn statement.....	1291
no-allow-link-events statement.....	1147	no-traps statement.....	1295
no-asynchronous-notification statement.....	916	usage guidelines.....	139, 159
no-auto-mdix statement.....	1148	no-unframed statement.....	1302
no-auto-negotiation statement		no-z0-increment statement.....	1338
Gigabit Ethernet.....	922	usage guidelines.....	850
usage guidelines.....	767	non-revertive statement.....	1150
J Series uPIM.....	589, 923	nonconfigurable interfaces.....	32
no-cbit-parity statement.....	943	normal (tracing flag)	
no-core-dump statement.....	963	arp aging timer.....	596
usage guidelines.....	138	notice icons defined.....	lxiv
no-feac-loop-respond statement.....	1025	NxDS0 (COC48/STM16) IQE interfaces	
no-flow-control statement.....	1028	usage guidelines.....	418
usage guidelines.....	594	NxDS0 interfaces	
no-gratuitous-arp-reply statement.....	1039	usage guidelines.....	439
usage guidelines.....	596	NxDS0 IQ interfaces	
no-gratuitous-arp-request statement.....	1149	example configuration.....	469, 482
usage guidelines.....	596	usage guidelines.....	444
no-keepalives statement.....	1149		
usage guidelines.....	126		
no-long-buildout statement.....	1111		
no-loop-timing statement.....	1116		
usage guidelines.....	852		
no-loopback statement.....	1112, 1114		
usage guidelines.....	593		
no-mac-learn-enable statement.....	1120		
usage guidelines.....	766		
no-partition statement.....	1151		
channelized E1 IQ interfaces			
usage guidelines.....	501		
channelized E1 IQE interfaces			
usage guidelines.....	501		
channelized OC12 IQ interfaces			
usage guidelines.....	435, 441		
channelized OC12/STM4 interfaces			
usage guidelines.....	427		
channelized OC3 IQ interfaces			
usage guidelines.....	457, 510		
channelized T3 (COC48/STM16 IQE) interfaces			
usage guidelines.....	409		

O

OAM	
configuration statements.....	25
E-LMI.....	690
OAM cells <i>See</i> keepalives	
oam statement.....	1156
oam-liveness statement.....	1158
usage guidelines.....	315, 329
oam-period statement.....	1159
usage guidelines.....	315
oc-slice statement.....	1160
usage guidelines.....	409, 427, 436, 441
oc12	
physical part of interface name.....	52
OC12 interfaces	
example configuration.....	445
traffic-shaping rates.....	320, 322, 327
oc3	
physical part of interface name.....	52

OC3 interfaces.....320, 322, 327
 example configuration.....881
 oc48
 physical part of interface name.....52
 OC48 interfaces
 SONET header bytes.....849
 OC48 PIC
 framing.....846
 open-timeout statement.....1160
 operating-mode statement.....1161
 usage guidelines.....357
 optics-options statement.....1162
 usage guidelines.....779
 otn-options.....773
 otn-options statement.....1163
 output option
 firewall filters.....203
 policers.....194
 output statement.....1165
 output-list statement.....1165
 usage guidelines.....203
 output-policer statement.....1166
 output-priority-map statement.....1166
 usage guidelines.....759
 output-three-color statement.....1167
 output-vlan-map statement.....1168
 usage guidelines.....641
 overflow option.....129
 overflow statement.....1170
 usage guidelines.....129, 876

P

p-bit-timeout statement.....1180
 packets (tracing flag).....596
 paired-group statement.....1171
 usage guidelines.....870
 PAP
 configuring default PAP password.....165
 pap statement.....1172
 pap-password statement.....1173
 parentheses, in syntax descriptions.....lxvi
 partition statement.....1174
 channelized AU-4 (on OC48 IQE) interfaces
 usage guidelines.....416
 channelized AU-4 interfaces
 usage guidelines.....436, 442
 channelized E1 (COC48/STM16) IQE interfaces
 usage guidelines.....418
 channelized E1 IQE interfaces
 usage guidelines.....438
 channelized E3 (COC48/STM16) IQE interfaces
 usage guidelines.....417
 channelized E3 interfaces
 usage guidelines.....437

channelized OC12 (COC48/STM16 IQE) interfaces
 usage guidelines.....407
 channelized OC3 (COC48/STM16 IQE) interfaces
 usage guidelines.....408
 channelized OC3 interfaces
 usage guidelines.....429
 channelized T1 (COC48/STM16 IQE) interfaces
 usage guidelines.....410
 channelized T1 interfaces
 usage guidelines.....429
 channelized T1 IQ interfaces
 usage guidelines.....443, 466
 channelized T3 IQ interfaces
 usage guidelines.....443
 clear channel STM1, STM4, and STM16 interface
 usage guide.....416
 E1 (COC48/STM16) IQE interfaces
 usage guidelines.....418
 NxDS0 (COC48/STM16 IQE) interfaces
 usage guidelines.....413
 NxDS0 (COC48/STM16) IQE interfaces
 usage guidelines.....418
 NxDS0 interfaces
 usage guidelines.....431, 439
 NxDS0 IQ interfaces
 usage guidelines.....444
 SDH (VC-4) interfaces
 usage guidelines.....436, 441
 passive access, configuring.....114, 117, 166
 passive monitoring flow
 Ethernet interfaces.....677
 SONET/SDH interfaces.....874
 passive statement.....1175
 usage guidelines.....112, 114, 627
 passive-monitor-mode statement.....1176
 usage guidelines.....677, 874
 path trace identifier.....853
 path-database-size statement.....1177
 path-trace statement.....1178
 10-Gigabit Ethernet interfaces
 usage guidelines.....872
 usage guidelines.....853
 payload-scrambler statement.....1179
 SONET interfaces
 usage guidelines.....854
 payload-size statement.....1180
 pd
 physical part of interface name.....52
 pdu-interval statement.....1181
 pdu-threshold statement.....1181
 pe
 physical part of interface name.....52
 peer statement.....1182
 peer-unit statement.....1182
 per-unit-scheduler statement.....1184
 performance-monitoring statement.....1183

