



J-series™ Services Router

Basic LAN and WAN Access Configuration Guide

Release 8.4

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, California 94089
USA

408-745-2000

www.juniper.net

Part Number: 530-020999-01, Revision 1

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, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. JUNOS and JUNOSe are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

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

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

J-series™ Services Router Basic LAN and WAN Access Configuration Guide

Release 8.4

Copyright © 2007, Juniper Networks, Inc.

All rights reserved. Printed in USA.

Writing: Nidhi Bhargava, Michael Bushong, Maya Devi, Taffy Everts, Walter Goralski, Joshua Kim, Jerry Isaac, Archana Maheshwari, Hareesh Kumar Kozhippurath Narayana Panicker, Laura Phillips, Frank Reade, Selvakumar T. S., and Alan Twigg

Editing: Taffy Everts and Stella Hackell

Illustration: Faith Bradford Brown and Nathaniel Woodward

Cover Design: Edmonds Design

Revision History

29 June 2007—Revision 1.

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.

SOFTWARE LICENSE

The terms and conditions for using this software are described in the software license contained in the acknowledgment to your purchase order or, to the extent applicable, to any reseller agreement or end-user purchase agreement executed between you and Juniper Networks. By using this software, you indicate that you understand and agree to be bound by those terms and conditions. Generally speaking, the software license restricts the manner in which you are permitted to use the software and may contain prohibitions against certain uses. The software license may state conditions under which the license is automatically terminated. You should consult the license for further details. For complete product documentation, please see the Juniper Networks Web site at www.juniper.net/techpubs.

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 Juniper Networks, Inc. and its subsidiaries (collectively "Juniper"), and 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, and updates and releases of such software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller. "Embedded Software" means Software which Juniper has embedded in the Juniper 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 the 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 software on multiple computers requires multiple licenses, regardless of whether such computers 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 the Embedded Software on non-Juniper equipment; (j) use the 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 for the Software are exclusive of taxes, withholdings, duties, or levies (collectively "Taxes"). Customer shall be responsible for paying Taxes arising from the purchase of the license, or importation or use of the Software.

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	xix
Part 1	Using the Configuration Interfaces	
Chapter 1	Using Services Router Configuration Tools	3
Part 2	Configuring Router Interfaces	
Chapter 2	Interfaces Overview	41
Chapter 3	Configuring Ethernet, DS1, DS3, and Serial Interfaces	103
Chapter 4	Configuring Channelized T1/E1/ISDN PRI Interfaces	139
Chapter 5	Configuring Digital Subscriber Line Interfaces	157
Chapter 6	Configuring Point-to-Point Protocol over Ethernet	191
Chapter 7	Configuring ISDN	211
Chapter 8	Configuring Link Services Interfaces	261
Chapter 9	Configuring VoIP	311
Chapter 10	Configuring uPIMs as Ethernet Switches	343
Part 3	Configuring Routing Protocols	
Chapter 11	Routing Overview	349
Chapter 12	Configuring Static Routes	383
Chapter 13	Configuring a RIP Network	395
Chapter 14	Configuring an OSPF Network	411
Chapter 15	Configuring the IS-IS Protocol	431
Chapter 16	Configuring BGP Sessions	441
Part 4	Index	
	Index	461

Table of Contents

About This Guide xix

Objectives	xix
Audience	xix
How to Use This Guide	xx
Document Conventions	xxi
Related Juniper Networks Documentation	xxii
Documentation Feedback	xxv
Requesting Support	xxv

Part 1

Using the Configuration Interfaces

Chapter 1

Using Services Router Configuration Tools 3

Configuration Tools Terms	3
Configuration Tools Overview	4
Editing and Committing a Configuration	4
J-Web Configuration Options	5
CLI Configuration Commands	5
Filtering Configuration Command Output	7
Before You Begin	7
Using J-Web Quick Configuration	7
Using the J-Web Configuration Editor	9
Viewing the Configuration Text	9
Editing and Committing the Clickable Configuration	10
Editing the Clickable Configuration	10
Discarding Parts of a Candidate Configuration	13
Committing a Clickable Configuration	14
Editing and Committing the Configuration Text	14
Uploading a Configuration File	15
Managing Configuration Files with the J-Web Interface	16
Configuration Database and History Overview	16
Displaying Users Editing the Configuration	18
Comparing Configuration Files	19
Downloading a Configuration File	20
Loading a Previous Configuration File	21
Setting, Viewing, or Deleting the Rescue Configuration	21
Using the CLI Configuration Editor	22
Entering and Exiting Configuration Mode	22
Navigating the Configuration Hierarchy	23

Modifying the Configuration	25
Adding or Modifying a Statement or Identifier	25
Using Search and Replace	26
Deleting a Statement or Identifier	27
Copying a Statement	27
Renaming an Identifier	28
Inserting an Identifier	28
Deactivating a Statement or Identifier	30
Committing a Configuration with the CLI	30
Verifying a Configuration	31
Committing a Configuration and Exiting Configuration Mode	31
Committing a Configuration That Requires Confirmation	31
Scheduling and Canceling a Commit	32
Loading a Previous Configuration File with the CLI	32
Setting or Deleting the Rescue Configuration with the CLI	33
Disabling the CONFIG or RESET CONFIG Button	34
Entering Operational Mode Commands During Configuration	34
Managing Configuration Files with the CLI	35
Loading a New Configuration File	35
Saving a Configuration File	37

Part 2

Configuring Router Interfaces

Chapter 2

Interfaces Overview	41
Interfaces Terms	42
Network Interfaces	46
Media Types	46
Network Interface Naming	47
J-series Interface Naming Conventions	47
Understanding CLI Output for J-series Interfaces	49
Data Link Layer Overview	52
Physical Addressing	52
Network Topology	52
Error Notification	52
Frame Sequencing	52
Flow Control	52
Data Link Sublayers	52
MAC Addressing	53
Ethernet Interface Overview	53
Ethernet Access Control and Transmission	54
Collisions and Detection	54
Collision Detection	54
Backoff Algorithm	55
Collision Domains and LAN Segments	55
Repeaters	55
Bridges and Switches	56

Broadcast Domains	56
Ethernet Frames	56
T1 and E1 Interfaces Overview	57
T1 Overview	57
E1 Overview	58
T1 and E1 Signals	58
Encoding	58
AMI Encoding	59
B8ZS and HDB3 Encoding	59
T1 and E1 Framing	59
Superframe (D4) Framing for T1	59
Extended Superframe (ESF) Framing for T1	60
T1 and E1 Loopback Signals	60
Channelized T1/E1/ISDN PRI Interfaces Overview	61
T3 and E3 Interfaces Overview	61
Multiplexing DS1 Signals	62
DS2 Bit Stuffing	62
DS3 Framing	63
M13 Asynchronous Framing	63
C-Bit Parity Framing	64
Serial Interface Overview	66
Serial Transmissions	67
Signal Polarity	68
Serial Clocking Modes	68
Serial Interface Transmit Clock Inversion	69
DTE Clock Rate Reduction	69
Serial Line Protocols	69
EIA-530	70
RS-232	70
RS-422/449	71
V.35	71
X.21	72
ADSL Interface Overview	72
ADSL Systems	73
ADSL2 and ADSL2 +	74
Asynchronous Transfer Mode	74
SHDSL Interface Overview	74
ISDN Interface Overview	75
ISDN Channels	75
ISDN Interfaces	75
Typical ISDN Network	75
NT Devices and S and T Interfaces	76
U Interface	76
ISDN Call Setup	77
Layer 2 ISDN Connection Initialization	77
Layer 3 ISDN Session Establishment	77
Interface Physical Properties	78
Bit Error Rate Testing	79
Interface Clocking	79
Data Stream Clocking	80
Explicit Clocking Signal Transmission	80

Frame Check Sequences	80
Cyclic Redundancy Checks and Checksums	80
Two-Dimensional Parity	81
MTU Default and Maximum Values	81
Physical Encapsulation on an Interface	83
Frame Relay	83
Virtual Circuits	84
Switched and Permanent Virtual Circuits	84
Data-Link Connection Identifiers	84
Congestion Control and Discard Eligibility	85
Point-to-Point Protocol	85
Link Control Protocol	85
PPP Authentication	86
Network Control Protocols	87
Magic Numbers	87
CSU/DSU Devices	88
Point-to-Point Protocol over Ethernet	88
PPPoE Discovery	88
PPPoE Sessions	89
High-Level Data Link Control	89
HDLC Stations	89
HDLC Operational Modes	90
Interface Logical Properties	90
Protocol Families	91
Common Protocol Suites	91
Other Protocol Suites	91
IPv4 Addressing	91
IPv4 Classful Addressing	92
IPv4 Dotted Decimal Notation	92
IPv4 Subnetting	93
IPv4 Variable-Length Subnet Masks	93
IPv6 Addressing	94
IPv6 Address Representation	94
IPv6 Address Types	95
IPv6 Address Scope	95
IPv6 Address Structure	95
Virtual LANs	96
Special Interfaces	97
Discard Interface	100
Loopback Interface	100
Management Interface	101
Services Interfaces	101
MLPPP and MLFR	101
MLFR Frame Relay Forum	102
CRTP	102

Chapter 3	Configuring Ethernet, DS1, DS3, and Serial Interfaces	103
	Before You Begin	103
	Configuring DS1, DS3, Ethernet, and Serial Interfaces with Quick Configuration	104
	Configuring an E1 Interface with Quick Configuration	105
	Configuring an E3 Interface with Quick Configuration	109
	Configuring a Fast Ethernet Interface with Quick Configuration	112
	Configuring Gigabit Ethernet Interfaces with Quick Configuration	115
	Configuring T1 Interfaces with Quick Configuration	120
	Configuring T3 Interfaces with Quick Configuration	124
	Configuring Serial Interfaces with Quick Configuration	127
	Configuring Network Interfaces with a Configuration Editor	131
	Adding a Network Interface with a Configuration Editor	131
	Configuring Static ARP Entries on Ethernet Interfaces	133
	Deleting a Network Interface with a Configuration Editor	134
	Verifying Interface Configuration	135
	Verifying the Link State of All Interfaces	136
	Verifying Interface Properties	137
 Chapter 4	 Configuring Channelized T1/E1/ISDN PRI Interfaces	 139
	Channelized T1/E1/ISDN PRI Terms	139
	Channelized T1/E1/ISDN PRI Overview	140
	Channelized T1/E1/ISDN PRI Interfaces	140
	Drop and Insert	141
	ISDN PRI Transmission on Channelized Interfaces	141
	Before You Begin	142
	Configuring Channelized T1/E1/ISDN PRI interfaces with a Configuration Editor	142
	Configuring Channelized T1/E1/ISDN PRI Interface as a Clear Channel	142
	Configuring Channelized T1/E1/ISDN PRI Interface to Drop and Insert Time Slots	145
	Configuring Channelized T1/E1/ISDN PRI Interfaces for ISDN PRI Operation	147
	Verifying Channelized T1/E1/ISDN PRI Interfaces	150
	Verifying Channelized Interfaces	151
	Verifying Clear-Channel Interfaces	152
	Verifying ISDN PRI Configuration on Channelized T1/E1/ISDN PRI Interfaces	153
	Frequently Asked Questions About Channelized T1/E1/ISDN PRI Interfaces	153
	What Clock Combinations Are Possible for Channelized T1/E1/ISDN PRI Drop and Insert?	153
 Chapter 5	 Configuring Digital Subscriber Line Interfaces	 157
	DSL Terms	157
	Before You Begin	158

Configuring ATM-over-ADSL Interfaces	159
Configuring an ATM-over-ADSL Interface with Quick Configuration	159
Adding an ATM-over-ADSL Network Interface with a Configuration Editor	164
Configuring ATM-over-SHDSL Interfaces	169
Configuring an ATM-over-SHDSL Interface with Quick Configuration	170
Adding an ATM-over-SHDSL Interface with a Configuration Editor	174
Configuring CHAP on DSL Interfaces (Optional)	179
Verifying DSL Interface Configuration	180
Verifying ADSL Interface Properties	181
Displaying a PPPoA Configuration for an ATM-over-ADSL Interface	185
Verifying an ATM-over-SHDSL Configuration	186

Chapter 6 Configuring Point-to-Point Protocol over Ethernet 191

PPPoE Terms	191
PPPoE Overview	192
PPPoE Interfaces	193
Ethernet Interface	193
ATM-over-ADSL or ATM-over-SHDSL Interface	193
PPPoE Stages	194
PPPoE Discovery Stage	194
PPPoE Session Stage	194
Optional CHAP Authentication	194
Before You Begin	195
Configuring PPPoE Interfaces with Quick Configuration	195
Configuring PPPoE with a Configuration Editor	198
Setting the Appropriate Encapsulation on the Interface (Required)	198
Configuring PPPoE Encapsulation on an Ethernet Interface	199
Configuring PPPoE Encapsulation on an ATM-over-ADSL or ATM-over-SHDSL Interface	200
Configuring PPPoE Interfaces (Required)	201
Configuring CHAP on a PPPoE Interface (Optional)	204
Verifying a PPPoE Configuration	205
Displaying a PPPoE Configuration for an Ethernet Interface	206
Displaying a PPPoE Configuration for an ATM-over-ADSL or ATM-over-SHDSL Interface	207
Verifying PPPoE Interfaces	208
Verifying PPPoE Sessions	209
Verifying the PPPoE Version	209
Verifying PPPoE Statistics	210

Chapter 7 Configuring ISDN 211

ISDN Terms	211
ISDN Overview	214
ISDN Interfaces	214
ISDN BRI Interface Types	214
ISDN PRI Interface Types	215
Dialer Interface	215

Before You Begin	215
Configuring ISDN BRI Interfaces with Quick Configuration	216
Configuring ISDN BRI Physical Interfaces with Quick Configuration	216
Configuring ISDN BRI Dialer Interfaces with Quick Configuration	219
Configuring ISDN Interfaces and Features with a Configuration Editor	223
Adding an ISDN BRI Interface (Required)	224
Configuring Dialer Interfaces (Required)	226
Configuring Dial Backup	228
Configuring Dialer Filters for Dial-on-Demand Routing Backup	229
Configuring the Dialer Filter	230
Applying the Dial-on-Demand Dialer Filter to the Dialer Interface	231
Configuring Dialer Watch	231
Adding a Dialer Watch Interface on the Services Router	232
Configuring the ISDN Interface for Dialer Watch	232
Configuring Dial-on-Demand Routing Backup with OSPF Support (Optional)	233
Configuring Bandwidth on Demand (Optional)	234
Configuring Dialer Interfaces for Bandwidth on Demand	234
Configuring an ISDN Interface for Bandwidth on Demand	238
Configuring Dial-In and Callback (Optional)	238
Configuring Dialer Interfaces for Dial-In and Callback	240
Configuring an ISDN Interface to Screen Incoming Calls	242
Configuring the Services Router to Reject Incoming ISDN Calls	243
Disabling Dialing Out Through Dialer Interfaces	243
Disabling ISDN Signaling	244
Verifying the ISDN Configuration	245
Displaying the ISDN Status	246
Verifying an ISDN BRI Interface	247
Verifying an ISDN PRI Interface and Checking B-Channel Interface Statistics	249
Checking D-Channel Interface Statistics	251
Displaying the Status of ISDN Calls	253
Verifying Dialer Interface Configuration	255

Chapter 8

Configuring Link Services Interfaces 261

Link Services Terms	261
Link Services Interfaces Overview	262
Services Available on J-series Link Services Interface	263
Link Services Exceptions on J-series Services Routers	264
Multilink Bundles Overview	264
Link Fragmentation and Interleaving Overview	265
Compressed Real-Time Transport Protocol Overview	266
Queuing with LFI on J-series Services Routers	267
Queuing on Q0s of Constituent Links	268
Queuing on Q2s of Constituent Links	268

Load Balancing with LFI	268
Configuring CoS Components with LFI	269
Shaping Rate	269
Scheduling Priority	270
Buffer Size	270
Before You Begin	270
Configuring the Link Services Interface with Quick Configuration	271
Configuring the Link Services Interface with a Configuration Editor	273
Configuring MLPPP Bundles and LFI on Serial Links	273
Configuring an MLPPP Bundle	274
Enabling Link Fragmentation and Interleaving	276
Defining Classifiers and Forwarding Classes	277
Defining and Applying Scheduler Maps	279
Applying Shaping Rates to Interfaces	283
Configuring MLFR FRF.15 Bundles	284
Configuring MLFR FRF.16 Bundles	287
Configuring CRTP	289
Verifying the Link Services Interface Configuration	291
Displaying Multilink Bundle Configurations	292
Displaying Link Services CoS Configurations	294
Verifying Link Services Interface Statistics	296
Verifying Link Services CoS	300
Frequently Asked Questions About the Link Services Interface	301
Which CoS Components Are Applied to the Constituent Links?	302
What Causes Jitter and Latency on the Multilink Bundle?	303
Are LFI and Load Balancing Working Correctly?	304
Why Are Packets Dropped on a PVC Between a J-series Router and Another Vendor?	310

Chapter 9

Configuring VoIP 311

VoIP Terms	311
VoIP Overview	313
About the Avaya IG550 Integrated Gateway	314
VoIP Interfaces	315
Avaya VoIP Modules Overview	316
Media Gateway Controller	317
Avaya Communication Manager	318
Dynamic Call Admission Control Overview	318
Supported Interfaces	318
Bearer Bandwidth Limit and Activation Priority	319
Rules for Determining Reported BBL	319
TGM550 Firmware Compatibility with JUNOS Internet Software	320
TGM550 IP Addressing Guidelines	320
VoIP Configuration Overview	321
Before You Begin	322
Configuring VoIP Interfaces with EPW and Disk-on-Key	323
Configuring VoIP Interfaces with Quick Configuration	325

Configuring VoIP with a Configuration Editor	327
Configuring the VoIP Interface (Required)	327
Configuring the Media Gateway Controller List (Required)	329
Configuring an MGC List and Adding Addresses	329
Clearing an MGC List	330
Configuring Dynamic Call Admission Control on WAN Interfaces (Optional)	330
Modifying the IP Address of the TGM550	332
Accessing and Administering the TGM550 CLI	333
TGM550 Access Requirements	334
Connecting Through the TGM550 Console Port	334
Connecting to the TGM550 with SSH	335
Accessing the TGM550 with Telnet	335
Enabling Telnet Service on the TGM550	336
Connecting to the TGM550 with Telnet	336
Disabling Telnet Service on the TGM550	336
Accessing the Services Router from the TGM550	337
Resetting the TGM550	337
Saving the TGM550 Configuration	338
Verifying the VoIP Configuration	338
Verifying the VoIP Interface	339
Verifying the Media Gateway Controller List	341
Verifying Bandwidth Available for VoIP Traffic	341
Frequently Asked Questions About the VoIP Interface	342
TGM550 Is Installed But the VoIP Interface Is Unavailable	342

Chapter 10

Configuring uPIMs as Ethernet Switches 343

Gigabit Ethernet uPIM Switch Overview	343
Joining uPIMs in a Daisy-Chain	344
Configuring Gigabit Ethernet uPIM Switches	344
Verifying Gigabit Ethernet uPIM Switch Configuration	345
Verifying Status of uPIM Switch Ports	346

Part 3

Configuring Routing Protocols

Chapter 11

Routing Overview 349

Routing Terms	349
Routing Overview	354
Networks and Subnetworks	354
Autonomous Systems	355
Interior and Exterior Gateway Protocols	355
Routing Tables	355
Forwarding Tables	356
Dynamic and Static Routing	357
Route Advertisements	357
Route Aggregation	358

RIP Overview	360
Distance-Vector Routing Protocols	360
Maximizing Hop Count	361
RIP Packets	361
Split Horizon and Poison Reverse Efficiency Techniques	362
Limitations of Unidirectional Connectivity	363
OSPF Overview	364
Link-State Advertisements	364
Role of the Designated Router	365
Path Cost Metrics	365
Areas and Area Border Routers	365
Role of the Backbone Area	366
Stub Areas and Not-So-Stubby Areas	367
IS-IS Overview	368
IS-IS Areas	368
Network Entity Titles and System Identifiers	369
IS-IS Path Selection	369
Protocol Data Units	369
IS-IS Hello PDU	369
Link-State PDU	370
Complete Sequence Number PDU	370
Partial Sequence Number PDU	370
BGP Overview	371
Point-to-Point Connections	371
BGP Messages for Session Establishment	372
BGP Messages for Session Maintenance	372
IBGP and EBGP	372
Route Selection	373
Local Preference	374
AS Path	375
Origin	376
Multiple Exit Discriminator	376
Default MED Usage	376
Additional MED Options for Path Selection	377
Scaling BGP for Large Networks	378
Route Reflectors—for Added Hierarchy	378
Confederations—for Subdivision	381

Chapter 12

Configuring Static Routes

383

Static Routing Overview	383
Static Route Preferences	384
Qualified Next Hops	384
Control of Static Routes	384
Route Retention	385
Readvertisement Prevention	385
Forced Rejection of Passive Route Traffic	385
Default Properties	385
Before You Begin	385
Configuring Static Routes with Quick Configuration	386

Configuring Static Routes with a Configuration Editor	387
Configuring a Basic Set of Static Routes (Required)	388
Controlling Static Route Selection (Optional)	389
Controlling Static Routes in the Routing and Forwarding Tables (Optional)	391
Defining Default Behavior for All Static Routes (Optional)	392
Verifying the Static Route Configuration	393
Displaying the Routing Table	393

Chapter 13**Configuring a RIP Network 395**

RIP Overview	395
RIP Traffic Control with Metrics	395
Authentication	396
Before You Begin	396
Configuring a RIP Network with Quick Configuration	396
Configuring a RIP Network with a Configuration Editor	398
Configuring a Basic RIP Network (Required)	398
Controlling Traffic in a RIP Network (Optional)	401
Controlling Traffic with the Incoming Metric	401
Controlling Traffic with the Outgoing Metric	403
Enabling Authentication for RIP Exchanges (Optional)	404
Enabling Authentication with Plain-Text Passwords	404
Enabling Authentication with MD5 Authentication	405
Verifying the RIP Configuration	406
Verifying the RIP-Enabled Interfaces	407
Verifying the Exchange of RIP Messages	408
Verifying Reachability of All Hosts in the RIP Network	409

Chapter 14**Configuring an OSPF Network 411**

OSPF Overview	411
Enabling OSPF	411
OSPF Areas	412
Path Cost Metrics	412
OSPF Dial-on-Demand Circuits	412
Before You Begin	412
Configuring an OSPF Network with Quick Configuration	413
Configuring an OSPF Network with a Configuration Editor	414
Configuring the Router Identifier (Required)	415
Configuring a Single-Area OSPF Network (Required)	415
Configuring a Multiarea OSPF Network (Optional)	417
Creating the Backbone Area	418
Creating Additional OSPF Areas	418
Configuring Area Border Routers	419
Configuring Stub and Not-So-Stubby Areas (Optional)	420
Tuning an OSPF Network for Efficient Operation	422
Controlling Route Selection in the Forwarding Table	422
Controlling the Cost of Individual Network Segments	423

Enabling Authentication for OSPF Exchanges	424
Controlling Designated Router Election	425
Verifying an OSPF Configuration	426
Verifying OSPF-Enabled Interfaces	427
Verifying OSPF Neighbors	428
Verifying the Number of OSPF Routes	429
Verifying Reachability of All Hosts in an OSPF Network	430

Chapter 15 Configuring the IS-IS Protocol 431

IS-IS Overview	431
ISO Network Addresses	431
System Identifier Mapping	432
Before You Begin	432
Configuring IS-IS with a Configuration Editor	433
Verifying IS-IS on a Services Router	434
Displaying IS-IS Interface Configuration	435
Displaying IS-IS Interface Configuration Detail	436
Displaying IS-IS Adjacencies	437
Displaying IS-IS Adjacencies in Detail	438

Chapter 16 Configuring BGP Sessions 441

BGP Overview	441
BGP Peering Sessions	441
IBGP Full Mesh Requirement	442
Route Reflectors and Clusters	442
BGP Confederations	442
Before You Begin	442
Configuring BGP Sessions with Quick Configuration	443
Configuring BGP Sessions with a Configuration Editor	444
Configuring Point-to-Point Peering Sessions (Required)	445
Configuring BGP Within a Network (Required)	447
Configuring a Route Reflector (Optional)	448
Configuring BGP Confederations (Optional)	451
Verifying a BGP Configuration	452
Verifying BGP Neighbors	454
Verifying BGP Groups	455
Verifying BGP Summary Information	456
Verifying Reachability of All Peers in a BGP Network	457

Part 4 Index

Index	461
-------------	-----

About This Guide

This preface provides the following guidelines for using the *J-series™ Services Router Basic LAN and WAN Access Configuration Guide*:

- Objectives on page xix
- Audience on page xix
- How to Use This Guide on page xx
- Document Conventions on page xxi
- Related Juniper Networks Documentation on page xxii
- Documentation Feedback on page xxv
- Requesting Support on page xxv

Objectives

This guide contains instructions for configuring the interfaces on a Services Router for basic IP routing with standard routing protocols. It also shows how to create backup ISDN interfaces and configure digital subscriber line (DSL) connections and link services.

J-series Services Router operations are controlled by the JUNOS Internet software. You direct the JUNOS software through either a Web browser or a command-line interface (CLI).



NOTE: This guide documents Release 8.4 of the JUNOS software. For additional information about J-series Services Routers—either corrections to or omissions from this guide—see the *J-series Services Router Release Notes* at <http://www.juniper.net>.

Audience

This guide is designed for anyone who installs and sets up a J-series Services Router or prepares a site for Services Router installation. The guide is intended for the following audiences:

- Customers with technical knowledge of and experience with networks and the Internet

- Network administrators who install, configure, and manage Internet routers but are unfamiliar with the JUNOS software
- Network administrators who install, configure, and manage products of Juniper Networks

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.

How to Use This Guide

J-series documentation explains how to install, configure, and manage J-series routers by providing information about JUNOS implementation specifically on J-series routers. (For comprehensive JUNOS information, see the JUNOS software manuals listed in “Related Juniper Networks Documentation” on page xxii.) Table 1 on page xx shows the location of J-series information, by task type, in Juniper Networks documentation.

Table 1: Location of J-series Information

J-series Tasks	Location of Instruction
Installing hardware and establishing basic connectivity	Getting Started Guide for your router
Configuring interfaces and routing protocols such as RIP, OSPF, BGP, and IS-IS	<i>J-series Services Router Basic LAN and WAN Access Configuration Guide</i>
Configuring advanced features such as virtual private networks (VPNs), IP Security (IPSec), multicast, routing policies, firewall filters, and class of service (CoS)	<i>J-series Services Router Advanced WAN Access Configuration Guide</i>
Managing users and operations, monitoring performance, upgrading software, and diagnosing common problems	<i>J-series Services Router Administration Guide</i>
Using the J-Web interface	<i>J-Web Interface User Guide</i>
Using the CLI	<i>JUNOS CLI User Guide</i>

Typically, J-series documentation provides both general and specific information—for example, a configuration overview, configuration examples, and verification methods. Because you can configure and manage J-series routers in several ways, you can choose from multiple sets of instructions to perform a task. To make best use of this information:

- *If you are new to the topic*—Read through the initial overview information, keep the related JUNOS guide handy for details about the JUNOS hierarchy, and follow the step-by-step instructions for your preferred interface.
- *If you are already familiar with the feature*—Go directly to the instructions for the interface of your choice, and follow the instructions. You can choose a J-Web method, the JUNOS CLI, or a combination of methods based on the level of complexity or your familiarity with the interface.