periodic statement.....	1183	time slots.....	468
usage guidelines.....	628	fractional E1 IQ and IQE interfaces.....	467
permanent interfaces.....	32	fractional E1 IQ interfaces.....	502
permanent virtual circuit encapsulation <i>See</i> ATM PVC		fractional T1 interfaces.....	433
pfc statement.....	1185	fractional T1 IQE interfaces.....	412
usage guidelines.....	121	NxDS0 IQ interfaces.....	502
physical interfaces.....	67	transitions, damping.....	138
accounting profiles.....	130	unidirectional mode	
example configuration.....	131	unidirectional link mode.....	139
APS.....	859	unidirectional mode, APS.....	867
byte encoding.....	938	VLAN tagging.....	599, 665, 1324
C-bit parity mode.....	943	<i>See also</i> names of specific interfaces	
clock rates.....	270	physical part of interface name.....	52
clock sources.....	64, 128, 875, 950	pic-type statement.....	1185
clocking mode.....	270	pimd	
configuration statements.....	67	physical part of interface name.....	52
DCE.....	380, 968	pime	
descriptive text.....	96, 976	physical part of interface name.....	52
descriptors.....	50	plp-to-clp statement.....	1187
disabling.....	140	plp1 statement.....	1186
DTR circuit.....	274	usage guidelines.....	341
dynamic VLANs.....	921	point-to-point connections	
encapsulation.....	106	logical interfaces.....	157
Ethernet link aggregation.....	625	unnumbered Ethernet interfaces.....	185
flow control.....	594	Point-to-Point Protocol encapsulation.....	106
frame checksums <i>See</i> frame checksums		encapsulation overhead.....	104
HDLC payload scrambling.....	854	keepalives.....	126
header byte values.....	849	on physical interfaces.....	106, 878
idle cycle flag.....	271	PPP CHAP.....	112, 114
keepalives.....	126	point-to-point statement.....	1187
leaky bucket properties.....	129, 876	usage guidelines.....	157
line encoding.....	277	policer statement	
line protocol.....	265	CoS.....	1188
link modes.....	97, 595	interface MAC.....	1190
loop timing.....	852	usage guidelines.....	194, 757
loopback capability		policers	
serial interfaces.....	275	applying.....	194
SONET interfaces.....	852	arp option.....	194
loopback mode.....	593	input option.....	194
MAC address filtering.....	591	logical interfaces.....	194
media MTU size.....	98, 873	output option.....	194
mixed VLAN tagging.....	603, 1027, 1144	policing	
MSP.....	859	Gigabit Ethernet IQ interfaces.....	757
path trace identifier.....	853	pool statement.....	1191
physical interface properties.....	67	usage guidelines.....	821
PPP CHAP.....	112, 114	pop statement	
revertive mode, APS.....	867	Gigabit Ethernet IQ interfaces.....	1192
RFC 2615 support.....	855	Gigabit Ethernet, IQ, IQ2 and IQ2-E interfaces	
signal polarity.....	274	usage guidelines.....	641
SNMP notifications.....	139	usage guidelines.....	649
SONET defect triggers.....	855	pop-all-labels statement.....	1193
configuration example.....	858	usage guidelines.....	874
hold timers.....	856	pop-pop statement	
SONET link aggregation.....	882	Gigabit Ethernet IQ interfaces.....	1194
speed.....	122, 123, 124, 847	usage guidelines.....	649
statements.....	77		

- pop-swap statement
 - Gigabit Ethernet IQ interfaces.....1194
 - usage guidelines.....650
- port statement
 - voice services.....1195
- Port-Based Network Access Control Protocol
 - IEEE 802.1x
 - dot1x.....741
- port-priority statement
 - LACP.....1195
- port-status-tlv statement.....1196
- post-service-filter statement.....1196
- PPP
 - address and control field compression.....120
 - dynamic-profile.....163
 - PPP clear loopback detected timer.....162
 - protocol field compression.....121
 - restart timer.....162
- PPP over ATM encapsulation
 - example configuration.....335
- PPP properties, assigning
 - IPCP.....179
- ppp-options
 - lcp-max-conf-req.....1088
 - ncp-max-conf-req.....1145
- ppp-options statement.....1198
 - usage guidelines.....112, 114
- PPPD processes, trace operations.....119
- PPPoE
 - example configuration.....795, 796
- PPPoE client
 - example configuration.....795
- PPPoE server
 - example configuration.....796
- pppoe-options statement.....1197
 - usage guidelines.....790
- preempt statement.....1199
- preferred interface addresses.....192, 194
- preferred statement.....1200
 - usage guidelines.....175, 194
- preferred-source-address statement.....1201
 - example.....190
 - usage guidelines.....187
- premium statement.....1202
 - usage guidelines.....757, 759
- preserve interfaces statement
 - usage guidelines.....869
- preserve-interface statement.....1204
- primary interface addresses.....192, 193
- primary router addresses.....192
- primary router interfaces.....193
- primary statement
 - address for interface.....1205
 - usage guidelines.....193
 - interface for router
 - usage guidelines.....193
- priority (IEEE 802.1ag OAM) statement.....1207
- priority statement.....1206
 - usage guidelines
 - ATM scheduler map.....341
- priority-cost statement.....1208
- promiscuous-mode statement.....1208
 - usage guidelines.....296
- protect circuit, APS
 - configuring.....861
 - load sharing between paired groups.....870
 - paired groups, switching between.....866
- protect-circuit statement.....1209
 - usage guidelines.....861
- protection-group
 - configuration statements.....27
- protection-group statement.....1210
- protocol families
 - configuration statements.....169
 - logical interfaces.....169, 172
 - protocol MTUs.....191
 - unnumbered interfaces.....185
- protocol field compression.....121
- protocol MTUs.....98
 - logical interfaces.....191
 - See also* MTU
- protocol redirect messages.....192
- protocol-down statement.....1210
- protocols
 - families.....169
 - line.....265
- protocols connections
 - configuration statements.....24, 27
- protocols dot1x
 - configuration statements.....25
- protocols Ethernet Ring Protection
 - configuration statements.....27
- protocols OAM
 - configuration statements.....25
- protocols statement.....1211
- protocols VRRP
 - configuration statements.....25, 27
- proxy statement.....1211
 - usage guidelines.....666
- proxy-arp statement.....1212
 - usage guidelines.....671
- push statement
 - Gigabit Ethernet IQ interfaces.....1212
 - Gigabit Ethernet, IQ, IQ2 and IQ2-E interfaces
 - usage guidelines.....641
 - usage guidelines.....648
- push-push statement
 - Gigabit Ethernet IQ interfaces.....1213
 - usage guidelines.....651
- PVC encapsulation *See* ATM PVC encapsulation

Q

queue-depth statement.....	1213
queue-length statement.....	1214
queues statement.....	1214
quiet-period statement.....	1215