For many J-series features, you can use J-Web Quick Configuration pages to configure the router quickly and easily without configuring each statement individually. For more extensive configuration, use the J-Web configuration editor or CLI configuration mode commands.

To monitor, diagnose, and manage a router, use the J-Web interface or CLI operational mode commands.

Document Conventions

Table 2 on page xxi defines the notice icons used in this guide.

Table 2: 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 3 on page xxi defines the text and syntax conventions used in this guide.

Table 3: Text and Syntax Conventions

Convention	Description	Examples
Bold sans serif typeface	Represents text that you type.	To enter configuration mode, type the <code>configure</code> command: user@host> configure
Fixed-width typeface	Represents output that appears on the terminal screen.	user@host> show chassis alarms No alarms currently active
<i>Italic typeface</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>

Table 3: Text and Syntax Conventions (*continued*)

Convention	Description	Examples
<i>Italic sans serif typeface</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>
Sans serif typeface	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> >;
(pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	broadcast multicast (<i>string1</i> <i>string2</i> <i>string3</i>)
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	rsvp { # Required for dynamic MPLS only
[] (square brackets)	Enclose a variable for which you can substitute one or more values.	community name members [<i>community-ids</i>]
Indentation and braces ({ })	Identify a level in the configuration hierarchy.	[edit] routing-options { static { route default { nexthop <i>address</i> ; retain; } } }
; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
J-Web GUI Conventions		
Bold typeface	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 .

Related Juniper Networks Documentation

J-series Services Routers are documented in multiple guides. Although the J-series guides provide instructions for configuring and managing a Services Router with the JUNOS CLI, they are not a comprehensive JUNOS software resource. For complete

documentation of the statements and commands described in J-series guides, see the JUNOS software manuals listed in Table 4 on page xxiii.

Table 4: J-series Guides and Related JUNOS Software Publications

Chapter in a J-series Guide	Corresponding JUNOS Software Manual
Getting Started Guide for Your Router	
“Services Router User Interface Overview”	■ <i>JUNOS CLI User Guide</i>
“Establishing Basic Connectivity”	■ <i>JUNOS System Basics Configuration Guide</i>
J-series Services Router Basic LAN and WAN Access Configuration Guide	
“Using Services Router Configuration Tools”	■ <i>JUNOS CLI User Guide</i> ■ <i>JUNOS System Basics Configuration Guide</i>
“Interfaces Overview”	■ <i>JUNOS Network Interfaces Configuration Guide</i> ■ <i>JUNOS Interfaces Command Reference</i>
“Configuring DS1, DS3, Ethernet, and Serial Interfaces”	
“Configuring Channelized T1/E1/ISDN PRI Interfaces”	
“Configuring Digital Subscriber Line Interfaces”	
“Configuring Point-to-Point Protocol over Ethernet”	
“Configuring ISDN”	
“Configuring Link Services Interfaces”	■ <i>JUNOS Services Interfaces Configuration Guide</i> ■ <i>JUNOS System Basics and Services Command Reference</i>
“Configuring VoIP”	■ <i>JUNOS Network Interfaces Configuration Guide</i> ■ <i>JUNOS Interfaces Command Reference</i>
“Configuring uPIMs as Ethernet Switches”	■ <i>JUNOS Network Interfaces Configuration Guide</i> ■ <i>JUNOS System Basics Configuration Guide</i> ■ <i>JUNOS System Basics and Services Command Reference</i>
“Routing Overview”	■ <i>JUNOS Routing Protocols Configuration Guide</i> ■ <i>JUNOS Routing Protocols and Policies Command Reference</i>
“Configuring Static Routes”	
“Configuring a RIP Network”	
“Configuring an OSPF Network”	
“Configuring the IS-IS Protocol”	
“Configuring BGP Sessions”	
J-series Services Router Advanced WAN Access Configuration Guide	

Table 4: J-series Guides and Related JUNOS Software Publications (continued)

Chapter in a J-series Guide	Corresponding JUNOS Software Manual
“Multiprotocol Label Switching Overview”	■ <i>JUNOS MPLS Applications Configuration Guide</i>
“Configuring Signaling Protocols for Traffic Engineering”	■ <i>JUNOS Routing Protocols and Policies Command Reference</i>
“Configuring Virtual Private Networks”	■ <i>JUNOS VPNs Configuration Guide</i>
“Configuring CLNS VPNs”	
“Configuring IPsec for Secure Packet Exchange”	■ <i>JUNOS System Basics Configuration Guide</i> ■ <i>JUNOS Services Interfaces Configuration Guide</i> ■ <i>JUNOS System Basics and Services Command Reference</i>
“Multicast Overview”	■ <i>JUNOS Multicast Protocols Configuration Guide</i>
“Configuring a Multicast Network”	■ <i>JUNOS Routing Protocols and Policies Command Reference</i>
“Configuring Data Link Switching”	■ <i>JUNOS Services Interfaces Configuration Guide</i> ■ <i>JUNOS System Basics and Services Command Reference</i>
“Policy Framework Overview”	■ <i>JUNOS Policy Framework Configuration Guide</i>
“Configuring Routing Policies”	■ <i>JUNOS Routing Protocols and Policies Command Reference</i>
“Configuring NAT”	■ <i>JUNOS Network Interfaces Configuration Guide</i>
“Configuring Stateful Firewall Filters and NAT”	■ <i>JUNOS Policy Framework Configuration Guide</i>
“Configuring Stateless Firewall Filters”	■ <i>JUNOS Services Interfaces Configuration Guide</i> ■ <i>Secure Configuration Guide for Common Criteria and JUNOS-FIPS</i> ■ <i>JUNOS System Basics and Services Command Reference</i> ■ <i>JUNOS Routing Protocols and Policies Command Reference</i>
“Class-of-Service Overview”	■ <i>JUNOS Class of Service Configuration Guide</i>
“Configuring Class of Service”	■ <i>JUNOS System Basics and Services Command Reference</i>
J-series Services Router Administration Guide	
“Managing User Authentication and Access”	■ <i>JUNOS System Basics Configuration Guide</i> ■ <i>Secure Configuration Guide for Common Criteria and JUNOS-FIPS</i>
“Configuring SNMP for Network Management”	<i>JUNOS Network Management Configuration Guide</i>
“Configuring the Router as a DHCP Server”	<i>JUNOS System Basics Configuration Guide</i>
“Configuring Autoinstallation”	
“Automating Network Operations and Troubleshooting”	<i>JUNOS Configuration and Diagnostic Automation Guide</i>

Table 4: J-series Guides and Related JUNOS Software Publications (continued)

Chapter in a J-series Guide	Corresponding JUNOS Software Manual
“Monitoring the Router and Routing Operations”	<ul style="list-style-type: none"> ■ <i>JUNOS System Basics and Services Command Reference</i> ■ <i>JUNOS Interfaces Command Reference</i> ■ <i>JUNOS Routing Protocols and Policies Command Reference</i>
“Monitoring Events and Managing System Log Files”	<ul style="list-style-type: none"> ■ <i>JUNOS System Log Messages Reference</i> ■ <i>Secure Configuration Guide for Common Criteria and JUNOS-FIPS</i>
“Configuring and Monitoring Alarms”	<i>JUNOS System Basics Configuration Guide</i>
“Performing Software Upgrades and Reboots”	<i>JUNOS Software Installation and Upgrade Guide</i>
“Using Services Router Diagnostic Tools”	<ul style="list-style-type: none"> ■ <i>JUNOS System Basics and Services Command Reference</i> ■ <i>JUNOS Interfaces Command Reference</i> ■ <i>JUNOS Routing Protocols and Policies Command Reference</i>
“Configuring Packet Capture”	<i>JUNOS Services Interfaces Configuration Guide</i>
“Configuring RPM Probes”	<i>JUNOS System Basics and Services Command Reference</i>

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 <http://www.juniper.net/techpubs/docbug/docbugreport.html>. 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

Requesting Support

For technical support, open a support case with the Case Manager link at <http://www.juniper.net/support/> or call 1-888-314-JTAC (from the United States, Canada, or Mexico) or 1-408-745-9500 (from elsewhere).

Part 1

Using the Configuration Interfaces

- Using Services Router Configuration Tools on page 3

Chapter 1

Using Services Router Configuration Tools

Use Services Router configuration tools to configure all services on a router, including system parameters, routing protocols, interfaces, network management, and user access.

This chapter contains the following topics:

- Configuration Tools Terms on page 3
- Configuration Tools Overview on page 4
- Before You Begin on page 7
- Using J-Web Quick Configuration on page 7
- Using the J-Web Configuration Editor on page 9
- Managing Configuration Files with the J-Web Interface on page 16
- Using the CLI Configuration Editor on page 22
- Managing Configuration Files with the CLI on page 35

Configuration Tools Terms

Before using the Services Router configuration tools, become familiar with the terms defined in Table 5 on page 3.

Table 5: Configuration Tools Terms

Term	Definition
candidate configuration	A working copy of the configuration that can be edited without affecting the Services Router until it is committed.
configuration group	Group of configuration statements that can be inherited by the rest of the configuration.
commit a configuration	Have the candidate configuration checked for proper syntax, activated, and marked as the current configuration file running on the Services Router.
configuration hierarchy	The JUNOS software configuration consists of a hierarchy of <i>statements</i> . There are two types of statements: <i>container statements</i> , which contain other statements, and <i>leaf statements</i> , which do not contain other statements. All the container and leaf statements together form the configuration hierarchy.

Table 5: Configuration Tools Terms (*continued*)

Term	Definition
rescue configuration	Configuration that recovers a Services Router from a configuration that denies management access. You set a current committed configuration through the J-Web interface or CLI for emergency use. To load and commit the rescue configuration, you press and release the CONFIG or RESET CONFIG button.
roll back a configuration	Return to a previously committed configuration.

Configuration Tools Overview

The J-Web interface provides a Quick Configuration tool for basic configuration and a configuration editor for complete configuration. You can also use the JUNOS CLI configuration mode as a configuration editor to create and modify a complete configuration hierarchy.

This section contains the following topics:

- Editing and Committing a Configuration on page 4
- J-Web Configuration Options on page 5
- CLI Configuration Commands on page 5

Editing and Committing a Configuration

When you edit a configuration, you work in a copy of the current configuration to create a candidate configuration. The changes you make to the candidate configuration are visible through the user interface immediately, but do not take effect on the Services Router until you *commit* the changes. When you commit the configuration, the candidate file is checked for proper syntax, activated, and marked as the current, operational software configuration file. If multiple users are editing the configuration when you commit the candidate configuration, all changes made by all the users take effect.

If you are editing the configuration with the CLI, you can edit an *exclusive* or *private* candidate configuration. For more information, see “Entering and Exiting Configuration Mode” on page 22.

When you commit a configuration, the Services Router saves the current operational version and the previous 49 versions of committed configurations. The most recently committed configuration is version 0 (the current operational version), and the oldest saved configuration is version 49. You can roll back the configuration to any saved version. Version 0 is stored in the file `juniper.conf`, and the last three committed configurations are stored in the files `juniper.conf.1.gz`, `juniper.conf.2.gz`, and `juniper.conf.3.gz`. These four files are located in the `/config` directory, and the remaining 46 previous versions of committed configurations—files `juniper.conf.4.gz` through `juniper.conf.49.gz`—are stored in the `/var/db/config` directory.



NOTE: You must assign a root password before committing a configuration. For more information, see the *JUNOS Software Installation and Upgrade Guide*.

J-Web Configuration Options

You access the J-Web interface configuration tools by selecting **Configuration** in the task bar. Table 6 on page 5 describes the J-Web configuration options.

Table 6: J-Web Configuration Options

Option	Purpose	Description
Quick Configuration	Basic configuration	Displays options for quick Services Router configuration— Set Up , Secure Access , Interfaces , Users , SNMP , Routing and Protocols , Class of Service , Firewall/NAT , DHCP , IPSec Tunnels , Realtime Performance Monitoring , and Firewall Filters . You can access these options in both the side and main panes. For more information, see “Using J-Web Quick Configuration” on page 7.
View and Edit	Complete configuration	Displays the configuration editor options— View Configuration Text , Edit Configuration , Edit Configuration Text , and Upload Configuration File . For more information, see “Using the J-Web Configuration Editor” on page 9.
History	File management	Displays the Services Router configuration history and a list of users currently editing the configuration. You can compare, roll back, or download specific versions of the configuration. For more information, see “Managing Configuration Files with the J-Web Interface” on page 16.
Rescue	Configuration recovery	Displays options for setting the current configuration as the rescue configuration, and for viewing and deleting the rescue configuration. For more information, see “Setting, Viewing, or Deleting the Rescue Configuration” on page 21.

CLI Configuration Commands

The CLI configuration commands allow you to perform the same configuration tasks you can perform using the J-Web interface. Instead of invoking the tools through a graphical interface, you enter configuration mode commands to perform the tasks.

Table 7 on page 6 provides a summary of the top-level CLI configuration commands.

Table 7: Top-Level CLI Configuration Commands

Command	Function
Managing the Configuration and Configuration Files	
commit	Commit the set of configuration changes in the candidate configuration to take operational effect.
load	Load a configuration from an ASCII configuration file or from terminal input.
rollback	Return to a previously committed configuration.
save	Save the configuration to an ASCII file.
Modifying the Configuration and Its Statements	
activate	Activate a previously deactivated statement or identifier.
annotate	Add a comment to a statement.
copy	Copy and add a statement to the configuration.
deactivate	Deactivate a statement or identifier.
delete	Delete a statement or identifier from the configuration.
insert	Insert an identifier into an existing hierarchy.
rename	Rename an existing statement or identifier.
replace	Replace an identifier or value with another in the existing hierarchy.
set	Create a statement hierarchy and set identifier values.
Navigating the Configuration Hierarchy	
edit	Move inside the specified statement hierarchy.
exit	Exit the current level of the statement hierarchy (same function as quit).
quit	Exit the current level of the statement hierarchy (same function as exit).
top	Return to the top level of configuration mode.
up	Move up one level in the statement hierarchy.
Miscellaneous	
help	Provide help about statements.
run	Issue an operational mode command without leaving configuration mode.
show	Display the current configuration.
status	Display the users currently editing the configuration.

For more information about CLI configuration mode commands, see the JUNOS software configuration guides.

Filtering Configuration Command Output

Certain configuration commands, such as **show** commands, display output. You can filter or redirect the output to a file by including a vertical bar (`|`), called a *pipe*, when you enter the command. For more information, see the *J-series Services Router Administration Guide*.

Before You Begin

To use the J-Web interface and CLI configuration tools, you must have the appropriate access privileges.

Using J-Web Quick Configuration

Use J-Web Quick Configuration to quickly and easily configure the Services Router for basic operation. To access Quick Configuration, select **Configuration > Quick Configuration**. You can select a Quick Configuration option from the side pane by pointer over the Quick Configuration option. Or you can select an option from the main pane (see Figure 1 on page 8).

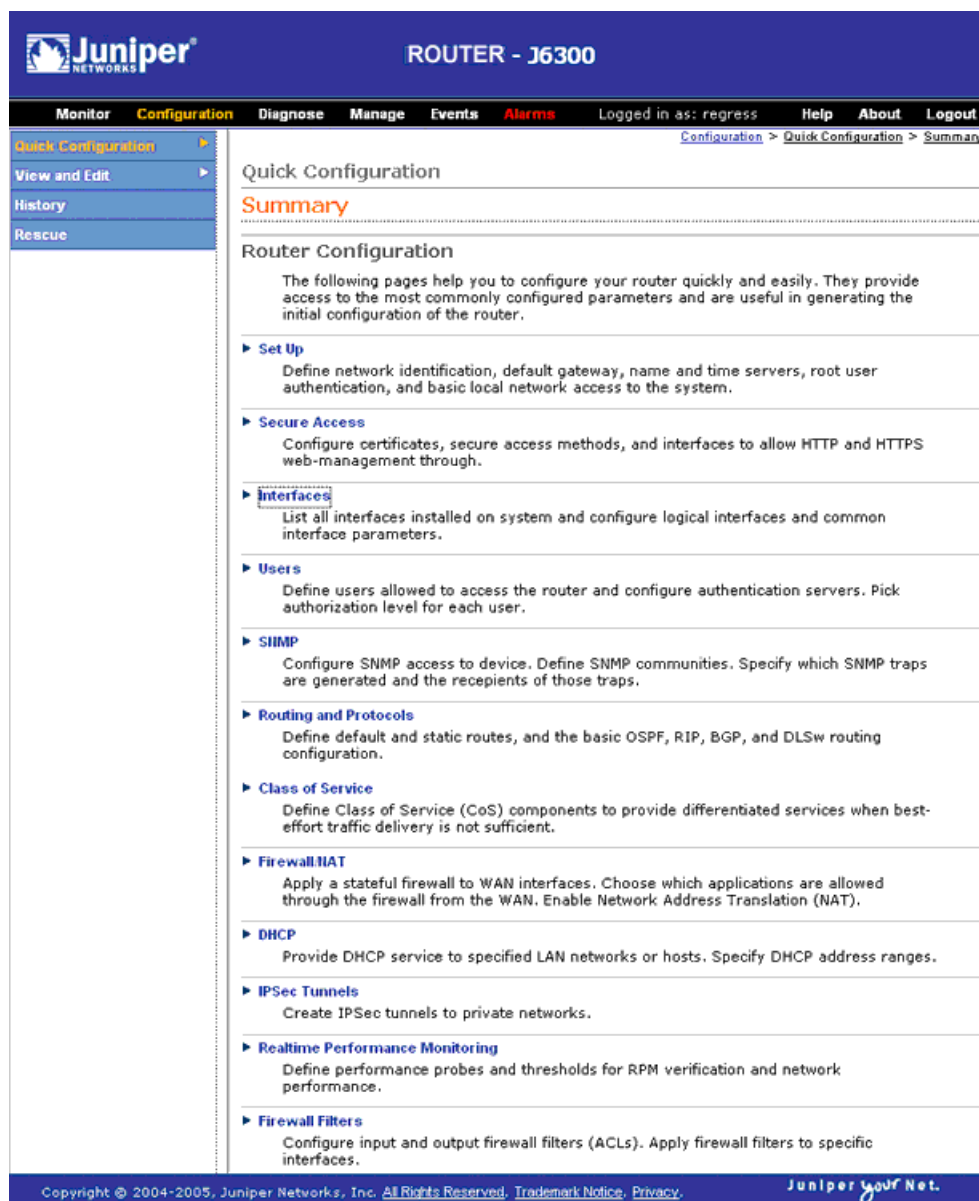
Figure 1: J-Web Quick Configuration Options

Table 8 on page 8 describes the functions of the buttons that appear in the J-Web Quick Configuration pages.

Table 8: J-Web Quick Configuration Buttons

Button	Function
Add	Adds statements or identifiers to the configuration.
Delete	Deletes statements or identifiers from the configuration.

Table 8: J-Web Quick Configuration Buttons *(continued)*

Button	Function
OK	Commits your entries into the configuration, and returns you one level up in the configuration hierarchy.
Cancel	Clears the entries you have not yet applied to the configuration, and returns you one level up in the configuration hierarchy.
Apply	Commits your entries into the configuration, and stays at the same level in the configuration hierarchy.

Using the J-Web Configuration Editor

You can use the J-Web configuration editor to perform the following tasks:

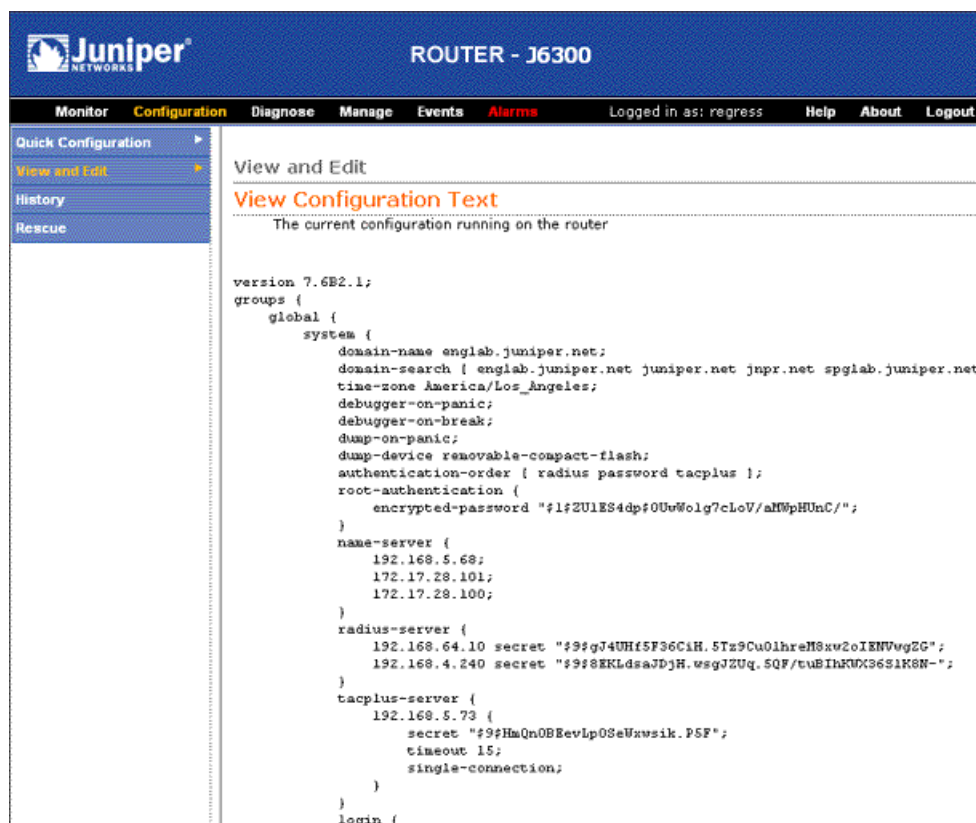
- Viewing the Configuration Text on page 9
- Editing and Committing the Clickable Configuration on page 10
- Editing and Committing the Configuration Text on page 14
- Uploading a Configuration File on page 15

Viewing the Configuration Text

To view the entire configuration in text format, select **Configuration > View and Edit > View Configuration Text**. The main pane displays the configuration in text format (see Figure 2 on page 10).

Each level in the hierarchy is indented to indicate each statement's relative position in the hierarchy. Each level is generally set off with braces, using an open brace ({) at the beginning of each hierarchy level and a closing brace (}) at the end. If the statement at a hierarchy level is empty, the braces are not displayed. Each leaf statement ends with a semicolon (;), as does the last statement in the hierarchy.

This indented representation is used when the configuration is displayed or saved as an ASCII file. However, when you load an ASCII configuration file, the format of the file is not so strict. The braces and semicolons are required, but the indentation and use of new lines are not required in ASCII configuration files.

Figure 2: View Configuration Text Page

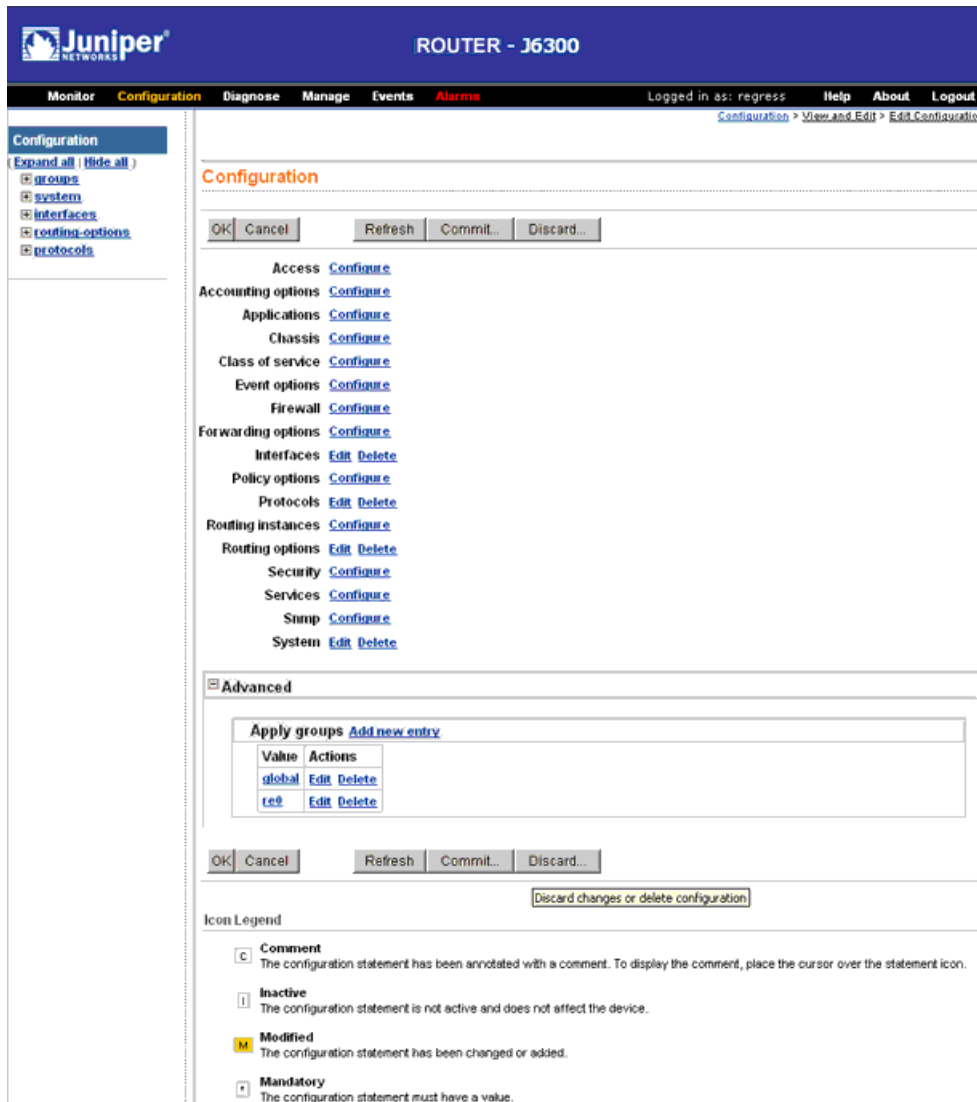
Editing and Committing the Clickable Configuration

Use the J-Web configuration editor's clickable interface to perform the following configuration tasks on a Services Router:

- Editing the Clickable Configuration on page 10
- Discarding Parts of a Candidate Configuration on page 13
- Committing a Clickable Configuration on page 14

Editing the Clickable Configuration

To edit the configuration on a series of pages of clickable options that steps you through the hierarchy, select **Configuration > View and Edit > Edit Configuration**. The side pane displays the top level of the configured hierarchy, and the main pane displays configured hierarchy options and the Icon Legend (see Figure 3 on page 11).

Figure 3: Edit Configuration Page (Clickable)

To expand or hide the hierarchy of all the statements in the side pane, click **Expand all** or **Hide all**. To expand or hide an individual statement in the hierarchy, click the expand (+) or collapse (–) icon to the left of the statement.



NOTE: Only those statements included in the committed configuration are displayed in the hierarchy.

The configuration information in the main pane consists of configuration options that correspond to configuration statements. Configuration options that contain subordinate statements are identified by the term *Nested configuration*.

To include, edit, or delete statements in the candidate configuration, click one of the links described in Table 9 on page 12 in the main pane. Then specify configuration

information by typing into a field, selecting a value from a list, or clicking a check box (toggle).

Table 9: J-Web Edit Clickable Configuration Links

Link	Function
Add new entry	Displays fields and lists for a statement identifier, allowing you to add a new identifier to a statement.
Configure	Displays information for a configuration option that has not been configured, allowing you to include a statement.
Delete	Deletes the corresponding statement or identifier from the configuration. All subordinate statements and identifiers contained within a deleted statement are also discarded.
Edit	Displays information for a configuration option that has already been configured, allowing you to edit a statement.
<i>identifier</i>	Displays fields and lists for an existing statement identifier, allowing you to edit the identifier.

As you navigate through the configuration, the hierarchy level is displayed at the top of the main pane. You can click a statement or identifier in the hierarchy to display the corresponding configuration options in the main pane.

The main pane includes icons that display information about statements and identifiers when you place your cursor over them. Table 10 on page 12 describes the meaning of these icons.

Table 10: J-Web Edit Clickable Configuration Icons

Icon	Meaning
C	Displays a comment about a statement.
I	Indicates that a statement is inactive.
M	Indicates that a statement has been added or modified, but has not been committed.
*	Indicates that the statement or identifier is required in the configuration.
?	Provides help information.



NOTE: You can annotate statements with comments or make them inactive only through the CLI. For more information, see “Deactivating a Statement or Identifier” on page 30 and the *JUNOS CLI User Guide*.

After typing or selecting your configuration edits, click a button in the main pane (described in Table 11 on page 13) to apply your changes or cancel them, refresh the display, or discard parts of the candidate configuration. An updated configuration does not take effect until you commit it.

Table 11: J-Web Edit Clickable Configuration Buttons

Button	Function
OK	Applies edits to the candidate configuration, and returns you one level up in the configuration hierarchy.
Cancel	Clears the entries you have not yet applied to the candidate configuration, and returns you one level up in the configuration hierarchy.
Refresh	Updates the display with any changes to the configuration made by other users.
Commit	Verifies edits and applies them to the current configuration file running on the Services Router. For details, see “Committing a Clickable Configuration” on page 14.
Discard	Removes edits applied to, or deletes existing statements or identifiers from, the candidate configuration. For details, see “Discarding Parts of a Candidate Configuration” on page 13.

Discarding Parts of a Candidate Configuration

Before committing a candidate configuration, you can discard changes you applied or delete existing statements or identifiers.

To discard parts of a candidate configuration:

1. Navigate to the level of the hierarchy you want to edit and click **Discard**.

The main pane displays a list of target statements based on the hierarchy level and the changes you have made.

2. Select an option button (also known as a radio button) to specify the appropriate discard operation or deletion. (Not all buttons appear in all situations.)
 - **Discard Changes Below This Point**—Discards changes made to the candidate configuration at the displayed hierarchy level and below. All subordinate statements and identifiers contained within a discarded statement are also discarded.
 - **Discard All Changes**—Discards all changes made to the candidate configuration.
 - **Delete Configuration Below This Point**—Deletes all changes and statements in the candidate configuration at the displayed hierarchy level and below. All subordinate statements and identifiers contained within a deleted statement are also deleted.
3. To confirm the discard operation or deletion, click **OK**.

To cancel a discard operation or deletion, click **Cancel**.

The updated candidate configuration does not take effect on the Services Router until you commit it.

Committing a Clickable Configuration

When you finish making changes to a candidate configuration with the J-Web configuration editor's clickable interface, you must commit the changes to use them in the current operational software running on the Services Router.

If another user is editing an exclusive candidate configuration with the CLI, you cannot commit a configuration until the user has committed the configuration. To display a list of users, see “Displaying Users Editing the Configuration” on page 18. For more information about editing an exclusive candidate configuration, see “Entering and Exiting Configuration Mode” on page 22.

To commit a candidate configuration:

1. In the J-Web configuration editor's clickable interface, click **Commit**.

The main pane displays a summary of your changes in statement form.

2. To confirm the commit operation, click **OK**.

To cancel a commit operation, click **Cancel**.

If multiple users are editing the configuration when you commit the candidate configuration, all changes made by all users take effect.

3. To display all the edits applied to the running configuration, click **Refresh**.

Editing and Committing the Configuration Text

To edit the entire configuration in text format:



CAUTION: We recommend that you use this method to edit and commit the configuration only if you have experience editing configurations through the CLI.

1. Select **Configuration > View and Edit > Edit Configuration Text**. The main pane displays the configuration in a text editor (see Figure 4 on page 15).

For more information about the format of an ASCII configuration file, see “Viewing the Configuration Text” on page 9.


2. Navigate to the hierarchy level you want to edit.

You can edit the candidate configuration using standard text editor operations—insert lines (by using the Enter key), delete lines, and modify, copy, and paste text.

3. Click **OK** to load and commit the configuration.

The Services Router checks the configuration for the correct syntax before committing it.

Figure 4: Edit Configuration Text Page



Uploading a Configuration File

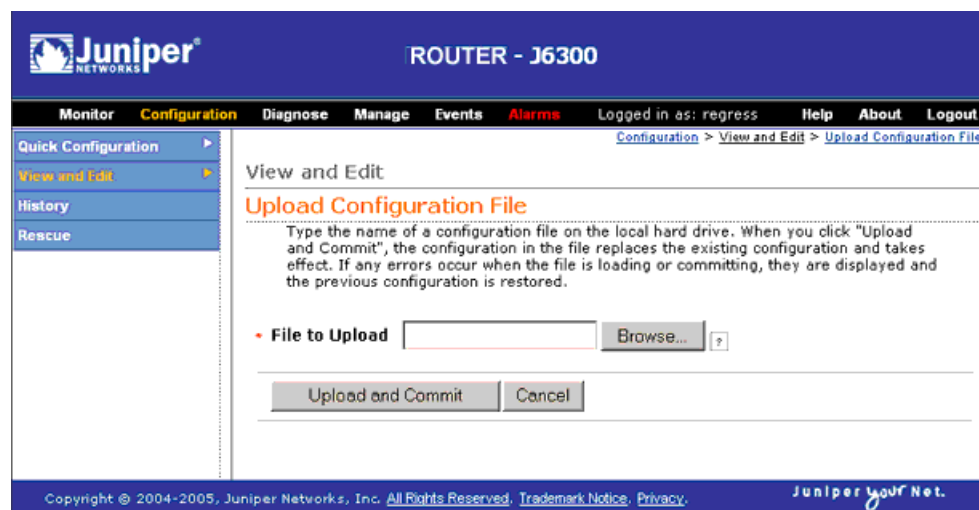
To upload a configuration file from your local system:

1. Select **Configuration > View and Edit > Upload Configuration File**.

The main pane displays the File to Upload box (see Figure 5 on page 16).

2. Specify the name of the file to upload using one of the following methods:
 - Type the absolute path and filename in the File to Upload box.
 - Click **Browse** to navigate to the file.
3. Click **OK** to upload and commit the configuration.

The Services Router checks the configuration for the correct syntax before committing it.

Figure 5: J-Web Upload Configuration File Page

Managing Configuration Files with the J-Web Interface

The J-Web interface provides configuration database and history information that allows you to manage configuration files. This section contains the following topics:

- Configuration Database and History Overview on page 16
- Displaying Users Editing the Configuration on page 18
- Comparing Configuration Files on page 19
- Downloading a Configuration File on page 20
- Loading a Previous Configuration File on page 21
- Setting, Viewing, or Deleting the Rescue Configuration on page 21

Configuration Database and History Overview

When you commit a configuration, the Services Router saves the current operational version and the previous 49 versions of committed configurations. To manage these configuration files with the J-Web interface, select **Configuration > History**. The main pane displays Database Information and Configuration History (see Figure 6 on page 17).

Table 12 on page 17 and Table 13 on page 17 summarize the contents of the display.

Figure 6: Configuration Database and History Page

History

Database Information

The following users are editing the configuration:

User Name	Start Time	Idle Time	Terminal	PID	Edit Flags	Edit Path
root	2005-01-18 14:57:05 PST	00:02:02	d0	2540	None	[edit groups]

Configuration History

The following table shows the router's commit history.

To view a configuration, click the revision number.

To compare configurations, select two and click "Compare".

Compare

	Number	Date/Time	User	Client	Comment	Log Message	Action
<input type="checkbox"/>	Current	2005-01-18 16:12:46 PST	root	cli			Download
<input type="checkbox"/>	1	2005-01-18 15:01:13 PST	root	cli			Download Rollback

Table 12: J-Web Configuration Database Information Summary

Field	Description
User Name	Name of user editing the configuration.
Start Time	Time of day the user logged in to the Services Router.
Idle Time	Elapsed time since the user issued a configuration command from the CLI.
Terminal	Terminal on which the user is logged in.
PID	Process identifier assigned to the user by the Services Router.
Edit Flags	Designates a private or exclusive edit.
Edit Path	Level of the configuration hierarchy that the user is editing.

Table 13: J-Web Configuration History Summary

Field	Description
Number	Version of the configuration file.
Date/Time	Date and time the configuration was committed.
User	Name of the user who committed the configuration.

Table 13: J-Web Configuration History Summary (*continued*)

Field	Description
Client	<p>Method by which the configuration was committed:</p> <ul style="list-style-type: none"> ■ cli—A user entered a JUNOS command-line interface command. ■ junoscript—A JUNOScript client performed the operation. Commit operations performed by users through the J-Web interface are identified in this way. ■ snmp—An SNMP set request started the operation. ■ button—The CONFIG or RESET CONFIG button on the router was pressed to commit the rescue configuration (if set) or to clear all configurations except the factory configuration. ■ autoinstall—Autoinstallation was performed. ■ other—Another method was used to commit the configuration.
Comment	Comment.
Log Message	<p>Method used to edit the configuration:</p> <ul style="list-style-type: none"> ■ Imported via paste—Configuration was edited and loaded with the Configuration > View and Edit > Edit Configuration Text option. For more information, see “Editing and Committing the Configuration Text” on page 14. ■ Imported upload [filename]—Configuration was uploaded with the Configuration > View and Edit > Upload Configuration File option. For more information, see “Uploading a Configuration File” on page 15. ■ Modified via quick-configuration—Configuration was modified using the J-Web Quick Configuration tool specified by <i>quick-configuration</i>. For more information, see “Using J-Web Quick Configuration” on page 7. ■ Rolled back via user-interface—Configuration was rolled back to a previous version through the user interface specified by <i>user-interface</i>, which can be Web Interface or CLI. For more information, see “Loading a Previous Configuration File” on page 21.
Action	Action to perform with the configuration file. The action can be Download or Rollback . For more information, see “Downloading a Configuration File” on page 20 and “Loading a Previous Configuration File” on page 21.

The configuration history display allows you to perform the following operations:

- View a configuration.
- Compare two configurations.
- Download a configuration file to your local system.
- Roll back the configuration to any of the previous versions stored on the Services Router.

For more information about saved versions of configuration files, see “Editing and Committing a Configuration” on page 4.

Displaying Users Editing the Configuration

To display a list of users editing the Services Router configuration, select **Configuration > History**. The list is displayed as Database Information in the main

pane (see Figure 6 on page 17). Table 12 on page 17 summarizes the Database Information display.

Comparing Configuration Files

To compare any two of the past 50 committed configuration files:

1. Select **Configuration > History**.

A list of the current and previous 49 configurations is displayed as Configuration History in the main pane (see Figure 6 on page 17). Table 13 on page 17 summarizes the Configuration History display.

2. Click two of the check boxes to the left of the configuration versions you want to compare.
3. Click **Compare**.

The main pane displays the differences between the two configuration files at each hierarchy level as follows (see Figure 7 on page 20):

- Lines that have changed are highlighted side by side in green.
- Lines that exist only in the more recent configuration file are displayed in red on the left.
- Lines that exist only in the least recent configuration file are displayed in blue on the right.

Figure 7: J-Web Configuration File Comparison Results

History	
Compare Rollback 49 Configuration to Current Configuration	
<div> <div>Legend:</div> <div>Removed from Rollback 49 Configuration</div> <div>changed lines</div> <div>Added in Current Configuration</div> </div>	
Rollback 49 Configuration	Current Configuration
[edit]	[edit]
version 8.2R2;	version "8.310 [builder]";
[edit groups global system radius-server]	[edit groups global system radius-server]
192.168.170.241 secret "\$9\$-Sd2aji..mfQoaGihqTQn/Cu1RcyeXNVApv8X-sY"; ## SECRET-DATA	192.168.170.241 secret "\$9\$Sa8yMXVb2goZLXNbwYJZjHqfz3/CpRor.P01RSKv"; ## SECRET-DATA
192.168.64.10 secret "\$9\$nMNB6pBcSeKML0BR Syl8LNdb2aZDqQ3/wYm5QnAt"; ## SECRET-DATA	192.168.64.10 secret "\$9\$8.xLdsaJDjH.wsgJZUq.5QF/tuBIhKWX36SIK8N-"; ## SECRET-DATA
192.168.4.240 secret "\$9\$McBWNb4oGUjkVbYoaZHkP5QnCt0BRlv8z3hylMx7"; ## SECRET-DATA	192.168.4.240 secret "\$9\$f5nCOEhSl9CpB1RrIM8X-wY4aGqPTxNDHqfF3"; ## SECRET-DATA
[edit groups global system tacplus-server 192.168.5.73]	[edit groups global system tacplus-server 192.168.5.73]
secret "\$9\$upvF1d7NboJDLxYoGifSp0BIEy"; ## SECRET-DATA	secret "\$9\$T3CuSyKxNbEcWxdsJZ5QFn/t"; ## SECRET-DATA
[edit]	[edit]
system { ... }	system { ... }
[edit system]	[edit system]
	<pre> services { web-management { http; control { max-child-process 15; } } } </pre>
<pre> syslog { time-format year millisecond; } chassis { aggregated-devices { ethernet { device-count 1; } } } </pre>	

Downloading a Configuration File

To download a configuration file from the Services Router to your local system:

1. Select **Configuration > History**.

A list of the current and previous 49 configurations is displayed as Configuration History in the main pane (see Figure 6 on page 17). Table 13 on page 17 summarizes the Configuration History display.

2. In the Action column, click **Download** for the version of the configuration you want to download.
3. Select the options your Web browser provides that allow you to save the configuration file to a target directory on your local system.

The file is saved as an ASCII file.

Loading a Previous Configuration File

To download a configuration file from the Services Router to your local system:

To load (roll back) and commit a previous configuration file stored on the Services Router:

1. Select **Configuration > History**.

A list of the current and previous 49 configurations is displayed as Configuration History in the main pane (see Figure 6 on page 17). Table 13 on page 17 summarizes the Configuration History display.

2. In the Action column, click **Rollback** for the version of the configuration you want to load.

The main pane displays the results of the rollback operation.



NOTE: When you click **Rollback**, the Services Router loads and commits the selected configuration. This behavior is different from entering the **rollback** configuration mode command from the CLI, where the configuration is loaded, but not committed.

Setting, Viewing, or Deleting the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to a Services Router, you can delete the invalid configuration and replace it with a rescue configuration by pressing the **CONFIG** or **RESET CONFIG** button on the router. You must have previously set the rescue configuration through the J-Web interface or the CLI. The rescue configuration is a previously committed, valid configuration.



CAUTION: Pressing and holding the **CONFIG** or **RESET CONFIG** button for 15 seconds or more—until the configuration LED blinks red—deletes *all* configurations on the router, including the backup configurations and rescue configuration, and loads and commits the factory configuration.

You can change the default behavior of the **CONFIG RESET CONFIG** button. For more information, see “Disabling the CONFIG or RESET CONFIG Button” on page 34.

To view, set, or delete the rescue configuration, select **Configuration > Rescue**. On the Rescue page, you can perform the following tasks:

- View the current rescue configuration—Click **View rescue configuration**.
- Set the current running configuration as the rescue configuration—Click **Set rescue configuration**.
- Delete the current rescue configuration—Click **Delete rescue configuration**.

Using the CLI Configuration Editor

You can use the CLI configuration editor to perform the following tasks:

- Entering and Exiting Configuration Mode on page 22
- Navigating the Configuration Hierarchy on page 23
- Modifying the Configuration on page 25
- Committing a Configuration with the CLI on page 30
- Disabling the CONFIG or RESET CONFIG Button on page 34
- Entering Operational Mode Commands During Configuration on page 34

Entering and Exiting Configuration Mode

You must have access privileges to edit the configuration. For more information, see “Before You Begin” on page 7.

To enter and exit configuration mode:

1. At the CLI prompt, enter the **configure** operational mode command.

Select the form of the **configure** command (see Table 14 on page 23) that is appropriate for the way you want to edit and commit the candidate configuration. For example:

```
user@host> configure
user@host#
```

2. To display the users currently editing the configuration, enter the **status** command:

```
user@host# status

Users currently editing the configuration:
  user1 terminal p1 (pid 66847) on since 2004-04-19 12:32:56 PDT
    [edit]
  user2 terminal p2 (pid 85743) on since 2004-04-19 11:44:06 PDT
    [edit interfaces]
```

For each user, the CLI displays the username, terminal, process identifier, login date and time, and hierarchy level being edited. You can specify the terminal and process identifier in the **request system logout** command.

3. To exit configuration mode and return to operational mode:

- For the top level, enter the following command:

```
user@host# exit
```

- From any level, enter the following command:

```
user@host# exit configuration-mode
```


For more information about the **configure** command, including restrictions on entering and exiting the various configuration modes, see the *JUNOS CLI User Guide*.

Table 14: Forms of the configure Command

Command	Edit Access	Commit Access
configure	<ul style="list-style-type: none"> ■ No one can lock the configuration. All users can make configuration changes. ■ When you enter configuration mode, the CLI displays the following information: <ul style="list-style-type: none"> ■ A list of the other users editing the configuration. ■ Hierarchy levels the users are viewing or editing. ■ Whether the configuration has been changed, but not committed. 	<ul style="list-style-type: none"> ■ No one can lock the configuration. All users can commit all changes to the candidate configuration. ■ If you and another user make changes and the other user commits changes, your changes are committed as well.
configure exclusive	<ul style="list-style-type: none"> ■ One user locks the configuration and makes changes without interference from other users. ■ Other users can enter and exit configuration mode, but they cannot change the configuration. ■ If you enter configuration mode while another user has locked the configuration, the CLI displays the user and the hierarchy level the user is viewing or editing. ■ If you enter configuration mode while another user has locked the configuration, you can forcibly log out that user with the request system logout user operational mode command. (For details, see the <i>JUNOS System Basics and Services Command Reference</i>.) 	
configure private	<ul style="list-style-type: none"> ■ Multiple users can edit the configuration at the same time. ■ Each user has a private candidate configuration to edit independently of other users. 	<ul style="list-style-type: none"> ■ When you commit the configuration, the Services Router verifies that the operational (running) configuration has not been modified by another user before accepting your private candidate configuration as the new operational configuration. ■ If the configuration has been modified by another user, you can merge the modifications into your private candidate configuration and attempt to commit again.

Navigating the Configuration Hierarchy

When you first enter configuration mode, you are at the top level of the configuration command hierarchy, which is indicated by the **[edit]** banner. To move down through an existing configuration command hierarchy, or to create a hierarchy and move down to that level, use the **edit** command, specifying the hierarchy level at which you want to be:

```
user@host# edit <statement-path> <identifier>
```

Replace *statement-path* with the hierarchy level and *identifier* with a string that identifies an instance of a statement. (Not all statements require identifiers.) If the identifier contains a space, you must enclose the identifier in quotation marks (" ").

After you enter an **edit** command, the banner changes to indicate your current level in the hierarchy:

```
[edit]
user@host# edit protocols ospf
[edit protocols ospf]
user@host#
```

To move back up to the previous hierarchy level, enter the **exit** command. This command is, in effect, the opposite of the **edit** command. For example:

```
[edit]
user@host# edit protocols ospf
[edit protocols ospf]
user@host# edit area 0.0.0.0
[edit protocols ospf area 0.0.0.0]
user@host# exit
[edit protocols ospf]
user@host# exit
[edit]
user@host#
```

To move up one level, enter the **up** command. For example:

```
[edit]
user@host# edit protocols ospf area 0.0.0.0
[edit protocols ospf area 0.0.0.0]
user@host# up
[edit protocols ospf]
user@host# up
[edit protocols]
user@host# up
[edit]
user@host#
```

To move directly to the top level of the hierarchy, enter the **top** command. For example:

```
[edit protocols ospf area 0.0.0.0]
user@host# top
[edit]
user@host#
```

To display the configuration, enter the **show** command:

```
show <statement-path>
```

The configuration at the current hierarchy level, or at the level specified by *statement-path*, is displayed. For example, entering the **show** command in each of the following cases displays the same level of the configuration:

```
[edit]
user@host# show interfaces ge-0/0/0
unit 0 {
  family inet {
    address 192.168.4.1/30;
```

```

    }
  }
[edit]
user@host# edit interfaces ge-0/0/0
[edit interfaces ge-0/0/0]
user@host# show
unit 0 {
  family inet {
    address 192.168.4.1/30;
  }
}

```

Modifying the Configuration

You can modify the configuration by performing the following operations:

- Adding or Modifying a Statement or Identifier on page 25
- Using Search and Replace on page 26
- Deleting a Statement or Identifier on page 27
- Copying a Statement on page 27
- Renaming an Identifier on page 28
- Inserting an Identifier on page 28
- Deactivating a Statement or Identifier on page 30

Adding or Modifying a Statement or Identifier

To add or modify statements in a configuration, use the **set** command:

```
set <statement-path> statement <identifier>
```

Replace *statement-path* with the path to the statement from the current hierarchy level, and *statement* with the statement itself. Replace *identifier* with a string that identifies an instance of a statement. (Not all statements require identifiers.) If the identifier contains a space, you must enclose the identifier in quotation marks (" ").

If the statement or identifier does not exist in the configuration hierarchy, it is added. If the statement or identifier already exists, it is modified (unless multiple occurrences of the same statement or identifier are allowed in the configuration, in which case another instance is added to the configuration). After you enter the **set** command, you remain at the same level in the hierarchy.

You can enter a single **set** command from the top level of the hierarchy. Alternatively, you can enter the **edit** command to move to the target hierarchy level, from which you can enter the **set** command. In either case, the CLI creates the hierarchy level if it does not exist. For example, to set the OSPF hello interval from the top level of the hierarchy, enter the **set** command as follows:

```

[edit]
user@host# set protocols ospf area 0.0.0.0 interface t1-0/0/0 hello-interval 5

```

Alternatively, use the **edit** command to create and move to the [edit protocols ospf area 0.0.0.0 interface t1-0/0/0] hierarchy level, then enter a **set** command to set the value of the hello-interval statement:

```
[edit]
user@host# edit protocols ospf area 0.0.0.0 interface t1-0/0/0
[edit protocols ospf area 0.0.0.0 interface t1-0/0/0]
user@host# set hello-interval 5
```

Using Search and Replace

Modifying a configuration often requires you to search for an identifier or value and replace it with another. For example, suppose you change the IP address of the ge-0/0/1 interface and need to update it throughout the configuration.

The **replace** configuration mode command allows you to find and replace identifiers or values without manually searching the configuration hierarchy on the router or uploading the configuration file into a text editor. The **replace** command supports regular expressions, and allows you to perform find and replace operations using complex search criteria. For example, you can use regular expressions to find and replace all instances of IP addresses starting with 192.



NOTE: The search-and-replace operation of the **replace** command is case-sensitive, and case-sensitivity cannot be disabled. In addition, the **replace** command does not support lazy quantifiers in regular expressions.

You can use the **replace** command at any level in the configuration hierarchy. Identifiers or values specified are replaced from that level downwards in the configuration hierarchy. For example, if you use the **replace** command at the [edit interfaces] hierarchy level to change an IP address from 10.1.1.1 to 10.2.1.1, the changes are made in the configuration hierarchy under **interfaces**, but not throughout the configuration.

To replace an identifier or value in the configuration, enter the **replace** command as follows:

replace pattern identifier 1 with identifier 2 upto number

The **upto** option does not specify the number of instances of the pattern to be replaced, but the number of configuration objects within which the search-and-replace operation is performed. For example, if you issue the **replace** command at the [edit interfaces] hierarchy level, and specify **3** for **upto**, the search-and-replace operation is performed within the first three configuration objects under **interfaces**—for example, ge-0/0/0, e1-5/0/1, and lo0.

If you do not specify the **upto** option, all occurrences of the specified pattern are replaced from the configuration hierarchy level downwards.

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

The examples in Table 15 on page 27 illustrate ways in which you can use the **replace** command.

Table 15: Sample replace Commands

Sample Command	Result
[edit] user@host# replace pattern 10.1.1.1 with 10.2.1.1	Replaces 10.1.1.1 with 10.2.1.1 throughout the configuration.
[edit interfaces ge-3/0/1 unit 0] user@host# replace pattern 1bf5 with 1bf4	Replaces all the instances of the last byte of the IPv6 address 1bf5 with 1bf4 under [edit interfaces ge-3/0/1 unit 0].
[edit] user@host# replace pattern "(.*)1bf5" with "\1bf5"	Replaces all instances of the IPv6 address 2000::c0a8::1bf5 with 2000::c0a8:1bf5.

Deleting a Statement or Identifier

To delete a statement or identifier from the configuration, enter the **delete** command:

```
delete <statement-path> <identifier>
```

When you delete a statement, the statement and all its subordinate statements and identifiers are removed from the configuration and revert to their default values. To delete the entire hierarchy starting at the current level, enter the **delete** command without specifying a statement or an identifier. You are prompted to confirm the deletion.

As with the **set** command, you can enter a single **delete** command from the top level of the hierarchy, or you can use the **edit** command to move to the target hierarchy level, from which you can enter the **delete** command.

Copying a Statement

To make a copy of an existing statement in the configuration, use the **copy** command:

```
copy existing-statement to new-statement
```

The existing statement and all its subordinate statements are copied and added to the configuration. After you enter the **copy** command, the configuration might not be valid. If necessary, modify the existing statement or the new statement to ensure the configuration is valid.

The following example shows how to copy a unit configured at the [edit interfaces ge-0/0/0] hierarchy level:

```
[edit interfaces ge-0/0/0]
user@host# show
unit 0 {
  family inet {
    address 10.14.1.1/24;
```

```

    }
  }
[edit interfaces ge-0/0/0]
user@host# copy unit 0 to unit 1
[edit interfaces ge-0/0/0]
user@host# show
unit 0 {
  family inet {
    address 10.14.1.1/24;
  }
}
unit 1 {
  family inet {
    address 10.14.1.1/24;
  }
}

```

In this example, after you enter the **copy** command, unit 0 and unit 1 have the same IP address in the candidate configuration. To modify the IP address of unit 1 before committing the configuration, use the **rename** command as described in “Renaming an Identifier” on page 28.

Renaming an Identifier

There are two ways to rename an identifier that already exists in a configuration:

- Delete the identifier with the **delete** command, then add it back into the configuration with the **set** command.
- Rename the identifier with the **rename** command:

```
rename <statement-path> identifier1 to identifier2
```

In the example provided in “Copying a Statement” on page 27, to rename the IP address of unit 1 from 10.14.1.1/24 to 10.14.2.1/24, enter the **rename** command as follows:

```
user@host# rename interfaces ge-0/0/0 unit 1 family inet address 10.14.1.1/24  
to address 10.14.2.1/24
```

Inserting an Identifier

To insert an identifier into a specific location within the configuration, use the **insert** command:

```
insert <statement-path> identifier1 (before | after) identifier2
```

Generally, you can add most identifiers into the configuration in any order. However, when you are inserting identifiers that must be analyzed in order—such as terms in a routing policy or firewall filter—you must specify **before** or **after**. If you do not specify where to insert an identifier with the **insert** command, the identifier is placed at the end of the list of similar identifiers.