R

ranges statement	
stacked VLAN.....	1216
VLAN.....	1216
rate statement.....	1217
usage guidelines.....	129, 876
reassemble-packets statement.....	1217
reauthentication statement.....	1218
receive (tracing flag modifier).....	137
receive-bucket statement.....	1218
usage guidelines.....	129, 876
receive-options-packets statement.....	1219
usage guidelines.....	294, 874
receive-ttl-exceeded statement.....	1219
usage guidelines.....	294, 874
red-differential-delay statement.....	1220
redial-delay statement.....	1220
redirect messages.....	192
redundancy-options statement.....	1221, 1222
remote statement.....	1223
usage guidelines.....	610, 665, 666
remote-loopback statement.....	1223
remote-loopback-respond statement.....	1224
remote-mep statement.....	1225
request statement.....	1225
required-depth statement.....	1226
usage guidelines.....	874
retries statement.....	1227
revert-time statement.....	1228
usage guidelines.....	867
revertive mode, APS.....	867
revertive statement.....	1227
rewrite VLAN tag on untagged frame	
usage guidelines.....	652
rewrite-on-egress statement	
usage guidelines.....	641
rewrite-on-ingress statement	
usage guidelines.....	641
RFC 2615 support, SONET interfaces.....	855
rfc-2615 statement.....	1228
usage guidelines.....	855
ring-protection	
, ethernet-ring.....	1013
ring-protection statement	
, west-interface.....	1337
east-interface.....	999
ring-protection-link-end statement.....	1229
ring-protection-link-owner statement.....	1229

rlsq	
physical part of interface name.....	52
rms	
physical part of interface name.....	52
routers	
default addresses.....	192
primary addresses.....	192
primary interfaces.....	193
Routing Instance.....	684
routing matrix	
chassis, interface naming.....	61
interface naming.....	57, 59
routing-instance.....	684
routing-instance statement.....	1230
rpf-check statement.....	1231
usage guidelines.....	208
rsp	
physical part of interface name.....	52
rtp statement.....	1232
rts statement.....	1232
usage guidelines.....	271
rts-polarity statement.....	1233
usage guidelines.....	274
rtvbr statement.....	1234

S

s1 SONET header byte.....	849
sampling on aggregated links.....	884
example configuration.....	884
sampling statement.....	1235
satop-options	
payload-size.....	1180
satop-options statement.....	1236
SCC.....	57
scheduler-maps statement.....	1237
schedulers statement.....	1237
SDH (VC-4) interfaces	
example configuration.....	436, 442
SDH (VC-4-4C) interface	
example configuration.....	435, 441
SDH interfaces.....	819, 843
configuration statements.....	844
framing.....	846
MTU sizes.....	98, 99, 100
overview.....	819, 843
physical properties.....	844
<i>See also</i> SONET interfaces	
se	
physical part of interface name.....	52
secondary statement.....	1238
send (tracing flag modifier).....	137
send-critical-event statement.....	1238
separators, in interface names.....	56

- serial interfaces
 - clock rates.....270
 - clocking mode.....269, 270
 - configuration statements.....263
 - control leads.....271
 - correcting phase shift.....270
 - default settings.....265
 - DTR circuit.....274
 - idle cycle flag.....271
 - invalid statements.....267
 - line encoding.....277
 - line protocol.....265
 - loopback capability.....275
 - physical interface properties.....263
 - signal polarity.....274
- serial-options statement.....1239
 - usage guidelines.....263
- server statement.....1240
- server-timeout statement.....1240
- service statement.....1241
- service-domain statement.....1241
- service-filter statement.....1242
- service-name statement.....1242
 - usage guidelines.....790
- service-set statement.....1243
- services interfaces.....36
- services statement.....1243
- services-options statement.....1244
- SFC.....59
- shaping statement.....1245
- SHDSL.....361
- shdsl-options statement.....1246
 - usage guidelines.....361
- short-name-format statement.....1246
- short-sequence statement.....1247
- signal polarity, serial interfaces.....274
- snext statement.....1247
 - usage guidelines.....361
- snr-margin statement.....1248
 - usage guidelines.....361
- so
 - physical part of interface name.....52
- SONET interfaces.....819, 843
 - APS.....819, 843
 - clock sources.....128, 875
 - configuration statements.....844
 - damping interface transitions.....877
 - defect triggers.....855
 - configuration example.....858
 - hold timers.....856
 - encapsulation.....106, 878
 - example configuration.....881
 - frame checksum.....851
 - framing.....846
 - HDLC payload scrambling.....854
 - header byte values.....849
 - incrementing STM ID.....850
 - leaky bucket properties.....876
 - link aggregation.....882
 - loopback capability.....852
 - MTU sizes.....98, 99, 100, 873
 - overview.....819, 843
 - path trace identifier.....853
 - physical interface properties.....844
 - RFC 2615 support.....855
 - See also* SDH interfaces
- SONET parameters
 - on channelized OC12 interfaces.....445
- sonet-options statement.....1249
 - usage guidelines.....445, 785, 844
- SONET/SDH interfaces
 - interface speed.....124, 847
 - passive monitoring.....874
- source address filtering
 - Fast Ethernet interfaces.....591, 592
 - Gigabit Ethernet interfaces.....591, 592
- source statement.....1251
- source-address-filter statement.....1252
 - usage guidelines.....592
- source-class usage
 - example configuration.....217
- source-class-usage statement.....1253
 - usage guidelines.....214
- source-filtering statement.....1254
 - usage guidelines.....591, 761
- sp
 - physical part of interface name.....52
- Specifying the Interface Over Which VPN Traffic Travels to the CE Router.....614
- Specifying the Interface to Handle Traffic for a CCC.....615
- Specifying the Interface to Handle Traffic for a CCC Connected to the Layer 2 Circuit.....616
- speed
 - MX Series DPC.....1256
- speed statement.....1255, 1257
 - usage guidelines.....122, 123, 124, 597, 847
- spid1 statement.....1257
- spid2 statement.....1258
- stacked VLAN-tag framing
 - VLAN ID list.....1328
- stacked-vlan-ranges statement.....1258
- stacked-vlan-tagging statement.....1259
 - usage guidelines.....641
- start-end-flag statement.....1260
- startup-silent-period statement.....754
- state (tracing flag).....596
- statement
 - usage guidelines.....790
- static ARP table entries
 - Ethernet interfaces.....669, 671
 - example configuration.....670

static routes	
unnumbered Ethernet interfaces.....	188
static-tei-val statement.....	1261
stm1	
physical part of interface name.....	52
STM1 interfaces	
example configuration.....	785
stm16	
physical part of interface name.....	52
stm4	
physical part of interface name.....	52
supplicant statement	
single.....	1261
supplicant-timeout statement.....	1262
support, technical <i>See</i> technical support	
swap statement	
Gigabit Ethernet IQ interfaces.....	1262
Gigabit Ethernet, IQ, IQ2 and IQ2-E interfaces	
usage guidelines.....	641
swap-push statement	
Gigabit Ethernet IQ interfaces.....	1263
usage guidelines.....	655
swap-swap statement	
Gigabit Ethernet IQ interfaces.....	1263
usage guidelines.....	655
switch-card chassis SCC.....	57
switch-fabric chassis SFC.....	59
switch-options statement.....	1264
switch-port statement	
access switching.....	1265
switch-type statement.....	1266
switching-mode statement.....	1264
usage guidelines.....	867
symbol-period statement.....	1267
Symmetrical Load Balancing	
on 802.3ad Link Aggregation on MX Series.....	636
Synchronous Digital Hierarchy <i>See</i> SDH interfaces	
Synchronous Optical Network <i>See</i> SONET interfaces	
syntax conventions.....	lxv
syslog statement.....	1268
usage guidelines.....	138
system processes	
configuration statements.....	27
system-priority statement	
LACP	
interface.....	1270
T	
t1	
physical part of interface name.....	52
T1 interfaces	
byte encoding.....	938
configuration statements.....	559
MTU sizes.....	98, 99, 100
overview.....	559
time slots	
example configuration.....	567
T1 IQ interfaces	
example configuration.....	480
t1-options statement.....	1271
usage guidelines.....	559
t1-time statement.....	1272
usage guidelines.....	180
T1600 routers	
configured in a routing matrix.....	59
t2-time statement.....	1272
usage guidelines.....	180
t3	
physical part of interface name.....	52
T3 interfaces	
C-bit parity mode.....	943
configuration statements.....	570
encapsulation.....	106
MTU sizes.....	98, 99, 100
overview.....	569
t3-options statement.....	1275
usage guidelines.....	570
t310 statement.....	1273
usage guidelines.....	821
t391 statement.....	1273
t392 statement.....	1274
T640 routers.....	57
tag-protocol-id statement.....	1276
usage guidelines.....	641, 645
tap	
physical part of interface name.....	52
TCC.....	223
encapsulation.....	110
technical support	
contacting JTAC.....	lxvii
tei-option statement.....	1277
then statement	
hierarchical policer.....	1278
threshold statement.....	1278
usage guidelines.....	129, 876
time slots	
channelized E1 IQ and IQE interfaces.....	467
channelized E1 IQE interfaces.....	502
channelized NxDS0 IQ interfaces.....	468, 502
E1 interfaces	
example configuration.....	549
fractional E1 IQ and IQE interfaces.....	467
fractional E1 IQ interfaces.....	502
fractional T1 (COC48/STM16 IQE)	
interfaces.....	412
fractional T1 interfaces.....	433
T1 interfaces	
example configuration.....	567
timer (tracing flag)	
Ethernet interface speed.....	596

- timers
 - APS.....868
- timeslots statement.....1279
 - channelized E1 IQ and IQE interfaces
 - usage guidelines.....467
 - channelized E1 IQ interfaces
 - usage guidelines.....502
 - channelized E1 IQE interfaces
 - usage guidelines.....502
 - channelized NxDS0 IQ interfaces
 - usage guidelines.....468, 502
 - fractional E1 IQ and IQE interfaces
 - usage guidelines.....467
 - fractional E1 IQ interfaces
 - usage guidelines.....502
 - fractional E1 IQE interfaces
 - usage guidelines.....502
 - fractional T1 (COC48/STM16 IQE) interfaces
 - usage guidelines.....412
 - fractional T1 interfaces
 - usage guidelines.....433
 - NxDS0 (COC48/STM16 IQE) interfaces
 - usage guidelines.....413
 - NxDS0 interfaces
 - usage guidelines.....431
- tm statement.....1280
 - usage guidelines.....271
- tm-polarity statement.....1280
 - usage guidelines.....274
- trace operations
 - VRRP.....596
- traceoptions (LACP) statement
 - usage guidelines.....631
- traceoptions statement
 - interface processes.....1283
 - interfaces.....1282
 - usage guidelines.....137
 - LACP.....1285
 - PPPD.....1287
 - PPPD processes
 - usage guidelines.....119
 - VRRP
 - usage guidelines.....596
- tracing flags
 - all.....137
 - event.....137
 - ipc.....137
 - media.....137
- tracing operations
 - LACP.....631
 - PPPD.....119
- track statement
 - DLSw.....1290
 - usage guidelines.....181
- transient interfaces.....32, 35
- transitions
 - damping.....138, 877
- transitions, damping.....138, 877
- translate-discard-eligible statement.....1291
- translate-fecn-and-becn statement.....1291
- translational cross-connect *See* TCC
- translational cross-connect encapsulation *See* TCC,
 - encapsulation
- transmit clock sources.....128, 875
- transmit-bucket statement.....1292
 - usage guidelines.....129, 876
- transmit-clock statement.....1292
 - usage guidelines.....270
- transmit-period statement.....1293
- transmit-weight statement.....1294
 - usage guidelines.....341
- traps statement.....1295
 - usage guidelines.....139, 159
- trej-time statement.....1296
 - usage guidelines.....180
- Tri-Rate Ethernet copper interfaces
 - speed.....597
- Tri-Rate Ethernet interfaces
 - interface speed.....124
- trigger statement.....1297
 - 10-Gigabit Ethernet interfaces
 - usage guidelines.....872
 - usage guidelines.....855
- trigger-link-failure statement.....1298
- Trivial Network Protocol family
 - interface addresses.....174
 - on logical interfaces.....172
- trunk interface.....1073
 - interface-mode statement.....620
 - usage guidelines.....621
 - vlan-id-list statement.....620
- trunk port.....620, 1073
 - VLAN ID list.....1319
- trunk-bandwidth statement.....1298
- trunk-id statement.....1299
- ttl statement.....1299
- tunnel statement.....1300
- TX Matrix Plus router
 - interface naming.....59
 - chassis.....61
- TX Matrix router
 - interface naming.....57
 - chassis.....61
- U**
- umd
 - physical part of interface name.....52
- underlying interfaces
 - specifying family type for.....253
- underlying-interface statement.....1301

unframed statement.....	1302
unicast RPF.....	208
example configuration.....	213
fail filters.....	212, 213
loose mode.....	210
routing asymmetry.....	212
strict mode.....	209
VPNs.....	212
example configuration.....	213
unidirectional link mode.....	139
unidirectional mode, APS.....	867
unidirectional statement.....	1302
unit numbers.....	155
unit statement	
logical interfaces.....	1303
usage guidelines.....	143, 155
unnumbered interfaces	
demux.....	186
Ethernet.....	186
preferred source address.....	1201
IPCP.....	179
point-to-point.....	185
unnumbered-address statement.....	1309
preferred source address	
usage guidelines.....	187
usage guidelines.....	179, 186
unnumbered-interface statement	
usage guidelines.....	179, 790
up-count statement.....	1311
uPIM Ethernet interfaces.....	589
USB modem	
configuring.....	93