In the following example, the firewall filter terms were added to the configuration in the following order: term1, term3, term2. The insert command is used to insert term2 before term3.

```
[edit]
user@host# show firewall
family inet {
  filter filter1 {
    term term1 {
      from {
        address {
          192.168.0.0/16;
        }
      }
      then {
        reject;
      }
    }
    term term3 {
      then {
        reject;
      }
    }
    term term2 {
      from {
        destination-port ssh;
      }
      then accept;
    }
  }
}
[edit]
user@host# insert firewall family inet filter filter1 term term2 before term term3
[edit]
user@host# show firewall
family inet {
  filter filter1 {
    term term1 {
      from {
        address {
          192.168.0.0/16;
        }
      }
      then {
        reject;
      }
    }
    term term2 {
      from {
        destination-port ssh;
      }
      then accept;
    }
    term term3 {
      then {
        reject;
      }
    }
  }
}
```

```

    }
  }
}

```

Deactivating a Statement or Identifier

You can deactivate a statement or identifier so that it does not take effect when you enter the `commit` command. Any deactivated statements and identifiers are marked with the `inactive:` tag and remain in the configuration.

To deactivate a statement or identifier, use the `deactivate` command:

```
deactivate (statement | identifier)
```

To reactivate a statement or identifier, use the `reactivate` command:

```
reactivate (statement | identifier)
```

Reactivate removes the `inactive:` tag so that a statement or identifier takes effect when you commit the configuration.

In both commands, *statement* or *identifier* must be at the current hierarchy level.

The following example shows how to deactivate interface `ge-0/0/0` at the `[edit interfaces]` hierarchy level:

```

[edit interfaces]
user@host# deactivate ge-0/0/0
[edit interfaces]
user@host# show
inactive: ge-0/0/0 {
  unit 0 {
    family inet {
      address 10.14.1.1/24;
    }
  }
}

```

Committing a Configuration with the CLI

To save candidate configuration changes to the configuration database and activate the configuration on the Services Router, enter the `commit` command from any hierarchy level:

```

[edit]
user@host# commit
commit complete

```

If more than one user is modifying the configuration, committing it saves and activates the changes made by all the users.

The Services Router checks the configuration for syntax errors. If the syntax is correct, the configuration is activated and becomes the current, operational configuration.

running on the Services Router. If the configuration contains syntax errors, the router sends a message indicating the location of the error and does not activate the configuration. The error message has the following format:

```
[edit edit-path]
  offending-statement;
  error-message
```

You can specify one or more options within the **commit** command—or use it with the **rollback** command—to perform the following operations:

- Verifying a Configuration on page 31
- Committing a Configuration and Exiting Configuration Mode on page 31
- Committing a Configuration That Requires Confirmation on page 31
- Scheduling and Canceling a Commit on page 32
- Loading a Previous Configuration File with the CLI on page 32
- Setting or Deleting the Rescue Configuration with the CLI on page 33

Verifying a Configuration

To verify that the syntax of a configuration is correct, enter the **commit check** command:

```
[edit]
user@host# commit check
configuration check succeeds
```

If the configuration contains syntax errors, a message indicates the location of the error.

Committing a Configuration and Exiting Configuration Mode

To save candidate configuration changes, activate the configuration on the Services Router, and exit configuration mode, enter the **commit and-quit** command:

```
[edit]
user@host# commit and-quit
commit complete
exiting configuration mode
user@host>
```

If the configuration contains syntax errors, a message indicates the location of the error.

Committing a Configuration That Requires Confirmation

You can commit the current candidate configuration but require an explicit confirmation for the committed configuration to become permanent. This commit process is useful for verifying that a configuration change works correctly and does not prevent management access to the Services Router. If the change prevents access

or causes other errors, an automatic rollback to the previous configuration restores access after the rollback confirmation timeout expires.

To commit the current candidate configuration, but require an explicit confirmation for the commit to become permanent, use the **commit confirmed** command:

commit confirmed <minutes>

Replace *minutes* with the number of minutes to allow for the timeout period. The default value is 10 minutes.

To make the new configuration permanent, enter the **commit** or **commit check** command within the timeout period specified in the **commit confirmed** command. If the commit is not confirmed within the timeout period, the Services Router automatically rolls back to the previous configuration.

If the configuration contains syntax errors, a message indicates the location of the error.

Scheduling and Canceling a Commit

To schedule a candidate configuration for a commit operation at a future time or the next time the Services Router is rebooted, use the **commit at** command:

commit at *string*

Replace *string* with **reboot** or the time at which the configuration is to be committed, in one of the following formats:

- *hh:mm* <:ss> —Hours, minutes, and seconds (optional), in 24-hour format. For example, 20:30 is 8:30 PM.
- *yyy-mm-dd hh:mm* <:ss> —Year, month, date, hours, minutes, and seconds (optional), in 24-hour format. For example, 2004-09-05 08:00 is September 5, 2004 at 8:00 AM.

The Services Router checks the configuration. If the result of the check is successful, the current user is logged out of configuration mode, and the configuration data is left in a read-only state. No other commit operation can be performed until the scheduled one is completed. If the configuration contains syntax errors, a message indicates the location of the error.

To cancel a pending commit operation, use the **clear system commit operational** mode command. For more information, see the *JUNOS System Basics and Services Command Reference*.

Loading a Previous Configuration File with the CLI

To load, or *roll back*, a previous configuration file stored on the Services Router without activating it, use the **rollback** command:

rollback <string>

Replace *string* with a value from 0 through 49, or **rescue** (if a rescue configuration exists). The default value is 0.

When you commit a configuration, the Services Router saves the current operational version and the previous 49 versions of committed configurations. The most recently committed configuration is version 0 (the current operational version), and the oldest saved configuration is version 49.

If you have defined a rescue configuration, you can roll back to this configuration by entering **rollback rescue**. (You can also roll back to the rescue configuration or the default factory configuration by pressing the **CONFIG** or **RESET CONFIG** button on the Services Router. For more information, see the Getting Started Guide for your router.)

To set the rescue configuration, see “Setting or Deleting the Rescue Configuration with the CLI” on page 33.

For more information about saved versions of configuration files, see “Editing and Committing a Configuration” on page 4.

To activate the configuration you loaded, you must commit it:

```
[edit]
user@host# rollback 2
load complete
[edit]
user@host# commit
```

To display previous configurations, including the rollback number, date, time, name of the user who committed changes, and commit method, use the **rollback ?** command:

```
user@host# rollback ?
Possible completions:
<[Enter]>      Execute this command
0              2004-05-27 14:50:05 PDT by root via junoscript
1              2004-05-27 14:00:14 PDT by root via cli
2              2004-05-27 13:16:19 PDT by snmpset via snmp
...
28             2004-05-21 16:56:25 PDT by root via cli
rescue         2004-05-27 14:30:23 PDT by root via cli
|              Pipe through a command
```

The access privilege level for using the **rollback** command is controlled by the **rollback** permission bit. Users for whom this permission bit is not set can return only to the most recently committed configuration. Users for whom this bit is set can return to any prior committed configuration. For more information, see the *JUNOS System Basics Configuration Guide*.

Setting or Deleting the Rescue Configuration with the CLI

If someone inadvertently commits a configuration that denies management access to the Services Router, you can delete the invalid configuration and replace it with a rescue configuration by pressing the **CONFIG** or **RESET CONFIG** button on the router.

You must have previously set the rescue configuration through the J-Web interface or the CLI. The rescue configuration is a previously committed, valid configuration.



CAUTION: Pressing and holding the **CONFIG** or **RESET CONFIG** button for 15 seconds or more—until the configuration LED blinks red—deletes *all* configurations on the router, including the backup configurations and rescue configuration, and loads and commits the factory configuration.

To set the current running configuration as the rescue configuration, use the following command:

```
user@host> request system configuration rescue save
```

To delete the current rescue configuration, use the following command:

```
user@host> request system configuration rescue delete
```

Disabling the CONFIG or RESET CONFIG Button

You can change the default behavior of the **CONFIG** or **RESET CONFIG** button by including the **config-button** statement at the [edit chassis] hierarchy level:

```
config-button <no-rescue> <no-clear>
```

The **no-rescue** option prevents the **CONFIG** or **RESET CONFIG** button from loading the rescue configuration. The **no-clear** option prevents the **CONFIG** or **RESET CONFIG** button from deleting all configurations on the router.

To return the function of the **CONFIG** or **RESET CONFIG** button to its default behavior, do not include the **config-button** statement in the router configuration.

Entering Operational Mode Commands During Configuration

While in configuration mode, you might need to enter an operational mode command, such as **show** or **request**. To enter a single operational mode command, first enter the **run** command and then specify the operational mode command as follows:

```
user@host# run operational-mode-command
```

For example, to display a pending system reboot while in configuration mode, enter the **show system reboot** operational mode command as follows:

```
[edit]
user@host# run show system reboot
No shutdown/reboot scheduled.
```

If you are in operational mode, the **show cli history** command displays the history of the operational mode commands issued. To display the history of the configuration mode commands issued, enter the **show cli history** command from configuration mode as follows:

```
[edit]
user@host# run show cli history
15:32:51 – exit
15:52:02 – load merge terminal
17:07:57 – run show ospf statistics
17:09:12 – exit
17:18:49 – run show cli history
```

Managing Configuration Files with the CLI

This section contains the following topics:

- Loading a New Configuration File on page 35
- Saving a Configuration File on page 37

Loading a New Configuration File

You can create a configuration file, copy the file to the Services Router, and then load the file into the CLI. After you load the file, you can commit it to activate the configuration on the router, or you can edit the configuration interactively with the CLI and commit it at a later time.

You can also create a configuration while typing at the terminal and then load it. Loading a configuration from the terminal is generally useful when you are cutting existing portions of the configuration and pasting them elsewhere in the configuration.

To load an existing configuration file that is located on the router, use the following version of the `load` command:

load (merge | override | patch | replace | update) filename <relative>

To load a configuration from the terminal, use the following version of the `load` command:

load (merge | override | patch | replace | update) terminal <relative>

Use the `load` command options provided in Table 16 on page 35. (The *incoming configuration* is the configuration in *filename* or the one that you type at the terminal). For more information about loading a configuration, see the *JUNOS CLI User Guide*.

Table 16: Load Configuration File Options

Option	Function
merge	Combines the current configuration and the incoming configuration. A merge operation is useful when you are adding a new section to an existing configuration. If the existing configuration and the incoming configuration contain conflicting statements, the statements in the incoming configuration override those in the existing configuration.
override	Discards the current candidate configuration and loads the incoming configuration.
patch	Changes part of the configuration with the incoming configuration and marks only those parts as changed.

Table 16: Load Configuration File Options (*continued*)

Option	Function
relative	Allows you to use the merge , replace , and update options without specifying the full hierarchy level.
replace	<p>Replaces portions of the configuration based on the replace: tags in the incoming configuration. The Services Router searches for the replace: tags, deletes the existing statements of the same name (if any), and replaces them with the incoming configuration. If no statement of the same name exists in the configuration, the replace operation adds it to the configuration.</p> <p>If you are performing a replace operation and the incoming configuration does not contain any replace: tags, the replace operation is equivalent to a merge operation. If you are running automated scripts and cannot know in advance whether the scripts need to perform a replace or a merge operation, the scripts can use the replace operation to cover either case.</p> <p>If you are performing an override or merge operation and the incoming configuration contains replace: tags, the tags are ignored and the override or merge operation is performed.</p>
update	Replaces only the configuration that has changed. An update operation compares the current configuration to the current candidate configuration, and loads only the changes between these configurations in the incoming configuration.

Figure 8 on page 36 through Figure 10 on page 37 show the results of override, replace, and merge operations.

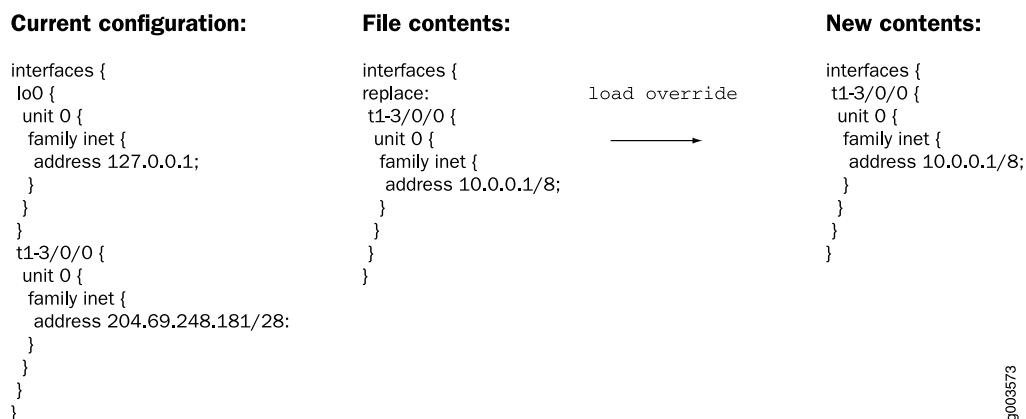
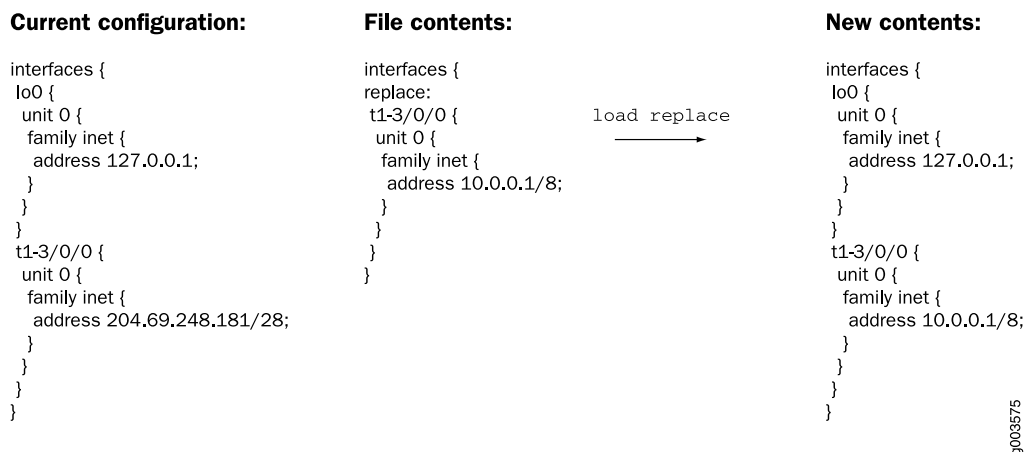
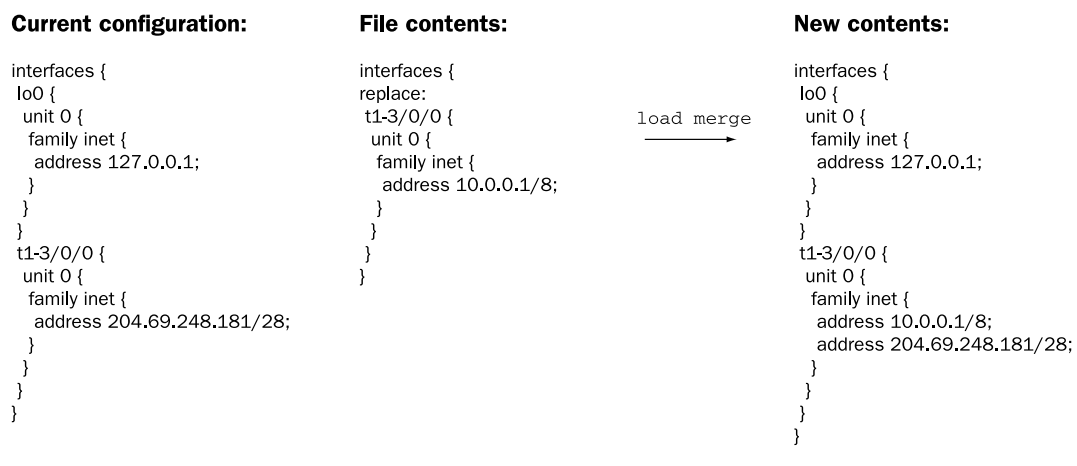
Figure 8: Loading a Configuration with the Override Operation

Figure 9: Loading a Configuration with the Replace Operation**Figure 10: Loading a Configuration with the Merge Operation**

Saving a Configuration File

To save your current configuration to an ASCII file, including any uncommitted changes made by you and all users, issue the **save** command:

save *filename*

By default, the configuration is saved to a file in your home directory. For information about specifying filenames, see the *JUNOS CLI User Guide*.

Part 2

Configuring Router Interfaces

- Interfaces Overview on page 41
- Configuring Ethernet, DS1, DS3, and Serial Interfaces on page 103
- Configuring Channelized T1/E1/ISDN PRI Interfaces on page 139
- Configuring Digital Subscriber Line Interfaces on page 157
- Configuring Point-to-Point Protocol over Ethernet on page 191
- Configuring ISDN on page 211
- Configuring Link Services Interfaces on page 261
- Configuring VoIP on page 311
- Configuring uPIMs as Ethernet Switches on page 343

Chapter 2

Interfaces Overview

J-series Services Routers support network interfaces for E1, E3, T1, T3, Fast Ethernet, Gigabit Ethernet, serial, Point-to-Point Protocol over Ethernet (PPPoE), and ISDN media. In addition, the router supports a set of special interfaces for such tasks as router identification and security services. Each type of interface has particular physical and logical characteristics.

To configure and monitor Services Router interfaces, you need to understand their media characteristics, as well as physical and logical properties such as IP addressing, link-layer protocols, and link encapsulation.

This chapter contains the following topics. For more information about interfaces, see the *JUNOS Network Interfaces Configuration Guide*, the *JUNOS Services Interfaces Configuration Guide*, and the *JUNOS Interfaces Command Reference*.

- Interfaces Terms on page 42
- Network Interfaces on page 46
- Data Link Layer Overview on page 52
- Ethernet Interface Overview on page 53
- T1 and E1 Interfaces Overview on page 57
- Channelized T1/E1/ISDN PRI Interfaces Overview on page 61
- T3 and E3 Interfaces Overview on page 61
- Serial Interface Overview on page 66
- ADSL Interface Overview on page 72
- SHDSL Interface Overview on page 74
- ISDN Interface Overview on page 75
- Interface Physical Properties on page 78
- Physical Encapsulation on an Interface on page 83
- Interface Logical Properties on page 90
- Special Interfaces on page 97

Interfaces Terms

To understand interfaces, become familiar with the terms defined in Table 17 on page 42.

Table 17: Network Interfaces Terms

Term	Definition
alternate mark inversion (AMI)	Original method of formatting T1 and E1 data streams.
asymmetric digital subscriber line (ADSL) interface	Physical WAN interface for connecting a Services Router to a digital subscriber line access multiplexer (DSLAM). An ADSL interface allocates line bandwidth asymmetrically with downstream (provider-to-customer) data rates of up to 8 Mbps for ADSL, 12 Mbps for ADSL2, and 25 Mbps for ADSL2 + , and upstream (customer-to-provider) rates of up to 800 Kbps for ADSL and 1 Mbps for ADSL2 and ADSL2 + , depending on the implementation.
ADSL2 interface	An ADSL interface that supports ITU-T Standards G.992.3 and G.992.4 and allocates downstream (provider-to-customer) data rates of up to 12 Mbps and upstream (customer-to-provider) rates of up to 1 Mbps.
ADSL2 + interface	An ADSL interface that supports ITU-T Standard G.992.5 and allocates downstream (provider-to-customer) data rates of up to 25 Mbps and upstream (customer-to-provider) rates of up to 1 Mbps.
Annex A	ITU-T Standard G.992.1 that defines how ADSL works over plain old telephone service (POTS) lines.
Annex B	ITU-T Standard G.992.1 that defines how ADSL works over Integrated Services Digital Network (ISDN) lines.
binary 8-zero substitution (B8ZS)	Improved method of formatting T1 and E1 data streams, in which a special code is substituted whenever 8 consecutive zeros are sent over the link.
Challenge Handshake Authentication Protocol (CHAP)	Protocol that authenticates remote users. CHAP is a server-driven, three-step authentication method that depends on a shared secret password residing on both the server and the client.
checksum	<i>See frame checksum sequence.</i>
channel group	Combination of DS0 interfaces partitioned from a channelized interface into a single logical bundle.
channel service unit (CSU)	Unit that connects a digital telephone line to a multiplexer or other signal service.
channelized E1	2.048-Mbps interface that can be configured as a single clear-channel E1 interface or channelized into as many as 31 discrete DS0 interfaces, or up to 30 ISDN PRI B-channels and 1 D-channel. On J-series channelized T1/E1/ISDN PRI interfaces, time slots are numbered from 1 through 31, and time slot 1 is reserved for framing. When the interface is configured for ISDN PRI service, time slot 16 is reserved for the D-channel.
channelized interface	Interface that is a subdivision of a larger interface, minimizing the number of Physical Interface Modules (PIMs) that an installation requires. On a channelized PIM, each port can be configured as a single clear channel or partitioned into multiple discrete T1, E1, and DS0 interfaces.

Table 17: Network Interfaces Terms (continued)

Term	Definition
channelized T1	1.544-Mbps interface that can be configured as a single clear-channel T1 interface or channelized into as many as 24 discrete DS0 interfaces, or up to 23 ISDN PRI B-channels and 1 D-channel. When the interface is configured for ISDN PRI service, time slot 24 is reserved for the D-channel.
Cisco HDLC	Cisco High-level Data Link Control protocol. Proprietary Cisco encapsulation for transmitting LAN protocols over a WAN. HDLC specifies a data encapsulation method on synchronous serial links by means of frame characters and checksums. Cisco HDLC enables the transmission of multiple protocols.
clock source	Source of the consistent, periodic signal used by a router to synchronize data communication and processing tasks.
CSU compatibility mode	Subrate on an E3 or T3 interface that allows a Services Router to connect to a channel service unit (CSU) with proprietary multiplexing at the remote end of the line. Subrating an E3 or T3 interface reduces the maximum allowable peak rate by limiting the payload encapsulated by the High-level Data Link Control protocol (HDLC).
data-link connection identifier (DLCI)	Identifier for a Frame Relay virtual connection, also called a logical interface.
data service unit (DSU)	Unit that connects a data terminal equipment (DTE) device—in this case, a Services Router—to a digital telephone line.
data terminal equipment (DTE)	RS-232 interface that a Services Router uses to exchange information with a serial device.
DS1	Digital signal 1, another name for a T1 interface.
DS3 interface	Digital signal 3, another name for a T3 interface.
data inversion	Transmission of all data bits in the data stream so that zeros are transmitted as ones and ones are transmitted as zeros. Data inversion is normally used only in alternate mark inversion (AMI) mode to guarantee ones density in the transmitted stream.
E1 interface	Physical WAN interface for transmitting signals in European digital transmission (E1) format. The E1 signal format carries information at a rate of 2.048 Mbps and can carry 32 channels of 64 Kbps each.
E3 interface	Physical WAN interface for transmitting 16 E1 circuits over copper wires using time-division multiplexing. E3 is widely used outside of North America and transfers traffic at the rate of 34.368 Mbps.
encapsulation type	Type of protocol header in which data is wrapped for transmission.
Fast Ethernet interface	Physical LAN interface for transmitting data at 100 Mbps. Fast Ethernet, also called 100Base-T, additionally supports standard 10Base-T Ethernet transmission. The two built-in ports on J2300, J4300, and J6300 Services Routers are Fast Ethernet interfaces. Fast Ethernet is also available in dual-port PIMs for these routers and in both dual-port and 4-port PIMs for the J4350 and J6350 Services Routers.
FPC	Logical identifier for a Physical Interface Module (PIM) installed on a Services Router. The FPC number used in the JUNOS command-line interface (CLI) and displayed in command output represents the chassis slot in which a PIM is installed.

Table 17: Network Interfaces Terms (continued)

Term	Definition
fractional E1	Interface that contains one or more of the 32 DS0 time slots that can be reserved from an E1 interface. (Time slot 0 is reserved.)
fractional T1	Interface that contains one or more of the 24 DS0 time slots that can be reserved from a T1 interface. (Time slot 0 is reserved.)
frame check sequence (FCS)	Calculation that is added to a frame to control errors in High-level Data Link Control (HDLC), Frame Relay, and other data link layer protocols.
Frame Relay	An efficient WAN protocol that does not require explicit acknowledgement of each frame of data. Frame Relay allows private networks to reduce costs by sharing facilities between the endpoint switches of a network managed by a Frame Relay service provider. Individual data link connection identifiers (DLCIs) are assigned to ensure that customers receive only their own traffic.
Gigabit Ethernet interface	Physical LAN or WAN interface for transmitting data at 1000 Mbps. The four built-in ports on J4350 and J6350 Services Routers are Gigabit Ethernet interfaces. Gigabit Ethernet is also available in a single-port copper or optical PIM for these routers.
High-Level Data Link Control (HDLC)	International Telecommunication Union (ITU) standard for a bit-oriented data link layer protocol on which most other bit-oriented protocols are based.
hostname	Name assigned to the Services Router during initial configuration.
ITU-T G.991.2	International Telecommunication Union standard describing a data transmission method for symmetric high-speed digital subscriber line (SHDSL) as a means for data transport in telecommunications access networks. The standard also describes the functionality required for interoperability of equipment from various manufacturers.
ITU-T G.992.1	International Telecommunication Union standard that requires the downstream (provider-to-customer) data transmission to consist of full-duplex low-speed bearer channels and simplex high-speed bearer channels. In the upstream (customer-to-provider) transmissions, only low-speed bearer channels are provided.
ITU-T G.994.1	International Telecommunication Union standard describing the types of signals, messages, and procedures exchanged between digital subscriber line (DSL) equipment when the operational modes of equipment need to be automatically established and selected.
ITU-T G.997.1	International Telecommunication Union standard describing the physical layer management for asymmetric digital subscriber line (ADSL) transmission systems. The standard specifies the means of communication on a transport transmission channel defined in the physical layer recommendations. In addition, the standard describes the content and syntax of network elements for configuration, fault management, and performance management.
logical interface	Virtual interface that you create on a physical interface to identify its connection. Creating multiple logical interfaces allows you to associate multiple virtual circuits, data line connections, or virtual LANs (VLANs) with a single interface device.
maximum transmission unit (MTU)	Maximum or largest segment size that a network can transmit.
Multilink Frame Relay (MLFR)	Protocol that allows multiple Frame Relay links to be aggregated by inverse multiplexing.