V

vbr statement.....	1312
vc-cos-mode statement.....	1313
vc4	
physical part of interface name.....	52
vci statement.....	1314
vci-range statement.....	1315
VCIs.....	281
VCS.....	297
virtual circuits <i>See</i> ATM VCs	
virtual paths <i>See</i> ATM VPs	
Virtual Router Redundancy Protocol <i>See</i> VRRP	
virtual-switch.....	1315
VLAN	
Configuring a Layer 2 Circuit on a VLAN-Bundled Logical Interface.....	615
Configuring a VLAN-Bundled Logical Interface.....	616
Configuring Logical Link-Layer Encapsulation to Support CCCs.....	612
Configuring VLAN ID List-Bundled Logical Interfaces That Connect CCCs.....	612

Specifying the Interface Over Which VPN Traffic Travels to the CE Router.....	614
Specifying the Interface to Handle Traffic for a CCC.....	615
VLAN CCC	
configuration guidelines.....	609
example configuration.....	610
VLAN CCC encapsulation	
encapsulation overhead.....	105
physical interfaces.....	106
VLAN IDs.....	1322
values, listed by Ethernet interface type.....	600
VLAN tag stacking and rewriting	
Gigabit Ethernet, IQ, IQ2 and IQ2-E interfaces.....	641
VLAN tagging.....	599, 665, 1324
VLAN VPLS	
configuration guidelines.....	609
example configuration.....	610
vlan-id statement	
802.1Q VLANs.....	1316
ATM-to-Ethernet cross-connect.....	1318
Ethernet interfaces.....	1316
usage guidelines.....	599, 665
Fast Ethernet interfaces	
usage guidelines.....	599, 665
Gigabit Ethernet interfaces	
usage guidelines.....	604
interface in bridge domain.....	1317
rewriting at ingress or egress.....	1317
usage guidelines.....	641
vlan-id-list statement	
bridge domain.....	1319
Ethernet VLAN circuit.....	1320
Gigabit Ethernet interfaces	
usage guidelines.....	604
usage guidelines.....	620
vlan-id-range statement.....	1322
Ethernet interfaces.....	1322
Gigabit Ethernet interfaces	
usage guidelines.....	604
vlan-ranges statement.....	1323
vlan-rewrite statement.....	1324
vlan-tagging statement.....	1324
Ethernet interfaces	
usage guidelines.....	599, 665
Fast Ethernet interfaces	
usage guidelines.....	599, 665
Gigabit Ethernet interfaces	
usage guidelines.....	599
vlan-tags statement	
dual-tag framing.....	1326
stacked VLAN tags.....	1328
usage guidelines.....	645
vlan-tags-outer statement.....	1329

vlan-vci-ccc encapsulation
 ATM-to-Ethernet interworking.....229
 vlan-vci-tagging statement.....1330
 ATM-to-Ethernet interworking.....229
 VLANs
 configuring VLAN ranges.....1258, 1323
 voice over IP
 dynamic CAC.....167
 vpi statement.....1331, 1333
 VPIs.....281
 vpls protocol family
 interface addresses.....174
 on logical interfaces.....172
 VPNs
 unicast RPF.....212
 VPs.....297
 VRRP.....753
 configuration statements.....25, 27
 trace operations.....596
 tracing flag.....596
 vrrp
 failover-delay.....1018
 vsp
 physical part of interface name.....52
 vt
 physical part of interface name.....52
 VT mapping
 channelized T1 (COC48/STM16 IQE)
 interfaces.....411
 channelized T1 interfaces.....431, 439
 channelized T1 IQ interfaces.....444, 460
 channelized T1 IQE interfaces.....419
 vtmapping statement.....1333
 usage guidelines.....472

W

WAN PHY.....872
 configuring.....781
 watch-list statement.....1334
 usage guidelines.....835
 wavelength statement.....1335
 usage guidelines.....779
 wavelength-division multiplexing *See* WDM
 WDM.....819, 843
 weighted random early detection.....598
 west-interface statement.....1337
 working circuit, APS
 configuring.....861
 load sharing between paired groups.....870
 paired groups, switching between.....866
 working-circuit statement.....1337
 usage guidelines.....861
 WRED.....598

X

xe
 physical part of interface name.....52
 xt
 physical part of interface name.....52

Y

yellow-differential-delay statement.....1338

Z

z0-increment statement.....1338
 usage guidelines.....850
 z3 SONET header byte.....849
 z4 SONET header byte.....849

Index of Statements and Commands

Symbols

802.3ad statement.....889

A

accept statement.....890
 accept-source-mac statement.....891
 access-concentrator statement.....892
 access-profile statement.....893
 accounting statement.....894
 accounting-profile statement.....894
 acfc statement.....895
 ack-delay-time statement.....895
 ack-max statement.....896
 acknowledge-retries statement.....896
 acknowledge-timer statement.....897
 action statement.....898
 action-profile statement
 applying to remote MEP.....899
 CFM.....899
 LFM.....900
 action-red-differential-delay statement.....901
 activation-delay statement.....901
 activation-priority statement.....902
 address statement.....903
 advertise-interval statement.....905
 age statement.....906
 aggregate statement.....907
 hierarchical policer.....908
 aggregate-ports statement.....909
 aggregated-ether-options statement.....910
 aggregated-sonet-options statement.....911
 allow-any-vci statement.....911
 allow-fragmentation statement.....912
 allow-remote-loopback statement.....912
 annex statement.....913
 aps statement.....914
 annex-b.....913
 arp statement.....915
 asynchronous-notification statement.....916
 atm-encapsulation statement.....916
 atm-options statement.....917
 atm-scheduler-map statement.....918
 authentication-key statement.....918

authentication-profile-name statement.....919
 authenticator statement.....920
 auto-configure statement.....921
 auto-discovery statement.....921
 auto-negotiation statement
 Gigabit Ethernet.....922
 J Series uPIM.....589, 923
 auto-reconnect statement.....923

B

backup-destination statement.....924
 backup-interface statement.....924
 backup-options statement.....925
 bandwidth statement.....925
 bandwidth-limit statement
 hierarchical policer.....926
 policer for Gigabit Ethernet interface.....927
 bchannel-allocation statement.....927
 bearer-bandwidth-limit statement.....928
 bert-algorithm statement.....929
 bert-error-rate statement.....931
 bert-period statement.....932
 bridge-domain statement.....933
 broadcast statement.....933
 buildout statement.....934
 bundle statement.....936
 burst-size-limit statement
 hierarchical policer.....937
 policer for Gigabit Ethernet interface.....937
 byte-encoding statement.....938
 bytes statement.....939

C

callback statement.....940
 callback-wait-period statement.....941
 caller statement.....942
 calling-number statement.....943
 cbit-parity statement.....943
 cbr statement.....944
 cell-bundle-size statement.....945
 chap statement.....946
 chap-secret statement.....947
 cisco-interoperability statement.....947
 classifier statement.....948

clear-dont-fragment-bit statement.....	948
client statement.....	949
clock-rate statement.....	952
clocking statement.....	950
clocking-mode statement.....	951
compatibility-mode statement.....	953
compression statement.....	954
compression-device statement.....	955
connections statement.....	956
connectivity-fault-management statement.....	957
container-devices statement.....	958
device-count.....	958
container-list statement.....	958
container-options statement.....	959
container-type statement.....	959
continuity-check statement.....	960
control-polarity statement.....	961
control-signal statement.....	962
copy-tos-to-outer-ip-header statement.....	962
core-dump statement.....	963
crc-major-alarm-threshold statement.....	963
crc-minor-alarm-threshold statement.....	964
cts statement.....	964
cts-polarity statement.....	965
current statement.....	965

D

data-input statement.....	966
dcd statement.....	967
dcd-polarity statement.....	967
dce statement.....	968
dce-options statement.....	968
deactivation-delay statement.....	969
default-action statement.....	969
default-chap-secret statement.....	970
default-pap-password statement.....	970
demux-destination statement.....	972
demux interfaces.....	973
demux-options statement.....	973
demux-source statement.....	974
demux interfaces.....	975
demux0 statement.....	971
description statement.....	976
destination statement.....	977
destination-class-usage statement.....	979
destination-profile statement.....	980
dial-options statement.....	982
dial-string statement.....	983
dialer statement.....	980
dialer-options statement.....	981
dialin statement.....	982
direction statement.....	983
disable statement.....	984
link protection.....	984
disable-mlppp-inner-ppp-pfc statement.....	985

dlci statement.....	985
dlsw statement.....	986
do-not-fragment statement.....	986
dot1x.....	987
authenticator.....	987
authentication-profile-name.....	987
interface.....	987
<i>See also</i> maximum-requests	
<i>See also</i> quiet-period	
<i>See also</i> reauthentication	
<i>See also</i> retries	
<i>See also</i> supplicant	
<i>See also</i> supplicant-timeout	
<i>See also</i> transmit-period	
interface.....	987
.....	987
<i>See also</i> server-timeout	
<i>See also</i> server-timeout	
down-count statement.....	988
drop-timeout statement.....	989
ds0-options statement.....	989
dsl-options statement.....	990
dsr statement.....	990
dsr-polarity statement.....	991
dte-options statement.....	991
dtr statement.....	992
dtr-circuit statement.....	993
dtr-polarity statement.....	993
dump-on-flow-control statement.....	994
dynamic-call-admission-control statement.....	994
dynamic-profile statement.....	
stacked VLAN ranges.....	995
VLAN ranges.....	996

E

e1-options statement.....	997
e3-options statement.....	998
east-interface statement.....	999
encapsulation statement.....	1000
encoding statement.....	1007
epd-threshold statement.....	1008
es-options statement.....	1009
ethernet statement.....	1010
ethernet-policer-profile statement.....	1012
ethernet-ring statement.....	1013
ethernet-switch-profile statement.....	762, 1014
eui-64 statement.....	1017
evcs statement.....	1015
event statement.....	1016
interface-status-tlv statement.....	1075
port-status-tlv statement.....	1196
event-thresholds statement.....	1016

F

f-max-period statement.....	1028
facility-override statement.....	1017
failover-delay statement.....	754, 1018
family statement.....	1019
fastether-options statement.....	1023
fcs statement.....	1024
feac-loop-respond statement.....	1025
filter statement.....	1026
firewall	
hierarchical-policer.....	1043
flexible-vlan-tagging statement.....	1027
flow-control statement.....	1028
force statement.....	1029
forwarding-class statement.....	1030
fragment-threshold statement (interfaces).....	1031
frame-error statement.....	1032
frame-period statement.....	1033
frame-period-summary statement.....	1034
framing statement.....	1035

G

gether-options statement.....	1038
gratuitous-arp-reply statement.....	1039

H

hardware-assisted-timestamping statement.....	1040
hello-timer statement.....	1041
hierarchical-policer.....	1043
high-plp-max-threshold statement.....	1041
high-plp-threshold statement.....	1042
hold-interval statement	
connectivity-fault management.....	1044
Ethernet ring protection switching.....	1044
hold-time statement.....	1045
host statement.....	1049

I

idle-cycle-flag statement.....	1050
idle-time statement.....	1051
idle-timeout statement.....	1051
ieee802.1p statement.....	1052
if-exceeding statement	
hierarchical policer.....	1052
ignore statement.....	1053
ignore-all statement.....	1053
ignore-l3-incompletes statement.....	1054
ilmi statement.....	1054
inactivity-timeout statement.....	1055
incoming-called-number statement.....	1055
incoming-map statement.....	1056
indication statement.....	1057
indication-polarity statement.....	1057

ingress-rate-limit statement.....	1058
init-command-string statement.....	1059
initial-route-check statement.....	1060
inner-tag-protocol-id statement.....	1060
inner-vlan-id statement.....	1061
inner-vlan-id-range statement.....	1062
input statement.....	1062
input-list statement.....	1063
input-policer statement.....	1063
input-priority-map statement.....	1064
input-three-color statement.....	1064
input-vlan-map statement.....	1065
instance statement.....	1066
interface.....	1069
maximum-requests.....	1069
quiet-period.....	1069
reauthentication.....	1069
retries.....	1069
server-timeout.....	1069
supplicant.....	1069
supplicant-timeout.....	1069
transmit-period.....	1069
interface statement.....	1067
CoS.....	1068
DLSw.....	1067
IEEE 802.1ag.....	1070
IEEE 802.3ah.....	1071
interface-down statement.....	969, 1072
interface-mode statement.....	1073
interface-set statement.....	1074
interface-status-tlv statement.....	1075
interface-switch statement.....	1075
interface-type statement.....	1076
interfaces statement.....	1072
interleave-fragments statement.....	1077
interval statement.....	1078
inverse-arp statement.....	1079
invert-data statement.....	1080
ipsec-sa statement.....	1080
isdn-options statement.....	1081