Table 17: Network Interfaces Terms (continued)

Term	Definition
Multilink Point-to-Point Protocol (MLPPP)	Protocol that allows you to bundle multiple Point-to-Point Protocol (PPP) links into a single logical unit. MLPPP improves bandwidth efficiency and fault tolerance and reduces latency.
Password Authentication Protocol (PAP)	Authentication protocol that uses a simple 2-way handshake to establish identity.
Physical Interface Module (PIM)	<p>Network interface card that is fixed or can be interchangeably installed on a Services Router to provide the physical connections to a LAN or WAN, receiving incoming packets and transmitting outgoing packets. A PIM contains <i>one</i> of the following interfaces or sets of interfaces:</p> <ul style="list-style-type: none"> ■ Single Gigabit Ethernet LAN or WAN interface (J4350 and J6350 models only) ■ Two or four Fast Ethernet LAN interfaces ■ Two T1 or two E1 WAN interfaces ■ Single E3 or T3 (DS3) WAN interface (J4350, J6300, and J6350 models only) ■ Single asynchronous digital subscriber line (ADSL) WAN interface—Annex A to support ADSL over plain old telephone service (POTS) lines or Annex B to support ADSL over ISDN (all J-series models except J2300) ■ Single ISDN S/T or U interface (J2300 model) or four ISDN S/T or U interfaces (all J-series models except J2300) ■ Two serial interfaces ■ Symmetric high-speed digital subscriber line (SHDSL) WAN interface—Annex A or Annex B to support ATM-over-SHDSL connections
Point-to-Point Protocol (PPP)	Link-layer protocol that provides multiprotocol encapsulation. PPP is used for link-layer and network-layer configuration.
serial interface	<p>Physical LAN interface for transmitting data between computing devices. A Services Router has two types of serial interfaces:</p> <ul style="list-style-type: none"> ■ Asynchronous serial interface—Console port, with speeds up to 110.5 Kbps. The console port supports an RS-232 (EIA-232) standard serial cable with a 25-pin (DB-25) connector. ■ Synchronous serial interface—Port that transmits packets to and from, for example, a T1 device or microwave link, at speeds up to 8 Mbps. You cannot use this serial interface to connect a console. Services Router synchronous serial interfaces support RS-232 (EIA-232), RS-422/449 (EIA-449), RS-530 (EIA-530), V.35, and X.21 cable types. For details, see “Serial Line Protocols” on page 69. <p>For cable details, see the Getting Started Guide for your router.</p>
symmetric high-speed digital subscriber line (G.SHDSL)	Physical WAN symmetric DSL interface capable of sending and receiving high-speed symmetrical data streams over a single pair of copper wires at rates between 192 Kbps and 2.31 Mbps. G.SHDSL incorporates features of other DSL technologies such as asymmetric DSL and transports T1, E1, ISDN, Asynchronous Transfer Mode (ATM), and IP signals.
symmetric high-speed digital subscriber line (SHDSL) transceiver unit-remote (STU-R)	Equipment that provides symmetric high-speed digital subscriber line (SHDSL) connections to remote user terminals such as data terminals or telecommunications equipment.

Table 17: Network Interfaces Terms (continued)

Term	Definition
T1 interface	Physical WAN interface for transmitting digital signals in the T-carrier system used in the United States, Japan, and Canada. The T1 signal format carries 24 pulse code modulation (PCM) signals using time-division multiplexing (TDM) at an overall rate of 1.544 Mbps.
T3 interface	Physical WAN interface for transmitting digital signals in the T-carrier system used in the United States, Japan, and Canada. T3 signals are formatted like T1 signals, but carry information at the higher rate of 44.736 Mbps. T3 is also called DS3.

Network Interfaces

Services Routers use network interfaces to make physical connections to other devices. A connection takes place along media-specific physical wires through a port on a Physical Interface Module (PIM) installed in the router. Each Services Router interface has a unique name that follows a naming convention.

This section contains the following topics:

- Media Types on page 46
- Network Interface Naming on page 47

Media Types

Each type of interface on a Services Router uses a particular medium to transmit data. The physical wires and data link layer protocols used by a medium determine how traffic is sent. Services Routers support the following media types:

- Asynchronous Transfer Mode over asymmetric digital subscriber line (ATM-over-ADSL) interface (all J-series models except J2300)



NOTE: Services Routers with ADSL PIMs can use PPP over Ethernet (PPPoE) and PPP over ATM (PPPoA) to connect through ADSL lines only, not for direct ATM connections.

- Asynchronous Transfer Mode over symmetrical high-speed digital subscriber line (ATM-over-SHDSL) interface



NOTE: Services Routers with SHDSL PIMs can connect through SHDSL lines only, not for direct ATM connections.

- Channelized E1 interface
- E1 interface
- E3 interface (J4350, J6300, and J6350 models only)
- Fast Ethernet interface

- Gigabit Ethernet interface (J4350 and J6350 models only)
- Integrated Services Digital Network (ISDN) BRI interface
- Serial interface (EIA-530, RS-449/422, RS-232, V.35, and X.21 line protocols)
- Channelized T1 interface
- T1 interface
- T3 interface (also called DS3) (J4350, J6300, and J6350 models only)

You must configure each network interface before it can operate on the router. Configuring an interface can define both the physical properties of the link and the logical properties of a logical interface on the link.

Network Interface Naming

The interfaces on a Services Router are used for networking and services. Most interfaces are configurable, but some internally generated interfaces are not configurable. If you are familiar with Juniper Networks M-series and T-series routing platforms, be aware that Services Router interface names are similar to but not identical with the interface names on those routing platforms.

This section contains the following topics:

- J-series Interface Naming Conventions on page 47
- Understanding CLI Output for J-series Interfaces on page 49

J-series Interface Naming Conventions

The unique name of each Services Router interface identifies its type and location and indicates whether it is a physical interface or an optional logical unit created on a physical interface:

- The name of each interface on the router has the following format, to identify the physical device that corresponds to a single physical network connector:

type-pim/0/port

- Network interfaces that are fractionalized into time slots include a channel number in the name, preceded by a colon (:):

type-pim/0/port:channel

- Each logical interface has an additional logical unit identifier, preceded by a period (.):

type-pim/0/port:<channel>.unit

The parts of an interface name are summarized in Table 18 on page 48.

Table 18: J-series Services Router Interface Names

Name Part	Meaning	Possible Values
<i>type</i>	Type of network medium that can connect to this interface.	<p>at—ATM-over-ADSL or ATM-over-SHDSL WAN interface</p> <p>bc—Bearer channel on an ISDN BRI</p> <p>br—Basic Rate Interface for establishing ISDN connections</p> <p>ce1—Channelized E1 interface</p> <p>ct1—Channelized T1 interface</p> <p>dc—Delta channel on an ISDN BRI</p> <p>dl—Dialer interface for initiating ISDN and USB modem connections</p> <p>e1—E1 WAN interface</p> <p>e3—E3 WAN interface</p> <p>fe—Fast Ethernet interface</p> <p>ge—Gigabit Ethernet interface</p> <p>se—Serial interface (either RS-232, RS-422/499, RS-530, V.35, or X.21)</p> <p>t1—T1 (also called DS1) WAN interface</p> <p>t3—T3 (also called DS3) WAN interface</p> <p>vp—VoIP interface</p> <p>In addition to these network types, J-series routers can have the following special interfaces: dsc, gr and gre, ip and ipip, lo, ls and lsi, lt, mt and mtun, pd and pimd, pe and pime, pp0, sp, tap, and umd0. For more information, see “Special Interfaces” on page 97.</p>
<i>pim</i>	Number of the chassis slot in which a PIM is installed.	<ul style="list-style-type: none"> ■ On a J2300 router, always 0. ■ On a J4300 or J6300 router, this number begins at 1 and increases from left to right, bottom to top to a maximum of 6. The PIM number 0 is reserved for the out-of-band management ports. (See “Management Interface” on page 101.) ■ On a J4350 or J6350 router, this number begins at 1 and increases from top to bottom, left to right to a maximum of 6. The PIM number 0 is reserved for the out-of-band management ports. (See “Management Interface” on page 101.)
0	Number of the PIM installed in a chassis slot.	<p>Always 0.</p> <p>Only one PIM can be installed in a slot.</p>

Table 18: J-series Services Router Interface Names (*continued*)

Name Part	Meaning	Possible Values
<i>port</i>	Number of the port on a PIM on which the physical interface is located.	<ul style="list-style-type: none"> ■ On a single-port PIM, always 0. ■ On a multiple-port PIM, this number begins at 0 and increases from left to right, bottom to top, to a maximum of 3. <p>Port numbers appear on the PIM faceplate.</p>
<i>channel</i>	Number of the channel (time slot) on a fractional or channelized T1 or E1 interface.	<ul style="list-style-type: none"> ■ On an E1 interface, a value from 1 through 31. The 1 time slot is reserved. ■ On a T1 interface, a value from 1 through 24.
<i>unit</i>	Number of the logical interface created on a physical interface.	<p>A value from 0 through 16384, except on a VoIP interface. A VoIP interface must have the logical interface number 0.</p> <p>If no logical interface number is specified, unit 0 is the default, but must be explicitly configured. For more information about logical interfaces, see “Interface Logical Properties” on page 90.</p>

For example, the interface name **e1-5/0/0:15.0** represents the following information:

- E1 WAN interface
- PIM slot 5
- PIM number 0 (always 0)
- Port 0
- Channel 15
- Logical interface, or unit, 0

Understanding CLI Output for J-series Interfaces

The JUNOS Internet software that operates J-series Services Routers was originally developed for Juniper Networks M-series and T-series routing platforms that support many ports, on interface cards called Physical Interface Cards (PICs). On these larger platforms, PICs are installed into slots on FPCs, and FPCs are installed into slots in the router chassis.

Because Services Routers have the same hardware and software architectures as the M-series and T-series routing platforms, PIM slots are detected internally by the JUNOS software as FPC slots, and the PIM in each slot is identified as a “PIC.” For example, in the following output, the three PIMs located in slots 0, 2, and 5 are reported as FPC 0, FPC 2, and FPC 5, and PIM 0 is reported as PIC 0:

```

user@host> show chassis hardware
Hardware inventory:
Item          Version  Part number  Serial number  Description
Chassis              REV 02.04  710-010001  JN000192AB    J4300
Midplane              CORE99563

```

System IO	REV 02.03	710-010003	CORE100885	System IO board
Routing Engine	RevX2.6	750-010005	IWGS40735451	RE-J.2
FPC 0				FPC
PIC 0				2x FE
FPC 2	RevX2.1	750-010355	CORE100458	FPC
PIC 0				2x T1
FPC 5	REV 04	750-010353	AF04451744	FPC
PIC 0				2x FE

To understand the abbreviations for PIMs, enhanced PIMs (ePIMs), and universal PIMs (uPIMs) that appear in JUNOS CLI output, see Table 19 on page 50 (for J2300 fixed PIMs) and Table 20 on page 50 (for PIMs in the other J-series models). For PIM details, see the Getting Started Guide for your router.

Table 19: J2300 Fixed PIM Abbreviations, Chassis Names, and Ports

PIM Abbreviation in JUNOS CLI	Chassis Name	Ports
2x FE, 2x E1	Dual-Port E1	2 Fast Ethernet ports and 2 E1 ports
2x FE, 2x E1, 1x BRI S/T	Dual-Port E1 with ISDN BRI S/T	2 Fast Ethernet ports, 2 E1 ports, and 1 ISDN BRI S/T port
2x FE, 2x T1	Dual-Port T1	2 Fast Ethernet ports and 2 T1 ports
2x FE, 2x T1, 1x BRI U	Dual-Port T1 with ISDN BRI U	2 Fast Ethernet ports, 2 T1 ports, and 1 ISDN BRI U port
2x FE, 2x Serial	Dual-Port Serial	2 Fast Ethernet ports and 2 serial ports
2x FE, 2x Serial, 1x BRI S/T	Dual-Port Serial with ISDN BRI S/T	2 Fast Ethernet ports, 2 serial ports, and 1 ISDN BRI S/T port
2x FE, 2x Serial, 1x BRI U	Dual-Port Serial with ISDN BRI U	2 Fast Ethernet ports 2 serial ports, and 1 ISDN BRI U port
2x FE, 2x SHDSL	Dual-Port G.SHDSL	2 Fast Ethernet ports and 2 two-wire G.SHDSL ports (in either 2-port two-wire mode or 1-port four-wire mode)
2x FE, 2x SHDSL, 1x BRI S/T	Dual-Port G.SHDSL with ISDN BRI S/T	2 Fast Ethernet ports, 2 two-wire G.SHDSL ports (in either 2-port two-wire mode or 1-port four-wire mode), and 1 ISDN BRI S/T port

Table 20: J4300, J4350, J6300, and J6350 PIM Abbreviations and Full Names

PIM Abbreviation in JUNOS CLI	PIM Name
2x FE	<ul style="list-style-type: none"> ■ 2 built-in Fast Ethernet ports on a J4300 or J6300 chassis (fixed PIM) ■ Dual-Port Fast Ethernet PIM
4x FE	4-Port Fast Ethernet ePIM

Table 20: J4300, J4350, J6300, and J6350 PIM Abbreviations and Full Names *(continued)*

PIM Abbreviation in JUNOS CLI	PIM Name
1x GE Copper	Copper Gigabit Ethernet ePIM (1 10-Mbps, 100-Mbps, or 1000-Mbps port)
1x GE SFP	SFP Gigabit Ethernet ePIM (1 fiber port)
1x SFP uPIM	1-Port Gigabit Ethernet uPIM
6x GE SFP uPIM	6-Port SFP Gigabit Ethernet uPIM
8x GE uPIM	8-Port Gigabit Ethernet uPIM
16x GE uPIM	16-Port Gigabit Ethernet uPIM
4x GE Base PIC	4 built-in Gigabit Ethernet ports on a J4350 or J6350 chassis (fixed PIM)
2x Serial	Dual-Port Serial PIM
2x T1	Dual-Port T1 PIM
2x E1	Dual-Port E1 PIM
2x CT1E1 /PRI	Dual-Port Channelized T1/E1/ISDN PRI PIM
1x T3	T3 PIM (1 port)
1x E3	E3 PIM (1 port)
4x BRI S/T	4-Port ISDN BRI S/T PIM
4x BRI U	4-Port ISDN BRI U PIM
1x ADSL Annex A	ADSL 2/2 + Annex A PIM (1 port, for POTS)
1x ADSL Annex B	ADSL 2/2 + Annex B PIM (1 port, for ISDN)
2x SHDSL (ATM)	G.SHDSL PIM (2-port two-wire mode or 1-port four-wire mode)
1x TGM550	TGM550 Telephony Gateway Module (Avaya VoIP Gateway Module with 1 console port, 2 analog LINE ports, and 2 analog TRUNK ports)
1x DS1 TIM510	TIM510 E1/T1 Telephony Interface Module (Avaya VoIP media module with 1 E1 or T1 trunk termination port)
4x FXS, 4xFXO TIM514	TIM514 Analog Telephony Interface Module (Avaya VoIP media module with 4 analog LINE ports and 4 analog TRUNK ports)
4x BRI TIM521	TIM521 BRI Telephony Interface Module (Avaya VoIP media module with 4 ISDN BRI ports)

Data Link Layer Overview

The data link layer is Layer 2 in the Open Systems Interconnection (OSI) model. The data link layer is responsible for transmitting data across a physical network link. Each physical medium has link-layer specifications for network and link-layer protocol characteristics such as physical addressing, network topology, error notification, frame sequencing, and flow control.

Physical Addressing

Physical addressing is different from network addressing. Network addresses differentiate between nodes or devices in a network, allowing traffic to be routed or switched through the network. In contrast, physical addressing identifies devices at the link-layer level, differentiating between individual devices on the same physical medium. The primary form of physical addressing is the media access control (MAC) address.

Network Topology

Network topology specifications identify how devices are linked in a network. Some media allow devices to be connected by a bus topology, while others require a ring topology. The bus topology is used by Ethernet technologies, which are supported on Services Routers.

Error Notification

The data link layer provides error notifications that alert higher-layer protocols that an error has occurred on the physical link. Examples of link-level errors include the loss of a signal, the loss of a clocking signal across serial connections, or the loss of the remote endpoint on a T1 or T3 link.

Frame Sequencing

The frame sequencing capabilities of the data link layer allow frames that are transmitted out of sequence to be reordered on the receiving end of a transmission. The integrity of the packet can then be verified by means of the bits in the Layer 2 header, which is transmitted along with the data payload.

Flow Control

Flow control within the data link layer allows receiving devices on a link to detect congestion and notify their upstream and downstream neighbors. The neighbor devices relay the congestion information to their higher-layer protocols so that the flow of traffic can be altered or rerouted.

Data Link Sublayers

The data link layer is divided into two sublayers: Logical Link Control (LLC) and Media Access Control (MAC). The LLC sublayer manages communications between devices

over a single link of a network. This sublayer supports fields in link-layer frames that enable multiple higher-layer protocols to share a single physical link.

The MAC sublayer governs protocol access to the physical network medium. Through the MAC addresses that are typically assigned to all ports on a router, multiple devices on the same physical link can uniquely identify one another at the data link layer. MAC addresses are used in addition to the network addresses that are typically configured manually on ports within a network.

MAC Addressing

A MAC address is the serial number permanently stored in a device adapter to uniquely identify the device. MAC addresses operate at the data link layer, while IP addresses operate at the network layer. The IP address of a device can change as the device is moved around a network to different IP subnets, but the MAC address remains the same, because it is physically tied to the device.

Within an IP network, devices match each MAC address to its corresponding configured IP address by means of the Address Resolution Protocol (ARP). ARP maintains a table with a mapping for each MAC address in the network.

Most Layer 2 networks use one of three primary numbering spaces—MAC-48, EUI-48 (Extended Unique Identifier), and EUI-64—which are all globally unique. MAC-48 and EUI-48 spaces each use 48-bit addresses, and EUI-64 spaces use a 64-bit addresses, but all three use the same numbering format. MAC-48 addresses identify network hardware, and EUI-48 addresses identify other devices and software.

The Ethernet and ATM technologies supported on Services Routers use the MAC-48 address space. IPv6 uses the EUI-64 address space.

MAC-48 addresses are the most commonly used MAC addresses in most networks. These addresses are 12-digit hexadecimal numbers (48 bits in length) that typically appear in one of the following formats:

- *MM:MM:MM:SS:SS:SS*
- *MM-MM-MM-SS-SS-SS*

The first three octets (*MM:MM:MM* or *MM-MM-MM*) are the ID number of the hardware manufacturer. Manufacturer ID numbers are assigned by the Institute of Electrical and Electronics Engineers (IEEE). The last three octets (*SS:SS:SS* or *SS-SS-SS*) make up the serial number for the device, which is assigned by the manufacturer. For example, an Ethernet interface card might have a MAC address of 00:05:85:c1:a6:a0.

Ethernet Interface Overview

Ethernet is a Layer 2 technology that operates in a shared bus topology. Ethernet supports broadcast transmission, uses best-effort delivery, and has distributed access control. Ethernet is a point-to-multipoint technology.

In a shared bus topology, all devices connect to a single, shared physical link through which all data transmissions are sent. All traffic is broadcast, so that all devices within

the topology receive every transmission. The devices within a single Ethernet topology make up a broadcast domain.

Ethernet uses best-effort delivery to broadcast traffic. The physical hardware provides no information to the sender about whether the traffic was received. If the receiving host is offline, traffic to the host is lost. Although the Ethernet data link protocol does not inform the sender about lost packets, higher-layer protocols like TCP/IP might provide this type of notification.

This section contains the following topics:

- Ethernet Access Control and Transmission on page 54
- Collisions and Detection on page 54
- Collision Domains and LAN Segments on page 55
- Broadcast Domains on page 56
- Ethernet Frames on page 56

Ethernet Access Control and Transmission

Ethernet's access control is distributed, because Ethernet has no central mechanism that grants access to the physical medium within the network. Instead, Ethernet uses carrier sense multiple access with collision detection (CSMA/CD). Because multiple devices on an Ethernet network can access the physical medium, or wire, simultaneously, each device must determine whether the physical medium is in use. Each host listens on the wire to determine if a message is being transmitted. If it detects no transmission, the host begins transmitting its own data.

The length of each transmission is determined by fixed Ethernet packet sizes. By fixing the length of each transmission and enforcing a minimum idle time between transmissions, Ethernet ensures that no pair of communicating devices on the network can monopolize the wire and block others from sending and receiving traffic.

Collisions and Detection

When a device on an Ethernet network begins transmitting data, the data takes a finite amount of time to reach all hosts on the network. Because of this delay, or latency, in transmitting traffic, a device might detect an idle state on the wire just as another device initially begins its transmission. As a result, two devices might send traffic across a single wire at the same time. When the two electrical signals collide, they become scrambled so that both transmissions are effectively lost.

Collision Detection

To handle collisions, Ethernet devices monitor the link while they are transmitting data. The monitoring process is known as collision detection. If a device detects a foreign signal while it is transmitting, it terminates the transmission and attempts to transmit again only after detecting an idle state on the wire. Collisions continue to occur if two colliding devices both wait the same amount of time before retransmitting. To avoid this condition, Ethernet devices use a binary exponential backoff algorithm.

Backoff Algorithm

To use the binary exponential backoff algorithm, each device that sent a colliding transmission randomly selects a value within a range. The value represents the number of transmission times that the device must wait before retransmitting its data. If another collision occurs, the range of values is doubled and retransmission takes place again. Each time a collision occurs, the range of values doubles, to reduce the likelihood that two hosts on the same network can select the same retransmission time. Table 21 on page 55 shows collision rounds up to round 10.

Table 21: Collision Backoff Algorithm Rounds

Round	Size of Set	Elements in the Set
1	2	{0,1}
2	4	{0,1,2,3}
3	8	{0,1,2,3,...,7}
4	16	{0,1,2,3,4,...,15}
5	32	{0,1,2,3,4,5,...,31}
6	64	{0,1,2,3,4,5,6,...,63}
7	128	{0,1,2,3,4,5,6,7,...,127}
8	256	{0,1,2,3,4,5,6,7,8,...,255}
9	512	{0,1,2,3,4,5,6,7,8,9,...,511}
10	1024	{0,1,2,3,4,5,6,7,8,9,10,...,1023}

Collision Domains and LAN Segments

Collisions are confined to a physical wire over which data is broadcast. Because the physical wires are subject to signal collisions, individual LAN segments are known as collision domains. Although the physical limitations on the length of an Ethernet cable restrict the length of a LAN segment, multiple collision domains can be interconnected by repeaters, bridges, and switches.

Repeaters

Repeaters are electronic devices that act on analog signals. Repeaters relay all electronic signals from one wire to another. A single repeater can double the distance between two devices on an Ethernet network. However, the Ethernet specification restricts the number of repeaters between any two devices on an Ethernet network to two, because collision detection with latencies increases in complexity as the wire length and number of repeaters increase.

Bridges and Switches

Bridges and switches combine LAN segments into a single Ethernet network by using multiple ports to connect the physical wires in each segment. Although bridges and switches are fundamentally the same, bridges generally provide more management and more interface ports. As Ethernet packets flow through a bridge, the bridge tracks the source MAC address of the packets and stores the addresses and their associated input ports in an interface table. As it receives subsequent packets, the bridge examines its interface table and takes one of the following actions:

- If the destination address does not match an address in the interface table, the bridge transmits the packet to all hosts on the network using the Ethernet broadcast address.
- If the destination address maps to the port through which the packet was received, the bridge or switch discards the packet. Because the other devices on the LAN segment also received the packet, the bridge does not need to retransmit it.
- If the destination address maps to a port other than the one through which the packet was received, the bridge transmits the packet through the appropriate port to the corresponding LAN segment.


Broadcast Domains

The combination of all the LAN segments within an Ethernet network is called a broadcast domain. In the absence of any signaling devices such as a repeater, bridge, or switch, the broadcast domain is simply the physical wire that makes up the connections in the network. If a bridge or switch is used, the broadcast domain consists of the entire LAN.

Ethernet Frames

Data is transmitted through an Ethernet network in frames. The frames are of variable length, ranging from 64 octets to 1518 octets, including the header, payload, and cyclic redundancy check (CRC) value. Figure 11 on page 56 shows the Ethernet frame format.

Figure 11: Ethernet Frame Format



Ethernet frames have the following fields:

- The preamble (PRE) in the frame is 7 octets of alternating 0s and 1s. The predictable format in the preamble allows receiving interfaces to synchronize themselves to the data being sent. The preamble is followed by a 1-octet start-of-frame delimiter (SFD).
- The destination address (DA) and source address (SA) fields contain the 6-octet (48-bit) MAC addresses for the destination and source ports on the network. These Layer 2 addresses uniquely identify the devices on the LAN.
- The length/type field is a 2-octet field that either indicates the length of the frame's data field or identifies the protocol stack associated with the frame. Following are some common frame types:
 - AppleTalk—0x809B
 - AppleTalk ARP—0x80F3
 - DECnet—0x6003
 - IP—0x0800
 - IPX—0x8137
 - Loopback—0x9000
 - XNS—0x0600
- The frame data is the packet payload.
- The frame check sequence (FCS) field is a 4-octet field that contains the calculated CRC value. This value is calculated by the originating host and appended to the frame. When it receives the frames, the receiving host calculates the CRC and checks it against this appended value to verify the integrity of the received frame.

T1 and E1 Interfaces Overview

T1 and E1 are equivalent digital data transmission formats that carry DS1 signals. T1 and E1 lines can be interconnected for international use. This section contains the following topics:

- T1 Overview on page 57
- E1 Overview on page 58
- T1 and E1 Signals on page 58
- Encoding on page 58
- T1 and E1 Framing on page 59
- T1 and E1 Loopback Signals on page 60

T1 Overview

T1 is a digital data transmission medium capable of handling 24 simultaneous connections running at a combined 1.544 Mbps. T1 combines these 24 separate connections, called channels or time slots, onto a single link. T1 is also called DS1.

The T1 data stream is broken into frames. Each frame consists of a single framing bit and 24 8-bit channels, totalling 193 bits per T1 frame. Frames are transmitted 8,000 times per second, at a data transmission rate of 1.544 Mbps ($8,000 \times 193 = 1.544$ Mbps).

As each frame is received and processed, the data in each 8-bit channel is maintained with the channel data from previous frames, enabling T1 traffic to be separated into 24 separate flows across a single medium. For example, in the following set of 4-channel frames (without a framing bit), the data in channel 1 consists of the first octet of each frame, the data in channel 2 consists of the second octet of each frame, and so on:

	Chan. 1	Chan. 2	Chan. 3	Chan. 4
Frame 1	[10001100]	[00110001]	[11111000]	[10101010]
Frame 2	[11100101]	[01110110]	[10001000]	[11001010]
Frame 3	[00010100]	[00101111]	[11000001]	[00000001]

E1 Overview

E1 is the European format for DS1 digital transmission. E1 links are similar to T1 links except that they carry signals at 2.048 Mbps. Each signal has 32 channels, and each channel transmits at 64 Kbps. E1 links have higher bandwidth than T1 links because they use all 8 bits of a channel. T1 links use 1 bit in each channel for overhead.

T1 and E1 Signals

T1 and E1 interfaces consist of two pairs of wires—a transmit data pair and a receive data pair. Clock signals, which determine when the transmitted data is sampled, are embedded in T1 and E1 transmissions.

Typical digital signals operate by sending either zeros (0s) or ones (1s), which are usually represented by the absence or presence of a voltage on the line. The receiving device need only detect the presence of the voltage on the line at the particular sampling edge to determine if the signal is 0 or 1. T1 and E1, however, use bipolar electrical pulses. Signals are represented by no voltage (0), positive voltage (1), or negative voltage (1). The bipolar signal allows T1 and E1 receivers to detect error conditions in the line, depending on the type of encoding that is being used. For more information, see “Encoding” on page 58.