K

keep-address-and-control statement.....	1082
keepalives statement.....	1083
key statement.....	1084

L

l2tp-interface-id statement.....	1084
lacp statement.....	1086
802.3ad.....	1085
layer2-policer statement.....	1087
lcp-max-conf-req.....	1088
lcp-restart-timer statement.....	1088
level statement.....	1089

line-encoding statement.....	1091
line-protocol statement.....	1091
line-rate statement.....	1092
linear-red-profile statement.....	1089
linear-red-profiles statement.....	1090
link-adjacency-loss statement.....	1092
link-discovery statement.....	1093
link-down statement.....	1093
link-event-rate statement.....	1094
link-fault-management statement.....	1095
link-layer-overhead statement.....	1096
link-mode statement.....	1097
link-protection statement.....	1098
link-speed statement.....	1099
linktrace statement.....	1100
llc2 statement.....	1101
lmi statement	
Ethernet OAM.....	1104
Frame Relay keepalives.....	1103
lmi-type statement.....	1105
load-interval statement.....	1105
load-threshold statement.....	1106
local-mac statement.....	1106
local-name statement.....	1107
local-password statement.....	1108
local-window statement.....	1108
lockout statement.....	1109
log-prefix statement.....	1110
logical-interface-policer statement.....	1109
logical-systems statement.....	1110
long-buildout statement.....	1111
loop-timing statement.....	1116
loopback statement.....	1112
loopback-clear-timer statement.....	1115
loss-priority statement.....	1116
loss-threshold statement.....	1117
low-plp-max-threshold statement.....	1117
low-plp-threshold statement.....	1118
lsq-failure-options statement.....	1118

M

mac statement.....	1119
mac-address statement.....	1119
mac-learn-enable statement.....	1120
mac-validate statement.....	1121
maintenance-association statement.....	1122
maintenance-domain	
mip-half-function.....	1131
virtual-switch.....	1315
maintenance-domain statement.....	1123
map statement.....	1124
master-only statement.....	1125
max-retry statement.....	1127
maximum-contexts statement.....	1125
maximum-requests statement.....	1126

maximum-vcs statement.....	1126
member-interface-speed statement.....	1127
member-interface-type statement.....	1128
member-interface-speed.....	1128
mep statement.....	1129
minimum-links statement.....	1130
mip-half-function.....	684, 1131
mlfr-uni-nni-bundle-options statement.....	1132
mode statement.....	1133
modem-options statement.....	1133
monitor-session statement.....	1134
mpls statement	
ATM interfaces.....	1134
mrru statement.....	1135
mtu statement.....	1136
multicast-dlci statement.....	1137
multicast-only statement.....	1138
multicast-vci statement.....	1138
multilink-max-classes statement.....	1139
multipoint statement.....	1139
multipoint-destination statement.....	1140
multiservice-options	
dump-on-flow-control.....	994
multiservice-options statement.....	1141

N

n391 statement.....	1141
n392 statement.....	1142
n393 statement.....	1142
name-format statement.....	1143
native-vlan-id statement.....	1144
ncp-max-conf-req.....	1145
ncp-restart-timer statement.....	1145
negotiate-address statement.....	1146
negotiation-options statement.....	1146
neighbor statement.....	1147
no-allow-link-events statement.....	1147
no-asynchronous-notification statement.....	916
no-auto-negotiation statement	
Gigabit Ethernet.....	922
J Series uPIM.....	589, 923
no-cbit-parity statement.....	943
no-core-dump statement.....	963
no-feac-loop-respond statement.....	1025
no-flow-control statement.....	1028
no-gratuitous-arp-reply statement.....	1039
no-gratuitous-arp-request statement.....	1149
no-keepalives statement.....	1149
no-long-buildout statement.....	1111
no-loop-timing statement.....	1116
no-loopback statement.....	1112, 1114
no-mac-learn-enable statement.....	1120
no-partition statement.....	1151
no-payload-scrambler statement.....	1179
no-preempt statement.....	1199

no-redirects statement.....	1154
no-source-filtering statement.....	1254
no-syslog statement.....	1269
no-termination-request statement.....	1155
no-translate-discard-eligible statement.....	1291
no-translate-fecn-and-becn statement.....	1291
no-unframed statement.....	1302
no-z0-increment statement.....	1338
non-revertive statement.....	1150

O

oam statement.....	1156
oam-liveness statement.....	1158
oam-period statement.....	1159
oc-slice statement.....	1160
open-timeout statement.....	1160
operating-mode statement.....	1161
optics-options statement.....	1162
otn-options.....	773
otn-options statement.....	1163
output statement.....	1165
output-list statement.....	1165
output-policer statement.....	1166
output-priority-map statement.....	1166
output-three-color statement.....	1167
output-vlan-map statement.....	1168
overflow statement.....	1170

P

p-bit-timeout statement.....	1180
paired-group statement.....	1171
pap statement.....	1172
pap-password statement.....	1173
partition statement.....	1174
passive statement.....	1175
passive-monitor-mode statement.....	1176
path-database-size statement.....	1177
path-trace statement.....	1178
payload-scrambler statement.....	1179
payload-size statement.....	1180
pdu-interval statement.....	1181
pdu-threshold statement.....	1181
peer statement.....	1182
peer-unit statement.....	1182
per-unit-scheduler statement.....	1184
performance-monitoring statement.....	1183
periodic statement.....	1183
pfc statement.....	1185
pic-type statement.....	1185
plp-to-clp statement.....	1187
plp1 statement.....	1186
point-to-point statement.....	1187