Encoding

Following are common T1 and E1 encoding techniques:

- Alternate mark inversion (AMI)—T1 and E1
- Bipolar with 8-zero substitution (B8ZS)—T1 only
- High-density bipolar 3 code (HDB3)—E1 only

AMI Encoding

AMI encoding forces the 1s signals on a T1 or E1 line to alternate between positive and negative voltages for each successive 1 transmission, as in this sample data transmission:

```
1 1 0 1 0 1 0 1
+ - 0 + 0 - 0 +
```

When AMI encoding is used, a data transmission with a long sequence of 0s has no voltage transitions on the line. In this situation, devices have difficulty maintaining clock synchronization, because they rely on the voltage fluctuations to constantly synchronize with the transmitting clock. To counter this effect, the number of consecutive 0s in a data stream is restricted to 15. This restriction is called the 1s density requirement, because it requires a certain number of 1s for every 15 0s that are transmitted.

On an AMI-encoded line, two consecutive pulses of the same polarity—either positive or negative—are called a bipolar violation (BPV), which is generally flagged as an error.

B8ZS and HDB3 Encoding

Both B8ZS and HDB3 encoding do not restrict the number of 0s that can be transmitted on a line. Instead, these encoding methods detect sequences of 0s and substitute bit patterns in their place to provide the signal oscillations required to maintain timing on the link.

The B8ZS encoding method for T1 lines detects sequences of eight consecutive 0 transmissions and substitutes a pattern of two consecutive BPVs (11110000). Because the receiving end uses the same encoding, it detects the BPVs as 0s substitutions, and no BPV error is flagged. A single BPV, which does not match the 11110000 substitution bit sequence, is likely to generate an error, depending on the configuration of the device.

The HDB3 encoding method for E1 lines detects sequences of four consecutive 0 transmissions and substitutes a single BPV (1100). Similar to B8ZS encoding, the receiving device detects the 0s substitutions and does not generate a BPV error.

T1 and E1 Framing

Services Router T1 interfaces use two types of framing: superframe (D4) and extended superframe (ESF). E1 interfaces use G.704 framing or G.704 with no CRC4 framing, or can be in unframed mode.

Superframe (D4) Framing for T1

A D4 frame consists of 192 data bits: 24 8-bit channels and a single framing bit. The single framing bit is part of a 12-bit framing sequence. The 193rd bit in each T1 frame is set to a value, and every 12 consecutive frames are examined to determine the framing bit pattern for the 12-bit superframe.

The following sample 12-frame sequence shows the framing pattern for D4 framing:

```
[data bits][framing bit]
[xxxxxxxx][1]
[xxxxxxxx][0]
[xxxxxxxx][0]
[xxxxxxxx][0]
[xxxxxxxx][1]
[xxxxxxxx][1]
[xxxxxxxx][0]
[xxxxxxxx][1]
[xxxxxxxx][1]
[xxxxxxxx][1]
[xxxxxxxx][0]
[xxxxxxxx][0]
```

The 100011011100 12-bit pattern is repeated in each successive superframe. The receiving device detects these bits to synchronize with the incoming data stream and determine when the framing pattern begins and ends.

D4 framing requires the 8th bit of every byte (of every channel) within the frame to be set to 1, a process known as bit robbing. The bit-robbing requirement ensures that the 1s density requirements are met, regardless of the data contents of the frames, but it reduces the bandwidth on the T1 link by an eighth.

Extended Superframe (ESF) Framing for T1

ESF extends the D4 superframe from 12 frames to 24 frames. By expanding the size of the superframe, ESF increases the number of bits in the superframe framing pattern from 12 to 24. The extra bits are used for frame synchronization, error detection, and maintenance communications through the facilities data link (FDL).

The ESF pattern for synchronization bits is 001011. Only the framing bits from frames 4, 8, 12, 16, 20, and 24 in the superframe sequence are used to create the synchronization pattern.

The framing bits from frames 2, 6, 10, 14, 18, and 22 are used to pass a CRC code for each superframe block. The CRC code verifies the integrity of the received superframe and detects bit errors with a CRC6 algorithm.

The framing bits for frames 1, 3, 5, 7, 9, 11, 13, 15, 17, 19, 21, and 23 are used for the data link channel. These 12 bits enable the operators at the network control center to query the remote equipment for information about the performance of the link.

T1 and E1 Loopback Signals

The control signal on a T1 or E1 link is the loopback signal. Using the loopback signal, the operators at the network control center can force the device at the remote end of a link to retransmit its received signals back onto the transmit path. The transmitting device can then verify that the received signals match the transmitted signals, to perform end-to-end checking on the link.

Two loopback signals are used to perform the end-to-end testing:

- The loop-up command signal sets the link into loopback mode, with the following command pattern:

....100001000010000100...

- The loop-down signal returns the link to its normal mode, with the following command pattern:

....100100100100100100....

While the link is in loopback mode, the operator can insert test equipment onto the line to test its operation.

Channelized T1/E1/ISDN PRI Interfaces Overview

Channelization enables routers to provide IP services to users with different access speeds and bandwidth requirements. Users share an interface that has been divided into discrete time slots, by transmitting in only their own time slot. On J-series Services Routers, a single channelized T1/E1/ISDN PRI interface can be partitioned into the following numbers of DS0 or ISDN PRI time slots, by means of software configuration:

- T1 interface—Up to 24 DS0 time slots (channels 1 through 24).
- E1 interface—Up to 31 DS0 time slots (channels 1 through 31).
- ISDN PRI—Up to 23 ISDN PRI B-channels and 1 D-channel when the parent interface is channelized T1, and up to 30 ISDN PRI B-channels and 1 D channel when the parent interface is channelized E1. Time slots on the interface unused by ISDN PRI can operate normally as DS0 interfaces.

For more information about ISDN, see “ISDN Interface Overview” on page 75.



NOTE: You cannot configure the channelized T1/E1/ISDN PRI PIM through a J-Web Quick Configuration page.

You can aggregate the channels on a channelized interface into bundles called channel groups to aggregate customer traffic.

A single channelized T1/E1/ISDN PRI interface also supports drop-and-insert multiplexing, to integrate voice and data channels on a single T1 or E1 link. The drop-and-insert feature allows you to remove the DS0 time slots of one T1 or E1 port and replace them by inserting the time slots of another T1 or E1 interface.

T3 and E3 Interfaces Overview

T3 is a high-speed data-transmission medium formed by multiplexing 28 DS1 signals into seven separate DS2 signals, and combining the DS2 signals into a single DS3 signal. T3 links operate at 43.736 Mbps. T3 is also called DS3.

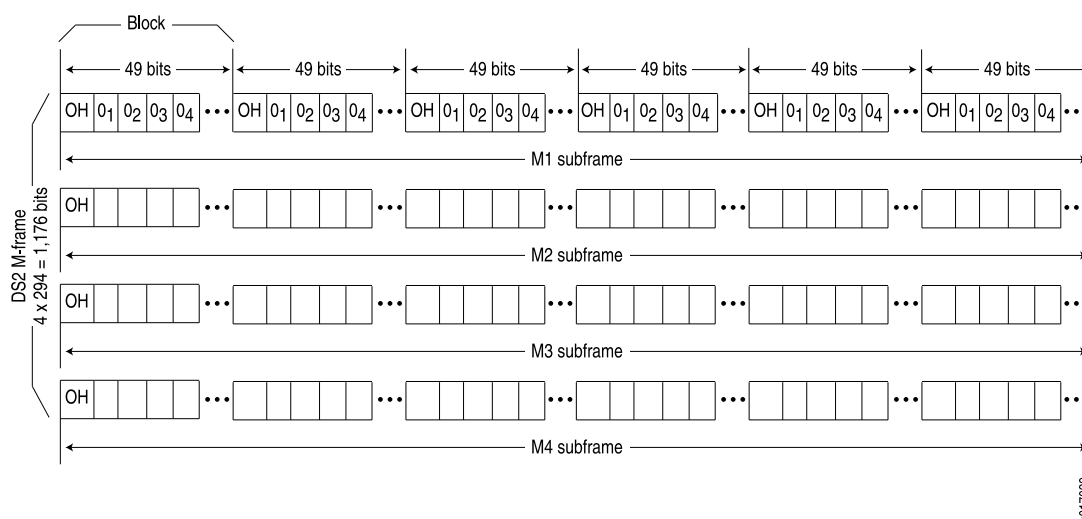
E3 is the equivalent European transmission format. E3 links are similar to T3 (DS3) links, but carry signals at 34.368 Mbps. Each signal has 16 E1 channels, and each channel transmits at 2.048 Mbps. E3 links use all 8 bits of a channel, whereas T3 links use 1 bit in each channel for overhead.

Multiplexing DS1 Signals

Four DS1 signals combine to form a single DS2 signal. The four DS1 signals form a single DS2 M-frame, which includes subframes M1 through M4. Each subframe has six 49-bit blocks, for a total of 294 bits per subframe. The first bit in each block is a DS2 overhead (OH) bit. The remaining 48 bits are DS1 information bits.

Figure 12 on page 62 shows the DS2 M-frame format.

Figure 12: DS2 M-Frame Format



The four DS2 subframes are not four DS1 channels. Instead, the DS1 data bits within the subframes are formed by data interleaved from the DS1 channels. The 0_n values designate time slots devoted to DS1 inputs as part of the bit-by-bit interleaving process. After every 48 DS1 information bits (12 bits from each signal), a DS2 OH bit is inserted to indicate the start of a subframe.

DS2 Bit Stuffing

Because the four DS1 signals are asynchronous signals, they might operate at different line rates. To synchronize the asynchronous streams, the multiplexers on the line use bit stuffing.

A DS2 connection requires a nominal transmit rate of 6.304 Mbps. However, because multiplexers increase the overall output rate to the intermediate rate of 6.312 Mbps, the output rate is higher than individual input rates on DS1 signals. The extra bandwidth is used to stuff the incoming DS1 signals with extra bits until the output rate of each signal equals the increased intermediate rate. These stuffed bits are

inserted at fixed locations in the DS2 M-frame. When DS2 frames are received and the signal is demultiplexed, the stuffing bits are identified and removed.


DS3 Framing

A set of four DS1 signals is multiplexed into seven DS2 signals, which are multiplexed into a single DS3 signal. The multiplexing occurs just as with DS1-to-DS2 multiplexing. The resulting DS3 signal uses either the standard M13 asynchronous framing format or the C-bit parity framing format. Although the two framing formats differ in their use of control and message bits, the basic frame structures are identical. The DS3 frame structures are shown in Figure 13 on page 63 and Figure 14 on page 65.

M13 Asynchronous Framing

A DS3 M-frame includes seven subframes, formed by DS2 data bits interleaved from the seven multiplexed DS2 signals. Each subframe has eight 85-bit blocks—a DS3 OH bit plus 84 data bits. The meaning of an OH bit depends on the block it precedes. Standard DS3 M13 asynchronous framing format is shown in Figure 13 on page 63.

Figure 13: DS3 M13 Frame Format



A DS3 M13 M-frame contains the following types of OH bits:

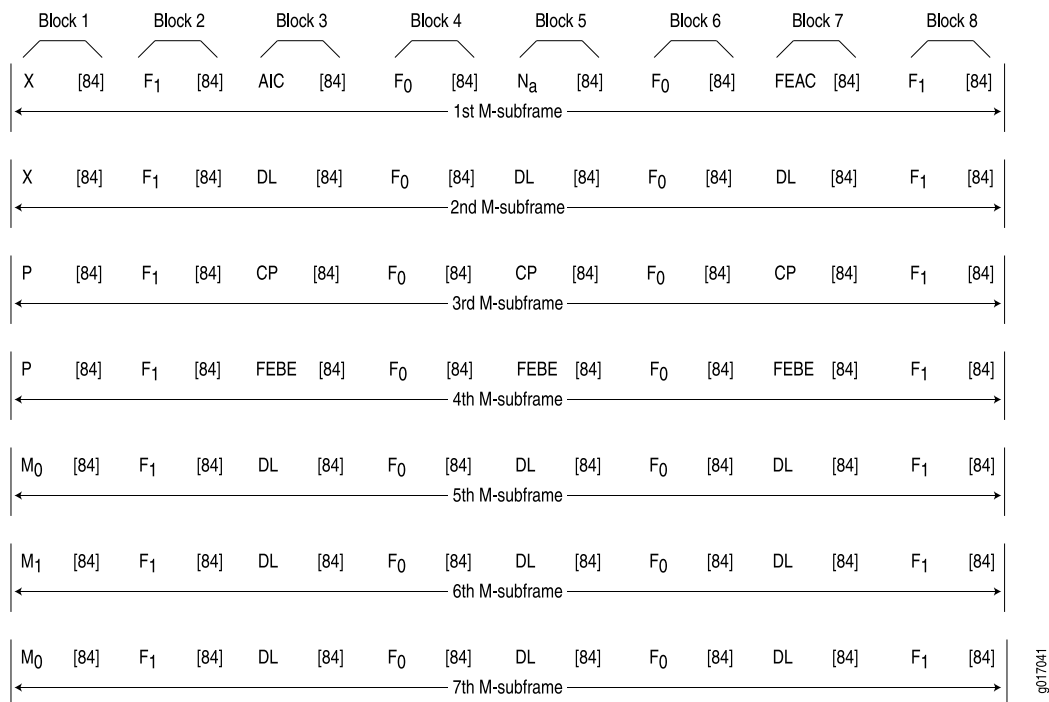
- Framing bits (F-bits)—Make up a frame alignment signal that synchronizes DS3 subframes. Each DS3 frame contains 28 F-bits (4 bits per subframe). F-bits are located at the beginning of blocks 2, 4, 6, and 8 of each subframe. When combined, the frame alignment pattern for each subframe is 1001. The pattern can be examined to detect bit errors in the transmission.
- Multiframe bits (M-bits)—Make up a multiframe alignment signal that synchronizes the M-frames in a DS3 signal. Each DS3 frame contains 3 M-bits, which are located at the beginning of subframes 5, 6, and 7. When combined, the multiframe alignment pattern for each M-frame is 010.
- Bit stuffing control bits (C-bits)—Serve as bit stuffing indicators for each DS2 input. For example, C_{11} , C_{12} , and C_{13} are indicators for DS2 input 1. Their values indicate whether DS3 bit stuffing has occurred at the multiplexer. If the three C-bits in a subframe are all 0s, no stuffing was performed for the DS2 input. If the three C-bits are all 1s, stuffing was performed.
- Message bits (X-bits)—Used by DS3 transmitters to embed asynchronous in-service messages in the data transmission. Each DS3 frame contains 2 X-bits, which are located at the beginning of subframes 1 and 2. Within an DS3 M-frame, both X-bits must be identical.
- Parity bits (P-bits)—Compute parity over all but 1 bit of the M-frame. (The first X-bit is not included.) Each DS3 frame contains 2 P-bits, which are located at the beginning of subframes 3 and 4. Both P-bits must be identical.

If the previous DS3 frame contained an odd number of 1s, both P-bits are set to 1. If the previous DS3 contained an even number of 1s, both P-bits are set to 0. If, on the receiving side, the number of 1s for a given frame does not match the P-bits in the following frame, it indicates one or more bit errors in the transmission.

C-Bit Parity Framing

In M13 framing, every C-bit in a DS3 frame is used for bit stuffing. However, because multiplexers first use bit stuffing when multiplexing DS1 signals into DS2 signals, the incoming DS2 signals are already synchronized. Therefore, the bit stuffing that occurs when DS2 signals are multiplexed is redundant.

C-bit parity framing format redefines the function of C-bits and X-bits, using them to monitor end-to-end path performance and provide in-band data links. The C-bit parity framing structure is shown in Figure 14 on page 65.

Figure 14: DS3 C-Bit Parity Framing

In C-bit parity framing, the X-bits transmit error conditions from the far end of the link to the near end. If no error conditions exist, both X-bits are set to 1. If an out-of-frame (OOF) or alarm indication signal (AIS) error is detected, both X-bits are set to 0 in the upstream direction for 1 second to notify the other end of the link about the condition.

The C-bits that control bit stuffing in M13 frames are typically used in the following ways by C-bit parity framing:

- Application identification channel (AIC)—The first C-bit in the first subframe identifies the type of DS3 framing used. A value of 1 indicates that C-bit parity framing is in use.
- N_a—A reserved network application bit.
- Far-end alarm and control (FEAC) channel—The third C-bit in the first subframe is used for the FEAC channel. In normal transmissions, the FEAC C-bit transmits all 1s. When an alarm condition is present, the FEAC C-bit transmits a code word in the format 0xxxxxx 11111111, in which x can be either 1 or 0. Bits are transmitted from right to left.

Table 22 on page 65 lists some C-bit code words and the alarm or status condition indicated.

Table 22: FEAC C-Bit Condition Indicators

Alarm or Status Condition	C-Bit Code Word
DS3 equipment failure requires immediate attention.	00110010 11111111

Table 22: FEAC C-Bit Condition Indicators (continued)

Alarm or Status Condition	C-Bit Code Word
DS3 equipment failure occurred—such as suspended, not activated, or unavailable service—that is non-service-affecting.	00011110 11111111
DS3 loss of signal.	00011100 11111111
DS3 out of frame.	00000000 11111111
DS3 alarm indication signal (AIS) received.	00101100 11111111
DS3 idle received.	00110100 11111111
Common equipment failure occurred that is non-service-affecting.	00011101 11111111
Multiple DS1 loss of signal.	00101010 11111111
DS1 equipment failure occurred that requires immediate attention.	00001010 11111111
DS1 equipment failure occurred that is non-service-affecting.	00000110 11111111
Single DS1 loss of signal.	00111100 11111111

- Data links—The 12 C-bits in subframes 2, 5, 6, and 7 are data link (DL) bits for applications and terminal-to-terminal path maintenance.
- DS3 parity—The 3 C-bits in the third subframe are DS3 parity C-bits (also called CP-bits). When a DS3 frame is transmitted, the sending device sets the CP-bits to the same value as the P-bits. When the receiving device processes the frame, it calculates the parity of the M-frame and compares this value to the parity in the CP-bits of the following M-frame. If no bit errors have occurred, the two values are typically the same.
- Far-end block errors (FEBEs)—The 3 C-bits in the fourth subframe make up the far-end block error (FEBE) bits. If a framing or parity error is detected in an incoming M-frame (via the CP-bits), the receiving device generates a C-bit parity error and sends an error notification to the transmitting (far-end) device. If an error is generated, the FEBE bits are set to 000. If no error occurred, the bits are set to 111.

Serial Interface Overview

Serial links are simple, bidirectional links that require very few control signals. In a basic serial setup, data communications equipment (DCE) installed in a user's premises is responsible for establishing, maintaining, and terminating a connection. A modem is a typical DCE device.

A serial cable connects the DCE to a telephony network where, ultimately, a link is established with data terminal equipment (DTE). DTE is typically where a serial link terminates.

The distinction between DCE and DTE is important because it affects the cable pinouts on a serial cable. A DTE cable uses a male 9-pin or 25-pin connector, and a DCE cable uses a female 9-pin or 25-pin connector.

To form a serial link, the cables are connected to each other. However, if the pins are identical, each side's transmit and receive lines are connected, which makes data transport impossible. To address this problem, each cable is connected to a null modem cable, which crosses the transmit and receive lines in the cable.

Serial Transmissions

In basic serial communications, nine signals are critical to the transmission. Each signal is associated with a pin in either the 9-pin or 25-pin connector. Table 23 on page 67 lists and defines serial signals and their sources.

Table 23: Serial Transmission Signals

Signal Name	Definition	Signal Source
TD	Transmitted data	DTE
RD	Received data	DCE
RTS	Request to send	DTE
CTS	Clear to send	DCE
DSR	Data set ready	DCE
Signal Ground	Grounding signal	–
CD	Carrier detect	–
DTR	Data terminal ready	DTE
RI	Ring indicator	–

When a serial connection is made, a serial line protocol—such as EIA-530, X.21, RS-422/449, RS-232, or V.35—begins controlling the transmission of signals across the line as follows:

1. The DCE transmits a DSR signal to the DTE, which responds with a DTR signal. After this handshake, the link is established and traffic can pass.
2. When the DTE device is ready to receive data, it sets its RTS signal to a marked state (all 1s) to indicate to the DCE that it can transmit data. (If the DTE is not able to receive data—because of buffer conditions, for example—it sets the RTS signal to all 0s.)

3. When the DCE device is ready to receive data, it sets its CTS signal to a marked state to indicate to the DTE that it can transmit data. (If the DCE is not able to receive data, it sets the CTS signal to all 0s.)
4. When the negotiation to send information has taken place, data is transmitted across the transmitted data (TD) and received data (RD) lines:
 - TD line—Line through which data from a DTE device is transmitted to a DCE device
 - RD line—Line through which data from a DCE device is transmitted to a DTE device

The name of the wire does not indicate the direction of data flow.

The DTR and DSR signals were originally designed to operate as a handshake mechanism. When a serial port is opened, the DTE device sets its DTR signal to a marked state. Similarly, the DCE sets its DSR signal to a marked state. However, because of the negotiation that takes place with the RTS and CTS signals, the DTR and DSR signals are not commonly used.

The carrier detect and ring indicator signals are used to detect connections with remote modems. These signals are not commonly used.

Signal Polarity

Serial interfaces use a balanced (also called differential) protocol signaling technique. Two serial signals are associated with a circuit: the A signal and the B signal. The A signal is denoted with a plus sign (for example, DTR+), and the B signal is denoted with a minus sign (for example, 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, but sometimes they might be reversed. For example, signals might be miswired as a result of reversed polarities.

Serial Clocking Modes


By default, a serial interface uses loop clocking to determine its timing source. For EIA-530 and V.35 interfaces, you can set each port independently to use one of the following clocking modes. X.21 interfaces can use only loop clocking mode.

- Loop clocking mode—Uses the DCE's receive (RX) clock to clock data from the DCE to the DTE.
- DCE clocking mode—Uses the transmit (TXC) clock, generated by the DCE specifically to be used by the DTE as the DTE's transmit clock.
- Internal clocking mode—Uses an internally generated clock. The speed of this clock is configured locally. Internal clocking mode is also known as line timing.

Both loop clocking mode and DCE clocking mode use external clocks generated by the DCE.

Figure 15 on page 69 shows the clock sources for loop, DCE, and internal clocking modes.

Figure 15: Serial Interface Clocking Modes



Serial Interface Transmit Clock Inversion

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.

DTE Clock Rate Reduction

Although the serial interface is intended for use at the default clock rate of 16.384 MHz, you might need to use a slower rate under any of the following conditions:

- 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.

The voltage must be measured differentially between the signal conductor and the point in the circuit from which all voltages are measured (“circuit common”) at the load end of the cable, with a 50-ohm resistor substituted for the generator.

- Interference with other signals must be minimized.
- Signals must be inverted.

Serial Line Protocols

Serial interfaces support the following line protocols:

- EIA-530 on page 70
- RS-232 on page 70
- RS-422/449 on page 71

- V.35 on page 71
- X.21 on page 72

EIA-530

EIA-530 is an Electronic Industries Association (EIA) standard for the interconnection of DTE and DCE using serial binary data interchange with control information exchanged on separate control circuits. EIA-530 is also known as RS-530.

The EIA-530 line protocol is a specification for a serial interface that uses a DB-25 connector and balanced equivalents of the RS-232 signals—also called V.24. The EIA-530 line protocol is equivalent to the RS-422 and RS-423 interfaces implemented on a 25-pin connector.

The EIA-530 line protocol supports both balanced and unbalanced modes. In unbalanced transmissions, voltages are transmitted over a single wire. Because only a single signal is transmitted, differences in ground potential can cause fluctuations in the measured voltage across the link. For example, if a 3V signal is sent from one endpoint to another, and the receiving endpoint has a ground potential 1V higher than the transmitter, the signal on the receiving end is measured as a 2V signal.

Balanced transmissions use two wires instead of one. Rather than sending a single signal across the wire and having the receiving end measure the voltage, the transmitting device sends two separate signals across two separate wires. The receiving device measures the difference in voltage of the two signals (balanced sampling) and uses that calculation to evaluate the signal. Any differences in ground potential affect both wires equally, and the difference in the signals is still the same.

The EIA-530 interface supports asynchronous and synchronous transmissions at rates ranging from 20 Kbps to 2 Mbps.

RS-232

RS-232 is a Recommended Standard (RS) describing the most widely used type of serial communication. The RS-232 protocol is used for asynchronous data transfer as well as synchronous transfers using HDLC, Frame Relay, and X.25. RS-232 is also known as EIA-232.

The RS-232 line protocol is very popular for low-speed data signals. RS-232 signals are carried as single voltages referred to a common ground signal. The voltage output level of these signals varies between -12V and $+12\text{V}$. Within this range, voltages between -3V and $+3\text{V}$ are considered inoperative and are used to absorb line noise. Control signals are considered operative when the voltage ranges from $+3$ to $+25\text{V}$.

The RS-232 line protocol is an unbalanced protocol, because it uses only one wire, and is susceptible to signal degradation. Degradation can be extremely disruptive, particularly when a difference in ground potential exists between the transmitting and receiving ends of a link.

The RS-232 interface is implemented in a 25-pin D-shell connector and supports line rates up to 200 Kbps over lines shorter than 98 feet (30 meters).



NOTE: RS-232 serial interfaces cannot function error-free with a clock rate greater than 200 KHz.

RS-422/449

RS-422 is a Recommended Standard (RS) describing the electrical characteristics of balanced voltage digital interface circuits that support higher bandwidths than traditional serial protocols like RS-232. RS-422 is also known as EIA-422.

The RS-449 standard (also known as EIA-449) is compatible with RS-422 signal levels. The EIA created RS-449 to detail the DB-37 connector pinout and define a set of modem control signals for regulating flow control and line status.

The RS-422/499 line protocol runs in balanced mode, allowing serial communications to extend over distances of up to 4,000 feet (1.2 km) and at very fast speeds of up to 10 Mbps.

In an RS-422/499-based system, a single master device can communicate with up to 10 slave devices in the system. To accommodate this configuration, RS-422/499 supports the following kinds of transmission:

- Half-duplex transmission—In half-duplex transmission mode, transmissions occur in only one direction at a time. Each transmission requires a proper handshake before it is sent. This operation is typical of a balanced system in which two devices are connected by a single connection.
- Full-duplex transmission—In full duplex transmission mode, multiple transmissions can occur simultaneously so that devices can transmit and receive at the same time. This operation is essential when a single master in a point-to-multipoint system must communicate with multiple receivers.
- Multipoint transmission—RS-422/449 allows only a single master in a multipoint system. The master can communicate to all points in a multipoint system, and the other points must communicate with each other through the master.

V.35

V.35 is 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 Europe.

The V.35 line protocol is a mixture of balanced (RS-422) and common ground (RS-232) signal interfaces. The V.35 control signals DTR, DSR, DCD, RTS, and CTS are single-wire common ground signals that are essentially identical to their RS-232 equivalents. Unbalanced signaling for these control signals is sufficient, because the control signals are mostly constant, varying at very low frequency, which makes single-wire transmission suitable. Higher-frequency data and clock signals are sent over balanced wires.

V.35 interfaces operate at line rates of 20 Kbps and above.

X.21

X.21 is an ITU-T standard for serial communications over synchronous digital lines. The X.21 protocol is used primarily in Europe and Japan.

The X.21 line protocol is a state-driven protocol that sets up a circuit-switched network using call setup. X.21 interfaces use a 15-pin connector with the following eight signals:

- Signal ground (G)—Reference signal used to evaluate the logic states of the other signals. This signal can be connected to the protective earth (ground).
- DTE common return (Ga)—Reference ground signal for the DCE interface. This signal is used only in unbalanced mode.
- Transmit (T)—Binary signal that carries the data from the DTE to the DCE. This signal can be used for data transfer or in call-control phases such as Call Connect or Call Disconnect.
- Receive (R)—Binary signal that carries the data from the DCE to the DTE. This signal can be used for data transfer or in call-control phases such as Call Connect or Call Disconnect.
- Control (C)—DTE-controlled signal that controls the transmission on an X.21 link. This signal must be on during data transfer, and can be on or off during call-control phases.
- Indication (I)—DCE-controlled signal that controls the transmission on an X.21 link. This signal must be on during data transfer, and can be on or off during call-control phases.
- Signal Element Timing (S)—Clocking signal that is generated by the DCE. This signal specifies when sampling on the line must occur.
- Byte Timing (B)—Binary signal that is on when data or call-control information is being sampled. When an 8-byte transmission is over, this signal switches to off.

Transmissions across an X.21 link require both the DCE and DTE devices to be in a ready state, indicated by an all 1s transmission on the T and R signals.

ADSL Interface Overview

Asymmetric digital subscriber line (ADSL) technology is part of the xDSL family of modem technologies that use existing twisted-pair telephone lines to transport high-bandwidth data. ADSL lines connect service provider networks and customer sites over the "last mile" of the network—the loop between the service provider and the customer site.

ADSL transmission is asymmetric because the downstream bandwidth is typically greater than the upstream bandwidth. A typical ADSL circuit uses bandwidths of 1.5 Mbps to 2.0 Mbps downstream and 16 Kbps upstream. Depending on the length of the copper wire, an ADSL link can have up to 6.1 Mbps downstream and 64 Kbps upstream.

All Services Routers except the J2300 router support ADSL, ADSL2, and ADSL2 + , which comply with the following standards:

- For Annex A and B—ITU G.992.1 (ADSL)
- For Annex A only—ANSI T1.413 Issue II, ITU G.992.3 (ADSL2) and ITU G.992.5 (ADSL2 +)
- For Annex B only—ETSI TS 101 388 V1.3



NOTE: Services Routers with ADSL PIMs can use PPP over Ethernet (PPPoE) and PPP over ATM (PPPoA) to connect through ADSL lines only, not for direct ATM connections.


ADSL Systems

ADSL links run across twisted-pair telephone wires. When ADSL modems are connected to each end of a telephone wire, a dual-purpose ADSL circuit can be created. Once established, the circuit can transmit lower-frequency voice traffic and higher-frequency data traffic.

To accommodate both types of traffic, ADSL modems are connected to plain old telephone service (POTS) splitters that filter out the lower-bandwidth voice traffic and the higher-bandwidth data traffic. The voice traffic can be directed as normal telephone voice traffic. The data traffic is directed to the ADSL modem, which is typically connected to the data network.

Because twisted-pair wiring has a length limit, ADSL modems are typically connected to multiplexing devices. DSL access multiplexers (DSLAMs) can process and route traffic from multiple splitters. This typical ADSL configuration is shown in Figure 16 on page 73.

Figure 16: Typical ADSL Topology



ADSL2 and ADSL2+

The ADSL2 and ADSL2+ standards were adopted by the ITU in July 2002. ADSL2 improves the data rate and reach performance, diagnostics, standby mode, and interoperability of ADSL modems.

ADSL2+ doubles the possible downstream data bandwidth, enabling rates of 20 Mbps on telephone lines shorter than 5,000 feet (1.5 km).

First-generation ADSL standards require fixed 32-bit overhead framing on all ADSL packets. On long lines with low rates of 128 Kbps, the overhead represents 25 percent of the available bandwidth. ADSL2 standards allow the overhead per frame to be a programmable value between 4 Kbps and 32 Kbps, to provide up to 28 Kbps more bandwidth for payload data.

ADSL2 uses seamless rate adaptation (SRA) to change the data rate of a connection during operation with no interruptions or bit errors. The ADSL2 transceiver detects changes in channel conditions—for example, the failure of another transceiver in a multicarrier link—and sends a message to the transmitter to initiate a data rate change. The message includes data transmission parameters such as the number of bits modulated and the power on each channel. When the transmitter receives the information, it transitions to the new transmission rate.

Asynchronous Transfer Mode

On a J-series Services Router, the ADSL link is employed over an Asynchronous Transfer Mode (ATM)-over-ADSL interface. Although the interface type is **at**, the physical interface is ADSL. ATM-over-ADSL and ATM-over-SHDSL interfaces can be configured with the properties associated with traditional ATM interfaces, including virtual circuit and path information and ATM encapsulation.

SHDSL Interface Overview

SHDSL interfaces on J-series Service Routers support a symmetric, high-speed digital subscriber line (SHDSL) multirate technology for data transfer between a single customer premises equipment (CPE) subscriber and a central office (CO). ITU-T G.991.2 is the officially designated standard describing SHDSL, also known as G.SHDSL.

Unlike ADSL, which delivers more bandwidth downstream than available upstream, SHDSL is symmetrical and delivers a bandwidth of up to 2.3 Mbps in both directions. Because business applications require higher-speed digital transportation methods, SHDSL is becoming very popular and gaining wide acceptance in the industry. Additionally, SHDSL is compatible with ADSL and therefore causes very little, if any, interference between cables.

SHDSL is deployed on a network in much the same manner as ADSL.

ISDN Interface Overview

The Integrated Services Digital Network (ISDN) technology is a design for a completely digital telecommunications network. ISDN can carry voice, data, images, and video across a telephony network, using a single interface for all transmissions.

ISDN Channels

ISDN uses separate channels to transmit voice and data over the network. Channels operate at bandwidths of either 64 Kbps or 16 Kbps, depending on the type of channel.

Bearer channels (B-channels) use 64 Kbps to transmit voice, data, video, or multimedia information. This bandwidth is derived from the fact that analog voice lines are sampled at a rate of 64 Kbps (8,000 samples per second using 8 bits per sample).

Delta channels (D-channels) are control channels that operate at either 16 Kbps or 64 Kbps. D-channels are used primarily for ISDN signaling between switching equipment in an ISDN network.

ISDN Interfaces


ISDN provides two basic types of service, Basic Rate Interface (BRI) and Primary Rate Interface (PRI). Services Routers support both ISDN BRI and ISDN PRI.

ISDN BRI is designed for high-bandwidth data transmissions through existing telephone lines. The copper wires that make up much of the existing telephony infrastructure can support approximately 160 Kbps, which provides enough bandwidth for two B-channels and a D-channel, leaving 16 Kbps for any data framing, maintenance, and link control overhead.

ISDN PRI is designed for users with greater capacity requirements than can be met with ISDN BRI. In the United States, the most common PRI supports 23 B-channels and 1 D-channel, totalling 1,536 Kbps, which is roughly equivalent to a T1 link. In Europe, the most common PRI supports 30 B-channels and 1 D-channel, totalling 1,984 Kbps, which is roughly equivalent to an E1 link.

Typical ISDN Network

Figure 17 on page 76 shows a typical ISDN network.

Figure 17: ISDN Network

In Figure 17 on page 76, two types of end-user devices are connected to the ISDN network:

- Terminal equipment type 1 (TE1) device—Designed to connect directly through an ISDN telephone line.
- Terminal equipment type 2 (TE2) device—Not designed for ISDN. TE2 devices—for example, analog telephones or modems—must connect to the ISDN network through a terminal adapter (TA).

A terminal adapter allows non-ISDN devices on the ISDN network.

NT Devices and S and T Interfaces

The interface between the ISDN network and a TE1 device or terminal adapter is called an S interface. The S interface connects to a network termination type 2 (NT2) device such as a PBX, or directly to the TE1 device or terminal adapter, as shown in Figure 17 on page 76. The NT2 device is then connected to a network termination type 1 (NT1) device through a T interface. The S and T interfaces are electrically equivalent.

An NT1 device is a physical layer device that connects a home telephone network to a service provider carrier network. ISDN devices that connect to an NT1 device from the home network side use a 4-wire S/T interface. The NT1 device converts the 4-wire S/T interface into the 2-wire U interface that telephone carriers use as their plain old telephone service (POTS) lines.

In the United States, NT1 devices are user owned. In many other countries, NT1 devices are owned by the telephone service providers.

U Interface

The U interface connects the ISDN network into the telephone switch through line termination (LT) equipment. The connection from LT equipment to other switches within the telephone network is called the exchange termination (ET).

ISDN Call Setup

Before traffic can pass through an ISDN network, an ISDN call must be set up. ISDN call setup requires a Layer 2 connection to be initialized and then a Layer 3 session to be established over the connection.

To specify the services and features to be provided by the service provider switch, you must set service profile identifiers (SPIDs) on TE1 devices before call setup and initialization. If you define SPIDs for features that are not available on the ISDN link, Layer 2 initialization takes place, but a Layer 3 connection is not established.

Layer 2 ISDN Connection Initialization

The TE device and the telephone network initialize a Layer 2 connection for ISDN as follows:

1. The TE device and the telephone network exchange Receive Ready (RR) frames, to indicate that they are available for data transmission. A call cannot be set up if either the TE device or telephone network does not transmit RR frames.
2. If both ends of the ISDN connection are available to receive data, the TE device sends an Unnumbered Information (UI) frame to all devices on the ISDN link.
3. When it receives the UI frame, the network responds with a message containing a unique terminal endpoint identifier (TEI) that identifies the endpoint on the ISDN link for all subsequent data transmissions.
4. When the TE device receives the TEI message, it sends out a call setup message.
5. The network sends an acknowledgement of the call setup message.
6. When the TE device receives the acknowledgement, a Layer 2 connection is initialized on the ISDN link.

Layer 3 ISDN Session Establishment

The caller, switch, and receiver establish a Layer 3 ISDN connection as follows:

1. When a Layer 2 connection is initialized, the caller sends a SETUP message to the switch in the telephone network.
2. If the setup is message is valid, the switch responds with a call proceeding (CALL PROC) message to the caller and a SETUP message to the receiver.
3. When the receiver receives the SETUP message, it responds with an ALERTING message to the telephone switch.
4. This ALERTING message is then forwarded to the caller.
5. The receiver then accepts the connection by sending a CONNECT message to the switch.
6. The switch forwards the CONNECT message to the caller.
7. The caller responds with an acknowledgement message (CONNECT ACK).
8. When the CONNECT ACK message is received by the receiver, the ISDN call is set up and traffic can pass.

Interface Physical Properties

The physical properties of a network interface are the characteristics associated with the physical link that affect the transmission of either link-layer signals or the data across the links. Physical properties include clocking properties, transmission properties, such as the maximum transmission unit (MTU), and encapsulation methods, such as point-to-point and Frame Relay encapsulation.

The default property values for an interface are usually sufficient to successfully enable a bidirectional link. However, if you configure a set of physical properties on an interface, those same properties must be set on all adjacent interfaces to which a direct connection is made.

Table 24 on page 78 summarizes some key physical properties of J-series Services Router interfaces.

Table 24: Interface Physical Properties

Physical Property	Description
bert-error-rate	Bit error rate (BER). The error rate specifies the number of bit errors in a particular bit error rate test (BERT) period required to generate a BERT error condition. See “Bit Error Rate Testing” on page 79.
bert-period	Bit error rate test (BERT) time period over which bit errors are sampled. See “Bit Error Rate Testing” on page 79.
chap	Challenge Handshake Authentication Protocol (CHAP). Specifying chap enables CHAP authentication on the interface. See “PPP Authentication” on page 86.
clocking	Clock source for the link. Clocking can be provided by the local system (internal) or a remote endpoint on the link (external). By default, all interfaces use the internal clocking mode. If an interface is configured to accept an external clock source, one adjacent interface must be configured to act as a clock source. Under this configuration, the interface operates in a loop timing mode, in which the clocking signal is unique for that individual network segment or loop. See “Interface Clocking” on page 79.
description	A user-defined text description of the interface, often used to describe the interface's purpose.
disable	Administratively disables the interface.
encapsulation	Type of encapsulation on the interface. Common encapsulation types include PPP, Frame Relay, Cisco HDLC, and PPP over Ethernet (PPPoE). See “Physical Encapsulation on an Interface” on page 83.
fcs	Frame check sequence (FCS). FCS is an error-detection scheme that appends parity bits to a digital signal and uses decoding algorithms that detect errors in the received digital signal. See “Frame Check Sequences” on page 80.
mtu	Maximum transmission unit (MTU) size. The MTU is the largest size packet or frame, specified in bytes or octets, that can be sent in a packet-based or frame-based network. The Transmission Control Protocol (TCP) uses the MTU to determine the maximum size of each packet in any transmission. For MTU values on J-series interfaces, see “MTU Default and Maximum Values” on page 81.

Table 24: Interface Physical Properties *(continued)*

Physical Property	Description
no-keepalives	Disabling of keepalive messages across a physical link. A keepalive message is sent between network devices to indicate that they are still active. Keepalives help determine whether the interface is operating correctly. Except for ATM-over-ADSL interfaces, all interfaces use keepalives by default.
pap	Password Authentication Protocol (PAP). Specifying pap enables PAP authentication on the interface. To configure PAP, use the CLI or J-Web configuration editor. PAP is not available in the J-Web Quick Configuration pages.
payload-scrambler	Scrambling of traffic transmitted out the interface. Payload scrambling randomizes the data payload of transmitted packets. Scrambling eliminates nonvariable bit patterns (strings of all 1s or all 0s) that generate link-layer errors across some physical links.

Bit Error Rate Testing

In telecommunication transmission, the bit error rate (BER) is the percentage of bits that have errors compared to the total number of bits received in a transmission, usually expressed as 10 to a negative power. For example, a transmission with a BER of 10^{-6} received 1 errored bit in 1,000,000 bits transmitted. The BER indicates how often a packet or other data unit must be retransmitted because of an error. If the BER is too high, a slower data rate might improve the overall transmission time for a given amount of data if it reduces the BER and thereby lowers the number of resent packets.

A bit error rate test (BERT) is a procedure or device that measures the BER for a given transmission. You can configure a Services Router to act as a BERT device by configuring the interface with a bit error rate and a testing period. When the interface receives a BERT request from a BER tester, it generates a response in a well-known BERT pattern. The initiating device checks the BERT-patterned response to determine the number of bit errors.

Interface Clocking

Clocking determines how individual routing nodes or entire networks sample transmitted data. As streams of information are received by a router in a network, a clock source specifies when to sample the data. In asynchronous networks, the clock source is derived locally, and synchronous networks use a central, external clock source. Interface clocking indicates whether the device uses asynchronous or synchronous clocking.



NOTE: Because truly synchronous networks are difficult to design and maintain, most synchronous networks are really plesiochronous networks. In a plesiochronous network, different timing regions are controlled by local clocks that are synchronized (with very narrow constraints). Such networks approach synchronicity and are generally known as synchronous networks.

Most networks are designed to operate as asynchronous networks. Each device generates its own clock signal, or devices use clocks from more than one clock source.

The clocks within the network are not synchronized to a single clock source. By default, Services Routers generate their own clock signals to send and receive traffic.

The system clock allows the router to sample (or detect) and transmit data being received and transmitted through its interfaces. Clocking enables the router to detect and transmit the 0s and 1s that make up digital traffic through the interface. Failure to detect the bits within a data flow results in dropped traffic.

Short-term fluctuations in the clock signal are known as clock jitter. Long-term variations in the signal are known as clock wander.

Asynchronous clocking can either derive the clock signal from the data stream or transmit the clocking signal explicitly.

Data Stream Clocking

Common in T1 links, data stream clocking occurs when separate clock signals are not transmitted within the network. Instead, devices must extract the clock signal from the data stream. As bits are transmitted across the network, each bit has a time slot of 648 nanoseconds. Within a time slot, pulses are transmitted with alternating voltage peaks and drops. The receiving device uses the period of alternating voltages to determine the clock rate for the data stream.

Explicit Clocking Signal Transmission

Clock signals that are shared by hosts across a data link must be transmitted by one or both endpoints on the link. In a serial connection, for example, one host operates as a clock master and the other operates as a clock slave. The clock master internally generates a clock signal that is transmitted across the data link. The clock slave receives the clock signal and uses its period to determine when to sample data and how to transmit data across the link.

This type of clock signal controls only the connection on which it is active and is not visible to the rest of the network. An explicit clock signal does not control how other devices or even other interfaces on the same device sample or transmit data.

Frame Check Sequences

All packets or frames within a network can be damaged by crosstalk or interference in the network's physical wires. The frame check sequence (FCS) is an extra field in each transmitted frame that can be analyzed to determine if errors have occurred. The FCS uses cyclic redundancy checks (CRCs), checksums, and two-dimensional parity bits to detect errors in the transmitted frames.

Cyclic Redundancy Checks and Checksums

On a link that uses CRCs for frame checking, the data source uses a predefined polynomial algorithm to calculate a CRC number from the data it is transmitting. The result is included in the FCS field of the frame and transmitted with the data. On the receiving end, the destination host performs the same calculation on the data it receives.

If the result of the second calculation matches the contents of the FCS field, the packet was sent and received without bit errors. If the values do not match, an FCS error is generated, the frame is discarded and the originating host is notified of the error.

Checksums function similarly to CRCs, but use a different algorithm.

Two-Dimensional Parity

On a link that uses two-dimensional parity bits for frame checking, the sending and receiving hosts examine each frame in the total packet transmission and create a parity byte that is evaluated to detect transmission errors.

For example, a host can create the parity byte for the following frame sequence by summing up each column (each bit position in the frame) and keeping only the least-significant bit:

```

Frame 1      0 1 0 1 0 0 1
Frame 2      1 1 0 1 0 0 1
Frame 3      1 0 1 1 1 1 0
Frame 4      0 0 0 1 1 1 0
Frame 5      0 1 1 0 1 0 0
Frame 6      1 0 1 1 1 1 1

```

```

Parity Byte  1 1 1 1 0 1 1

```

If the sum of the bit values in a bit position is even, the parity bit for the position is 0. If the sum is odd, the parity bit is 1. This method is called even parity. Matching parity bytes on the originating and receiving hosts indicate that the packet was received without error.

MTU Default and Maximum Values

Table 25 on page 81 lists maximum transmission unit (MTU) sizes for interfaces on J2300, J4300, and J6300 Services Routers. Table 26 on page 82 lists MTU values for J4350 and J6350 Services Routers.

Table 25: MTU Values for J2300, J4300, and J6300 Interfaces

Model and Interface	Default Media MTU (bytes)	Maximum MTU (bytes)	Default IP MTU (bytes)
J2300 Built-In Interface			
Fast Ethernet (10/100)	1514	9192	1500
T1 or E1	1504	9192	1500
Serial	1504	9192	1500
ISDN BRI	1504	4092	1500
G.SHDSL	4482	9192	4470

Table 25: MTU Values for J2300, J4300, and J6300 Interfaces *(continued)*

Model and Interface	Default Media MTU (bytes)	Maximum MTU (bytes)	Default IP MTU (bytes)
J4300 and J6300 Interfaces			
Fast Ethernet (10/100) built-in interface	1514	9192	1500
Dual-Port Fast Ethernet (10/100) PIM	1514	9192	1500
Dual-Port Serial PIM	1504	9192	1500
Dual-Port T1 or E1 PIM	1504	9192	1500
Dual-Port Channelized T1/E1/ISDN PRI PIM (channelized to DS0s)	1504	4500	1500
Dual-Port Channelized T1/E1/ISDN PRI PIM (clear-channel T1 or E1)	1504	9150	1500
Dual-Port Channelized T1/E1/ISDN PRI PIM (ISDN PRI dialer interface)	1504	4098	1500
T3 (DS3) or E3 PIM	4474	9192	4470
4-Port ISDN BRI PIM	1504	4092	1500
ADSL + 2 PIM	4482	9192	4470
G.SHDSL PIM	4482	9192	4470

Table 26: MTU Values for J4350 and J6350 Interfaces

J4350 and J6350 Interfaces	Default Media MTU (bytes)	Maximum MTU (bytes)	Default IP MTU (bytes)
Gigabit Ethernet (10/100/1000) built-in interface	1514	9018	1500
6-Port, 8-Port, and 16-Port Gigabit Ethernet uPIMs	1514	9014	1500
Gigabit Ethernet (10/100/1000) ePIM	1514	9018	1500
Gigabit Ethernet (10/100/1000) SFP ePIM	1514	9018	1500
4-Port Fast Ethernet (10/100) ePIM	1514	1514	1500
Dual-Port Fast Ethernet (10/100) PIM	1514	9192	1500
Dual-Port Serial PIM	1504	9192	1500
Dual-Port T1 or E1 PIM	1504	9192	1500
Dual-Port Channelized T1/E1/ISDN PRI PIM (channelized to DS0s)	1504	4500	1500

Table 26: MTU Values for J4350 and J6350 Interfaces *(continued)*

J4350 and J6350 Interfaces	Default Media MTU (bytes)	Maximum MTU (bytes)	Default IP MTU (bytes)
Dual-Port Channelized T1/E1/ISDN PRI PIM (clear-channel T1 or E1)	1504	9150	1500
Dual-Port Channelized T1/E1/ISDN PRI PIM (ISDN PRI dialer interface)	1504	4098	1500
T3 (DS3) or E3 PIM	4474	9192	4470
4-Port ISDN BRI PIM	1504	4092	1500
ADSL + 2 PIM	4482	9192	4470
G.SHDSL PIM	4482	9192	4470

Physical Encapsulation on an Interface

Encapsulation is the process by which a lower-level protocol accepts a message from a higher-level protocol and places it in the data portion of the lower-level frame. As a result, datagrams transmitted through a physical network have a sequence of headers: the first header for the physical network (or data link layer) protocol, the second header for the network layer protocol (IP, for example), the third header for the transport protocol, and so on.

The following encapsulation protocols are supported on J-series Services Router physical interfaces:

- Frame Relay on page 83
- Point-to-Point Protocol on page 85
- Point-to-Point Protocol over Ethernet on page 88
- High-Level Data Link Control on page 89

Frame Relay

The Frame Relay packet-switching protocol operates at the physical and data link layers in a network to optimize packet transmissions by creating virtual circuits between hosts. Figure 18 on page 84 shows a typical Frame Relay network.


Figure 18: Frame Relay Network

Figure 18 on page 84 shows multiple paths from Host A to Host B. In a typical routed network, traffic is sent from device to device with each device making routing decisions based on its own routing table. In a packet-switched network, the paths are predefined. Devices switch a packet through the network according to predetermined next-hops established when the virtual circuit is set up.

Virtual Circuits

A virtual circuit is a bidirectional path between two hosts in a network. Frame Relay virtual circuits are logical connections between two hosts that are established either by a call setup mechanism or by explicit configuration.

A virtual circuit created through a call setup mechanism is known as a switched virtual circuit (SVC). A virtual circuit created through explicit configuration is called a permanent virtual circuit (PVC).

Switched and Permanent Virtual Circuits

Before data can be transmitted across an SVC, a signaling protocol like ISDN must set up a call by the exchange of setup messages across the network. When a connection is established, data is transmitted across the SVC. After data transmission, the circuit is torn down and the connection is lost. For additional traffic to pass between the same two hosts, a subsequent SVC must be established, maintained, and terminated.

Because PVCs are explicitly configured, they do not require the setup and teardown of SVCs. Data can be switched across the PVC whenever a host is ready to transmit. SVCs are useful in networks where data transmission is sporadic and a permanent circuit is not needed.

Data-Link Connection Identifiers

An established virtual circuit is identified by a data-link connection identifier (DLCI). The DLCI is a value from 16 through 1022. (Values 1 through 15 are reserved.) The

DLCI uniquely identifies a virtual circuit locally so that routers can switch packets to the appropriate next-hop address in the circuit. Multiple paths that pass through the same transit routers have different DLCIs and associated next-hop addresses.

Congestion Control and Discard Eligibility

Frame Relay uses the following types of congestion notification to control traffic within a Frame Relay network. Both are controlled by a single bit in the Frame Relay header.

- Forward-explicit congestion notification (FECN)
- Backward-explicit congestion notification (BECN)

Traffic congestion is typically defined in the buffer queues on a router. When the queues reach a predefined level of saturation, traffic is determined to be congested. When traffic congestion occurs in a virtual circuit, the router experiencing congestion sets the congestion bits in the Frame Relay header to 1. As a result, transmitted traffic has the FECN bit set to 1, and return traffic on the same virtual circuit has the BECN bit set to 1.

When the FECN and BECN bits are set to 1, they provide a congestion notification to the source and destination devices. The devices can respond in either of two ways: to control traffic on the circuit by sending it through other routes, or to reduce the load on the circuit by discarding packets.

If devices discard packets as a means of congestion (flow) control, Frame Relay uses the discard eligibility (DE) bit to give preference to some packets in discard decisions. A DE value of 1 indicates that the frame is of lower importance than other frames and more likely to be dropped during congestion. Critical data (such as signaling protocol messages) without the DE bit set is less likely to be dropped.

Point-to-Point Protocol

The Point-to-Point Protocol (PPP) is an encapsulation protocol for transporting IP traffic across point-to-point links. PPP is made up of three primary components:

- Link control protocol (LCP)—Establishes working connections between two points.
- Authentication protocols—Enable secure connections between two points.
- Network control protocols (NCPs)—Initialize the PPP protocol stack to handle multiple network layer protocols, such as IPv4, IPv6, and Connectionless Network Protocol (CLNP).

Link Control Protocol

LCP is responsible for establishing, maintaining, and tearing down a connection between two endpoints. LCP also tests the link and determines whether it is active. LCP establishes a point-to-point connection as follows:

1. LCP must first detect a clocking signal on each endpoint. However, because the clocking signal can be generated by a network clock and shared with devices on the network, the presence of a clocking signal is only a preliminary indication that the link might be functioning.
2. When a clocking signal is detected, a PPP host begins transmitting PPP Configure-Request packets.
3. If the remote endpoint on the point-to-point link receives the Configure-Request packet, it transmits a Configure-Acknowledgement packet to the source of the request.
4. After receiving the acknowledgement, the initiating endpoint identifies the link as established. At the same time, the remote endpoint sends its own request packets and processes the acknowledgement packets. In a functioning network, both endpoints treat the connection as established.

During connection establishment, LCP also negotiates connection parameters such as FCS and HDLC framing. By default, PPP uses a 16-bit FCS, but you can configure PPP to use either a 32-bit FCS or a 0-bit FCS (no FCS). Alternatively, you can enable HDLC encapsulation across the PPP connection.

After a connection is established, PPP hosts generate Echo-Request and Echo-Response packets to maintain a PPP link.

PPP Authentication

PPP's authentication layer uses a protocol to help ensure that the endpoint of a PPP link is a valid device. Authentication protocols include the Password Authentication Protocol (PAP), the Extensible Authentication Protocol (EAP), and the Challenge Handshake Authentication Protocol (CHAP). CHAP is the most commonly used.



NOTE: EAP is not currently supported on J-series Services Routers. PAP is supported, but must be configured from the CLI or J-Web configuration editor. PAP is not configurable from the J-Web Quick Configuration pages.

CHAP ensures secure connections across PPP links. After a PPP link is established by LCP, the PPP hosts at either end of the link initiate a three-way CHAP handshake. Two separate CHAP handshakes are required before both sides identify the PPP link as established.