policer statement	
CoS.....	1188
interface MAC.....	1190
pool statement.....	1191
pop statement	
Gigabit Ethernet IQ interfaces.....	1192
pop-all-labels statement.....	1193
pop-pop statement	
Gigabit Ethernet IQ interfaces.....	1194
pop-swap statement	
Gigabit Ethernet IQ interfaces.....	1194
port statement	
voice services.....	1195
port-priority statement	
LACP.....	1195
port-status-tlv statement.....	1196
post-service-filter statement.....	1196
ppp-options	
lcp-max-conf-req.....	1088
ncp-max-conf-req.....	1145
ppp-options statement.....	1198
pppoe-options statement.....	1197
preempt statement.....	1199
preferred statement.....	1200
preferred-source-address statement.....	1201
premium statement.....	1202
preserve-interface statement.....	1204
primary statement	
address for interface.....	1205
priority (IEEE 802.1ag OAM) statement.....	1207
priority statement.....	1206
priority-cost statement.....	1208
promiscuous-mode statement.....	1208
protect-circuit statement.....	1209
protection-group statement.....	1210
protocol-down statement.....	1210
protocols statement.....	1211
proxy statement.....	1211
proxy-arp statement.....	1212
push statement	
Gigabit Ethernet IQ interfaces.....	1212
push-push statement	
Gigabit Ethernet IQ interfaces.....	1213

Q

queue-depth statement.....	1213
queue-length statement.....	1214
queues statement.....	1214
quiet-period statement.....	1215

R

ranges statement	
stacked VLAN.....	1216
VLAN.....	1216

rate statement.....	1217
reassemble-packets statement.....	1217
reauthentication statement.....	1218
receive-bucket statement.....	1218
receive-options-packets statement.....	1219
receive-ttl-exceeded statement.....	1219
red-differential-delay statement.....	1220
redial-delay statement.....	1220
redundancy-options statement.....	1221, 1222
remote statement.....	1223
remote-loopback statement.....	1223
remote-loopback-respond statement.....	1224
remote-mep statement.....	1225
request statement.....	1225
required-depth statement.....	1226
retries statement.....	1227
revert-time statement.....	1228
revertive statement.....	1227
rfc-2615 statement.....	1228
ring-protection	
, ethernet-ring.....	1013
ring-protection statement	
, west-interface.....	1337
east-interface.....	999
ring-protection-link-end statement.....	1229
ring-protection-link-owner statement.....	1229
routing-instance statement.....	1230
rpf-check statement.....	1231
rtp statement.....	1232
rts statement.....	1232
rts-polarity statement.....	1233
rtvbr statement.....	1234

S

sampling statement.....	1235
satop-options statement.....	1236
scheduler-maps statement.....	1237
schedulers statement.....	1237
secondary statement.....	1238
send-critical-event statement.....	1238
serial-options statement.....	1239
server statement.....	1240
server-timeout statement.....	1240
service statement.....	1241
service-domain statement.....	1241
service-filter statement.....	1242
service-name statement.....	1242
service-set statement.....	1243
services statement.....	1243
services-options statement.....	1244
shaping statement.....	1245
shdsl-options statement.....	1246
short-name-format statement.....	1246
short-sequence statement.....	1247
snext statement.....	1247

snr-margin statement.....	1248
sonet-options statement.....	1249
source statement.....	1251
source-address-filter statement.....	1252
source-class-usage statement.....	1253
source-filtering statement.....	1254
speed	
MX Series DPC.....	1256
speed statement.....	1255, 1257
spid1 statement.....	1257
spid2 statement.....	1258
stacked-vlan-ranges statement.....	1258
stacked-vlan-tagging statement.....	1259
start-end-flag statement.....	1260
startup-silent-period statement.....	754
static-tei-val statement.....	1261
supplicant statement	
single.....	1261
supplicant-timeout statement.....	1262
swap statement	
Gigabit Ethernet IQ interfaces.....	1262
swap-push statement	
Gigabit Ethernet IQ interfaces.....	1263
swap-swap statement	
Gigabit Ethernet IQ interfaces.....	1263
switch-options statement.....	1264
switch-port statement	
access switching.....	1265
switch-type statement.....	1266
switching-mode statement.....	1264
symbol-period statement.....	1267
syslog statement.....	1268
system-priority statement	
LACP	
interface.....	1270

T

t1-options statement.....	1271
t1-time statement.....	1272
t2-time statement.....	1272
t3-options statement.....	1275
t310 statement.....	1273
t391 statement.....	1273
t392 statement.....	1274
tag-protocol-id statement.....	1276
tei-option statement.....	1277
then statement	
hierarchical policer.....	1278
threshold statement.....	1278
timeslots statement.....	1279
tm statement.....	1280
tm-polarity statement.....	1280
traceoptions statement	
interface processes.....	1283
interfaces.....	1282

LACP.....	1285
PPPD.....	1287
track statement	
DLSw.....	1290
translate-discard-eligible statement.....	1291
translate-fecn-and-becn statement.....	1291
transmit-bucket statement.....	1292
transmit-clock statement.....	1292
transmit-period statement.....	1293
transmit-weight statement.....	1294
traps statement.....	1295
trej-time statement.....	1296
trigger statement.....	1297
trigger-link-failure statement.....	1298
trunk-bandwidth statement.....	1298
trunk-id statement.....	1299
ttl statement.....	1299
tunnel statement.....	1300

U

underlying-interface statement.....	1301
unframed statement.....	1302
unidirectional statement.....	1302
unit statement	
logical interfaces.....	1303
unnumbered-address statement.....	1309
up-count statement.....	1311

V

vbr statement.....	1312
vc-cos-mode statement.....	1313
vci statement.....	1314
vci-range statement.....	1315
virtual-switch.....	1315
vlan-id statement	
802.1Q VLANs.....	1316
ATM-to-Ethernet cross-connect.....	1318
Ethernet interfaces.....	1316
interface in bridge domain.....	1317
rewriting at ingress or egress.....	1317
vlan-id-list statement	
bridge domain.....	1319
Ethernet VLAN circuit.....	1320
vlan-id-range statement.....	1322
vlan-ranges statement.....	1323
vlan-rewrite statement.....	1324
vlan-tagging statement.....	1324
vlan-tags statement	
dual-tag framing.....	1326
stacked VLAN tags.....	1328
vlan-tags-outer statement.....	1329
vlan-vci-tagging statement.....	1330
vpi statement.....	1331, 1333

vrrp	
failover-delay.....	1018
vtmapping statement.....	1333

W

watch-list statement.....	1334
wavelength statement.....	1335
west-interface statement.....	1337
working-circuit statement.....	1337

Y

yellow-differential-delay statement.....	1338
--	------

Z

z0-increment statement.....	1338
-----------------------------	------