CHAP configuration requires each endpoint on a PPP link to use a shared secret (password) to authenticate challenges. The shared secret is never transmitted over the wire. Instead, the hosts on the PPP connection exchange information that enables both to determine that they share the same secret. Challenges consist of a hash function calculated from the secret, a numeric identifier, and a randomly chosen challenge value that changes with each challenge. If the response value matches the challenge value, authentication is successful. Because the secret is never transmitted and is required to calculate the challenge response, CHAP is considered very secure.

PAP authentication protocol uses a simple 2-way handshake to establish identity. PAP is used after the link establishment phase (LCP up), during the authentication

phase. JUNOS software can support PAP in one direction (egress or ingress), and CHAP in the other.

Network Control Protocols

After authentication is completed, the PPP connection is fully established. At this point, any higher-level protocols (for example, IP protocols) can initialize and perform their own negotiations and authentication.

PPP NCPs include support for the following protocols. IPCP and IPV6CP are the most widely used on J-series Services Routers.

- ATCP—AppleTalk Control Protocol
- BCP—Bridging Control Protocol
- BVCP—Banyan Vines Control Protocol
- DNCP—DECnet Phase IV Control Protocol
- IPCP—IP Control Protocol
- IPV6CP—IPv6 Control Protocol
- IPXCP—Novell IPX Control Protocol
- LECP—LAN Extension Control Protocol
- NBFCP—NetBIOS Frames Control Protocol
- OSINLCP—OSI Network Layer Control Protocol (includes IS-IS, ES-IS, CLNP, and IDRP)
- SDTP—Serial Data Transport Protocol
- SNACP—Systems Network Architecture (SNA) Control Protocol
- XNSCP—Xerox Network Systems (XNS) Internet Datagram Protocol (IDP) Control Protocol

Magic Numbers

Hosts running PPP can create “magic” numbers for diagnosing the health of a connection. A PPP host generates a random 32-bit number and sends it to the remote endpoint during LCP negotiation and echo exchanges.

In a typical network, each host's magic number is different. A magic number mismatch in an LCP message informs a host that the connection is not in loopback mode and traffic is being exchanged bidirectionally. If the magic number in the LCP message is the same as the configured magic number, the host determines that the connection is in loopback mode, with traffic looped back to the transmitting host.

Looping traffic back to the originating host is a valuable way to diagnose network health between the host and the loopback location. To enable loopback testing, telecommunications equipment typically supports channel service unit/data service unit (CSU/DSU) devices.

CSU/DSU Devices

A channel service unit (CSU) connects a terminal to a digital line. A data service unit (DSU) performs protective and diagnostic functions for a telecommunications line. Typically, the two devices are packaged as a single unit. A CSU/DSU device is required for both ends of a T1 or T3 connection, and the units at both ends must be set to the same communications standard.

A CSU/DSU device enables frames sent along a link to be looped back to the originating host. Receipt of the transmitted frames indicates that the link is functioning correctly up to the point of loopback. By configuring CSU/DSU devices to loop back at different points in a connection, network operators can diagnose and troubleshoot individual segments in a circuit.

Point-to-Point Protocol over Ethernet

Point-to-Point Protocol over Ethernet (PPPoE) combines PPP, which is typically run over broadband connections, with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator. PPPoE enables service providers to maintain access control through PPP connections and also manage multiple hosts at a remote site.

To provide a PPPoE connection, each PPP session must learn the Ethernet address of the remote peer and establish a unique session identifier during the PPPoE discovery and session stages.

PPPoE Discovery

To initiate a PPPoE session, a host must first identify the Ethernet MAC address of the remote peer and establish a unique PPPoE session ID for the session. Learning the remote Ethernet MAC address is called PPPoE discovery.

During the PPPoE discovery process, the host does not discover a remote endpoint on the Ethernet network. Instead, the host discovers the access concentrator through which all PPPoE sessions are established. Discovery is a client/server relationship, with the host (a J-series Services Router) acting as the client and the access concentrator acting as the server.

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 to 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 Sessions

The PPPoE session stage starts after the PPPoE discovery stage is over. 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. Magic numbers, echo requests, and all other PPP traffic behave exactly as in normal PPP sessions. 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.

High-Level Data Link Control

High-Level Data Link Control (HDLC) is a bit-oriented, switched and nonswitched link-layer protocol. HDLC is widely used because it supports half-duplex and full-duplex connections, point-to-point and point-to-multipoint networks, and switched and nonswitched channels.

HDLC Stations

Nodes within a network running HDLC are called stations. HDLC supports three types of stations for data link control:

- Primary stations—Responsible for controlling the secondary and combined other stations on the link. Depending on the HDLC mode, the primary station is responsible for issuing acknowledgement packets to allow data transmission from secondary stations.
- Secondary stations—Controlled by the primary station. Under normal circumstances, secondary stations cannot control data transmission across the link with the primary station, are active only when requested by the primary station, and can respond to the primary station only (not to other secondary stations). All secondary station frames are response frames.

- Combined stations—A combination of primary and secondary stations. On an HDLC link, all combined stations can send and receive commands and responses without any permission from any other stations on the link and cannot be controlled by any other station.

HDLC Operational Modes

HDLC runs in three separate modes:

- Normal Response Mode (NRM)—The primary station on the HDLC link initiates all information transfers with secondary stations. A secondary station on the link can transmit a response of one or more information frames only when it receives explicit permission from the primary station. When the last frame is transmitted, the secondary station must wait for explicit permission before it can transmit more frames.

NRM is used most widely for point-to-multipoint links, in which a single primary station controls many secondary stations.

- Asynchronous Response Mode (ARM)—The secondary station can transmit either data or control traffic at any time, without explicit permission from the primary station. The primary station is responsible for error recovery and link setup, but the secondary station can transmit information at any time.

ARM is used most commonly with point-to-point links, because it reduces the overhead on the link by eliminating the need for control packets.

- Asynchronous Balance Mode (ABM)—All stations are combined stations. Because no other station can control a combined station, all stations can transmit information without explicit permission from any other station. ABM is not a widely used HDLC mode.

Interface Logical Properties

The logical properties of an interface are the characteristics that do not apply to the physical interface or the wires connected to it. Logical properties include the protocol families running on the interface (including any protocol-specific MTUs), the IP address or addresses associated with the interface, virtual LAN (VLAN) tagging, and any firewall filters or routing policies that are operating on the interface.

The IP specification requires a unique address on every interface of each system attached to an IP network, so that traffic can be correctly routed. Individual hosts such as home computers must have a single IP address assigned. Routers must have a unique IP address for every interface.

This section contains the following topics:

- Protocol Families on page 91
- IPv4 Addressing on page 91
- IPv6 Addressing on page 94
- Virtual LANs on page 96

Protocol Families

A protocol family is a group of logical properties within an interface configuration. Protocol families include all the protocols that make up a protocol suite. To use a protocol within a particular suite, you must configure the entire protocol family as a logical property for an interface. The protocol families include common and not-so-common protocol suites.

Common Protocol Suites

JUNOS protocol families include the following common protocol suites:

- **Inet**—Supports IP protocol traffic, including Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), and Internet Control Message Protocol (ICMP).
- **Inet6**—Supports IPv6 protocol traffic, including Routing Information Protocol for IPv6 (RIPng), Intermediate System-to-Intermediate System (IS-IS), and BGP.
- **ISO**—Supports IS-IS traffic.
- **MPLS**—Supports Multiprotocol Label Switching (MPLS).

Other Protocol Suites

In addition to the common protocol suites, JUNOS protocol families sometimes use the following protocol suites:

- **ccc**—Circuit cross-connect (CCC).
- **mlfr-uni-nni**—Multilink Frame Relay (MLFR) FRF.16 user-to-network network-to-network (UNI NNI).
- **mlfr-end-to-end**—Multilink Frame Relay end-to-end.
- **mlppp**—Multilink Point-to-Point Protocol.
- **tcc**—Translational cross-connect (TCC).
- **tnp**—Trivial Network Protocol. This Juniper Networks proprietary protocol provides communication 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.

IPv4 Addressing

IPv4 addresses are 32-bit numbers that are typically displayed in dotted decimal notation. A 32-bit address contains two primary parts: the network prefix and the host number.

All hosts within a single network share the same network address. Each host also has an address that uniquely identifies it. Depending on the scope of the network and the type of device, the address is either globally or locally unique. Devices that are visible to users outside the network (Web servers, for example) must have a globally unique IP address. Devices that are visible only within the network (routers, for example) must have locally unique IP addresses.

IP addresses are assigned by a central numbering authority called the Internet Assigned Numbers Authority (IANA). IANA ensures that addresses are globally unique where needed and has a large address space reserved for use by devices not visible outside their own networks.

IPv4 Classful Addressing

To provide flexibility in the number of addresses distributed to networks of different sizes, 4-octet (32-bit) IP addresses were originally divided into three different categories or classes: class A, class B, and class C. Each address class specifies a different number of bits for its network prefix and host number:

- Class A addresses use only the first byte (octet) to specify the network prefix, leaving 3 bytes to define individual host numbers.
- Class B addresses use the first 2 bytes to specify the network prefix, leaving 2 bytes to define host addresses.
- Class C addresses use the first 3 bytes to specify the network prefix, leaving only the last byte to identify hosts.

In binary format, with an x representing each bit in the host number, the three address classes can be represented as follows:

```
00000000 xxxxxxxx xxxxxxxx xxxxxxxx (Class A)
00000000 00000000 xxxxxxxx xxxxxxxx (Class B)
00000000 00000000 00000000 xxxxxxxx (Class C)
```

Because each bit (x) in a host number can have a 0 or 1 value, each represents a power of 2. For example, if only 3 bits are available for specifying the host number, only the following host numbers are possible:

```
111 110 101 100 011 010 001 000
```

In each IP address class, the number of host-number bits raised to the power of 2 indicates how many host numbers can be created for a particular network prefix. Class A addresses have 2^{24} (or 16,777,216) possible host numbers, class B addresses have 2^{16} (or 65,536) host numbers, and class C addresses have 2^8 (or 256) possible host numbers.

IPv4 Dotted Decimal Notation

The 32-bit IPv4 addresses are most often expressed in dotted decimal notation, in which each octet (or byte) is treated as a separate number. Within an octet, the rightmost bit represents 2^0 (or 1), increasing to the left until the first bit in the octet is 2^7 (or 128). Following are IP addresses in binary format and their dotted decimal equivalents:

```
11010000 01100010 11000000 10101010 = 208.98.192.170
01110110 00001111 11110000 01010101 = 118.15.240.85
00110011 11001100 00111100 00111011 = 51.204.60.59
```

IPv4 Subnetting

Because of the physical and architectural limitations on the size of networks, you often must break large networks into smaller subnetworks. Within a network, each wire or ring requires its own network number and identifying subnet address.

Figure 19 on page 93 shows two subnets in a network.

Figure 19: Subnets in a Network




Figure 19 on page 93 shows three devices connected to one subnet and three more devices connected to a second subnet. Collectively, the six devices and two subnets make up the larger network. In this example, the network is assigned the network prefix **192.14.0.0**, a class B address. Each device has an IP address that falls within this network prefix.

In addition to sharing a network prefix (the first two octets), the devices on each subnet share a third octet. The third octet identifies the subnet. All devices on a subnet must have the same subnet address. In this case, the alpha subnet has the IP address **192.14.126.0** and the beta subnet has the IP address **192.14.17.0**.

The subnet address **192.14.17.0** can be represented as follows in binary notation:

11000000 . 00001110 . 00010001 . xxxxxxxx

Because the first 24 bits in the 32-bit address identify the subnet, the last 8 bits are not significant. To indicate the subnet, the address is written as **192.14.17.0/24** (or just **192.14.17/24**). The /24 is the subnet mask (sometimes shown as 255.255.255.0).

IPv4 Variable-Length Subnet Masks

Traditionally, subnets were divided by address class. Subnets had either 8, 16, or 24 significant bits, corresponding to 2^8 , 2^{16} , or 2^{24} possible hosts. As a result, an

entire /16 subnet had to be allocated for a network that required only 400 addresses, wasting 65,136 ($2^{16} - 400 = 65,136$) addresses.

To help allocate address spaces more efficiently, variable-length subnet masks (VLSMs) were introduced. Using VLSM, network architects can allocate more precisely the number of addresses required for a particular subnet.

For example, suppose a network with the prefix **192.14.17/24** is divided into two smaller subnets, one consisting of 18 devices and the other of 46 devices.

To accommodate 18 devices, the first subnet must have 2^5 (32) host numbers. Having 5 bits assigned to the host number leaves 27 bits of the 32-bit address for the subnet. The IP address of the first subnet is therefore **192.14.17.128/27**, or the following in binary notation:

```
11000000 . 00001110 . 00010001 . 100xxxxx
```

The subnet mask includes 27 significant digits.

To create the second subnet of 46 devices, the network must accommodate 2^6 (64) host numbers. The IP address of the second subnet is **192.14.17.64/26**, or

```
11000000 . 00001110 . 00010001 . 01xxxxxx
```

By assigning address bits within the larger /24 subnet mask, you create two smaller subnets that use the allocated address space more efficiently.

IPv6 Addressing

To create a much larger address space and relieve a projected future shortage of IP addresses, IPv6 was created. IPv6 addresses consist of 128 bits, instead of 32 bits, and include a scope field that identifies the type of application suitable for the address. IPv6 does not support broadcast addresses, but instead uses multicast addresses for broadcast. In addition, IPv6 defines a new type of address called anycast.

IPv6 Address Representation

IPv6 addresses consist of 8 groups of 16-bit hexadecimal values separated by colons (:). IPv6 addresses have the following format:

```
aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa:aaaa
```

Each **aaaa** is a 16-bit hexadecimal value, and each **a** is a 4-bit hexadecimal value. Following is a sample IPv6 address:

```
3FFE:0000:0000:0001:0200:F8FF:FE75:50DF
```

You can omit the leading zeros of each 16-bit group, as follows:

```
3FFE:0:0:1:200:F8FF:FE75:50DF
```


You can compress 16-bit groups of zeros to double colons (::) as shown in the following example, but only once per address:

```
3FFE::1:200:F8FF:FE75:50DF
```

IPv6 Address Types

IPv6 has three types of addresses:

- Unicast—For a single interface.
- Multicast—For a set of interfaces on the same physical medium. A packet is sent to all interfaces associated with the address.
- Anycast—For a set of interfaces on different physical media. A packet is sent to only one of the interfaces associated with this address, not to all the interfaces.

IPv6 Address Scope

Unicast and multicast IPv6 addresses support feature called address scoping that identifies the application suitable for the address.

Unicast addresses support global address scope and two types of local address scope:

- Link-local unicast addresses—Used only on a single network link. The first 10 bits of the prefix identify the address as a link-local address. Link-local addresses cannot be used outside the link.
- Site-local unicast addresses—Used only within a site or intranet. A site consists of multiple network links. Site-local addresses identify nodes inside the intranet and cannot be used outside the site.

Multicast addresses support 16 different types of address scope, including node, link, site, organization, and global scope. A 4-bit field in the prefix identifies the address scope.

IPv6 Address Structure

Unicast addresses identify a single interface. Each unicast address consists of n bits for the prefix, and $128 - n$ bits for the interface ID.

Multicast addresses identify a set of interfaces. Each multicast address consists of the first 8 bits of all 1s, a 4-bit flags field, a 4-bit scope field, and a 112-bit group ID:

```
11111111 | flgs | scop | group ID
```

The first octet of 1s identifies the address as a multicast address. The flags field identifies whether the multicast address is a well-known address or a transient multicast address. The scope field identifies the scope of the multicast address. The 112-bit group ID identifies the multicast group.

Similar to multicast addresses, anycast addresses identify a set of interfaces. However, packets are sent to only one of the interfaces, not to all interfaces. Anycast addresses


are allocated from the normal unicast address space and cannot be distinguished from a unicast address in format. Therefore, each member of an anycast group must be configured to recognize certain addresses as anycast addresses.

Virtual LANs

A local area network (LAN) is a single broadcast domain. When traffic is broadcast, all hosts within the LAN receive the broadcast traffic. A LAN is determined by the physical connectivity of devices within the domain.

Within a traditional LAN, hosts are connected by a hub or repeater that propagates any incoming traffic throughout the network. Each host and its connecting hubs or repeaters make up a LAN segment. LAN segments are connected through switches and bridges to form the broadcast domain of the LAN. Figure 20 on page 96 shows a typical LAN topology.

Figure 20: Typical LAN




Virtual LANs (VLANs) allow network architects to segment LANs into different broadcast domains based on logical groupings. Because the groupings are logical, the broadcast domains are not determined by the physical connectivity of the devices in the network. Hosts can be grouped according to a logical function, to limit the traffic broadcast within the VLAN to only the devices for which the traffic is intended.

Suppose a corporate network has three major organizations: engineering, sales, and support. Using VLAN tagging, hosts within each organization can be tagged with a different VLAN identifier. Traffic sent to the broadcast domain is then checked against

the VLAN identifier and broadcast to only the devices in the appropriate VLAN. Figure 21 on page 97 shows a typical VLAN topology.

Figure 21: Typical VLAN



Special Interfaces

In addition to the configured network interfaces associated with the physical ports and wires that make up much of the network, J-series Services Routers have special interfaces. Table 27 on page 97 lists each special interface and briefly describes its use.

For information about interface names, see “Network Interface Naming” on page 47.

Table 27: Special Interfaces on a Services Router

Interface Name	Description
dsc	Discard interface. See “Discard Interface” on page 100.
fxp0	<p>This interface is not supported on a J-series Services Router. (On an M-series or T-series router, <code>fxp0</code> is used for out-of-band management.) For more information about the J-series Services Router management port interface, see “Management Interface” on page 101.</p> <p>For information about the interfaces supported on J-series, M-series, and T-series routers, see the <i>JUNOS Interfaces Command Reference</i>.</p>
gr-0/0/0	<p>Configurable generic routing encapsulation (GRE) interface. GRE allows the encapsulation of one routing protocol over another routing protocol.</p> <p>Within a Services Router, packets are routed to this internal interface, where they are first encapsulated with a GRE packet and then re-encapsulated with another protocol packet to complete the GRE. The GRE interface is an internal interface only and is not associated with a physical medium or PIM. You must configure the interface for it to perform GRE.</p>
gre	Internally generated GRE interface. This interface is generated by the JUNOS software to handle GRE. It is not a configurable interface.

Table 27: Special Interfaces on a Services Router (continued)

Interface Name	Description
ip-0/0/0	<p>Configurable IP-over-IP encapsulation (also called IP tunneling) interface. IP tunneling allows the encapsulation of one IP packet over another IP packet.</p> <p>Generally, IP routing allows packets to be routed directly to a particular address. However, in some instances you might need to route an IP packet to one address and then encapsulate it for forwarding to a different address. In a mobile environment in which the location of the end device changes, a different IP address might be used as the end device migrates between networks.</p> <p>Within a Services Router, packets are routed to this internal interface where they are encapsulated with an IP packet and then forwarded to the encapsulating packet's destination address. The IP-IP interface is an internal interface only and is not associated with a physical medium or PIM. You must configure the interface for it to perform IP tunneling.</p>
ipip	Internally generated IP-over-IP interface. This interface is generated by the JUNOS software to handle IP-over-IP encapsulation. It is not a configurable interface.
lo0	Loopback address. The loopback address has several uses, depending on the particular JUNOS feature being configured. See “Loopback Interface” on page 100.
lo0.16385	Internal loopback address. The internal loopback address is a particular instance of the loopback address with the logical unit number 16385. It is created by the JUNOS software as the loopback interface for the internal routing instance. This interface prevents any filter on lo0.0 from disrupting internal traffic.
ls-0/0/0	<p>Configurable link services interface. Link services include the multilink services MLPPP, MLFR, and Compressed Real-Time Transport Protocol (CRTP).</p> <p>Within a Services Router, packets are routed to this internal interface for link bundling or compression. The link services interface is an internal interface only and is not associated with a physical medium or PIM. You must configure the interface for it to perform multilink services.</p> <p>For more information about multilink services, see “Services Interfaces” on page 101.</p>
lsi	Internally generated link services interface. This interface is generated by the JUNOS software to handle multilink services like MLPPP, MLFR, and CRTP. It is not a configurable interface.
lt-0/0/0	<p>Interface used to provide class-of-service (CoS) support for data link switching (DLSw) traffic and real-time performance monitoring (RPM) probe packets.</p> <p>Within a Services Router, packets are routed to this internal interface for services. The lt interface is an internal interface only and is not associated with a physical medium or PIM. You must configure the interface for it to perform CoS for DLSW and RPM services.</p> <p>NOTE: The lt interface on the M-series and T-series routing platforms supports configuration of logical routers—the capability to partition a single physical router into multiple logical devices that perform independent routing tasks. However, the lt interface on the J-series Services Router does not support logical routers.</p>
mt-0/0/0	<p>Configurable multicast tunnel interface. Multicast tunnels filter all unicast packets; if an incoming packet is not destined for a 224/8-or-greater prefix, the packet is dropped and a counter is incremented.</p> <p>Within a Services Router, packets are routed to this internal interface for multicast filtering. The multicast tunnel interface is an internal interface only and is not associated with a physical medium or PIM. You must configure the interface for it to perform multicast tunneling.</p>

Table 27: Special Interfaces on a Services Router (continued)

Interface Name	Description
mtun	Internally generated multicast tunnel interface. This interface is generated by the JUNOS software to handle multicast tunnel services. It is not a configurable interface.
pd-0/0/0	<p>Configurable Protocol Independent Multicast (PIM) de-encapsulation interface. In PIM sparse mode, the first-hop routing platform encapsulates packets destined for the rendezvous point router. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree.</p> <p>Within a Services Router, packets are routed to this internal interface for de-encapsulation. The PIM de-encapsulation interface is an internal interface only and is not associated with a physical medium or Physical Interface Module (PIM). You must configure the interface for it to perform PIM de-encapsulation.</p>
pe-0/0/0	<p>Configurable Protocol Independent Multicast (PIM) encapsulation interface. In PIM sparse mode, the first-hop routing platform encapsulates packets destined for the rendezvous point router. The packets are encapsulated with a unicast header and are forwarded through a unicast tunnel to the rendezvous point. The rendezvous point then de-encapsulates the packets and transmits them through its multicast tree.</p> <p>Within a Services Router, packets are routed to this internal interface for encapsulation. The PIM encapsulation interface is an internal interface only and is not associated with a physical medium or Physical Interface Module (PIM). You must configure the interface for it to perform PIM encapsulation.</p>
pimd	Internally generated Protocol Independent Multicast (PIM) de-encapsulation interface. This interface is generated by the JUNOS software to handle PIM de-encapsulation. It is not a configurable interface.
pime	Internally generated Protocol Independent Multicast (PIM) encapsulation interface. This interface is generated by the JUNOS software to handle PIM encapsulation. It is not a configurable interface.
pp0	<p>Configurable PPPoE encapsulation interface. PPP packets being routed in an Ethernet network use PPPoE encapsulation.</p> <p>Within a Services Router, packets are routed to this internal interface for PPPoE encapsulation. The PPPoE encapsulation interface is an internal interface only and is not associated with a physical medium or PIM. You must configure the interface for it to forward PPPoE traffic. For more information about PPPoE interfaces, see “Configuring Point-to-Point Protocol over Ethernet” on page 191.</p>
sp-0/0/0	<p>Configurable services interface. The services interface is used to enable a number of routing services such as stateful firewall filters, IPSec, and Network Address Translation (NAT).</p> <p>Within a Services Router, packets are routed to this internal interface for encapsulation or processing, depending on the services configured. The configurable services interface is an internal interface only and is not associated with a physical medium or PIM. You must configure the interface for it to enable service sets.</p>
tap	Internally generated interface. This interface is generated by the JUNOS software to monitor and record traffic during passive monitoring. When packets are discarded by the Packet Forwarding Engine, they are placed on this interface. It is not a configurable interface.

Table 27: Special Interfaces on a Services Router (*continued*)

Interface Name	Description
umd0	Configurable USB modem physical interface. This interface is detected when an USB modem is connected to the USB port on the Services Router. NOTE: The J4350 and J6350 Services Routers have two USB ports. However, you can connect only one USB modem to the USB ports on these routers. If you connect USB modems to both the USB ports, only the first USB modem connected to the router is recognized.

Discard Interface

The discard (**dsc**) interface is not a physical interface, but a virtual interface that discards packets. You can configure one discard interface. This interface allows you to identify the ingress (inbound) 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. Traffic routed out the discard interface is silently discarded.

Loopback Interface

The Internet Protocol (IP) specifies a loopback network with the (IPv4) address **127.0.0.0/8**. Most IP implementations support a loopback interface (**lo0**) to represent the loopback facility. Any traffic that a computer program sends on the loopback network is addressed to the same computer. The most commonly used IP address on the loopback network is **127.0.0.1** for IPv4 and **::1** for IPv6. The standard domain name for the address is **localhost**.

The loopback interface can perform the following functions:

- Router identification—The loopback interface is used to identify the router. While any interface address can be used to determine if the router is online, the loopback address is the preferred method. Whereas interfaces might be removed or addresses changed based on network topology changes, the loopback address *never changes*.

When you ping an individual interface address, the results do not always indicate the health of the router. For example, a subnet mismatch in the configuration of two endpoints on a point-to-point link makes the link appear to be inoperable. Pinging the interface to determine whether the router is online provides a misleading result. An interface might be unavailable because of a problem unrelated to the router's configuration or operation.

- Routing information—The loopback address is used by protocols such as OSPF to determine protocol-specific properties for the router or network. Further, some commands such as **ping mpls** require a loopback address to function correctly.
- Packet filtering—Stateless firewall filters can be applied to the loopback address to filter packets originating from, or destined for, the Routing Engine.

Management Interface

The management interfaces (also called the out-of-band management interfaces) on a J-series Services Router are located on the front panel of the router chassis. The number and type of management interfaces depend on the model of the Services Router as follows:

- On J2300, J4300, and J6300 routers—Two Fast Ethernet interfaces designated `fe-0/0/0` and `fe-0/0/1` from left to right
- On J4350 and J6350 routers—Four Gigabit Ethernet interfaces designated `ge-0/0/0`, `ge-0/0/1`, `ge-0/0/2`, and `ge-0/0/3` from left to right

The management interfaces are the primary interfaces for accessing the router remotely. Typically, a management interface is not connected to the in-band network, but is connected instead to the router's internal network. Through a management interface you can access the router over the network and configure it from anywhere, regardless of its physical location.

As a security feature, users cannot log in as `root` through a management interface. To access the router as `root`, you must use the console port.

Services Interfaces

On Juniper Networks M-series and T-series routing platforms, individual services such as IP-over-IP encapsulation, link services such as multilink protocols, adaptive services such as stateful firewall filters and NAT, and sampling and logging capabilities are implemented by services Physical Interface Cards (PICs). On a J-series Services Router, these same features are implemented by the general-purpose CPU on the main circuit board.

Although the same JUNOS Internet software image supports the services features across all routing platforms, on a Services Router no Physical Interface Module (PIM) is associated with services features.

To configure services on a Services Router, you must configure one or more internal interfaces by specifying PIM slot 0 and port 0—for example, `sp-0/0/0` for stateful firewall filters and NAT or `gr-0/0/0` for GRE.

Services Routers support multilink protocol services on the `ls-0/0/0` interface. At the logical level, the `ls-0/0/0` interface supports the Multilink Point-to-Point Protocol (MLPPP) and Multilink Frame Relay (MLFR) FRF.15 encapsulation types, and at the physical level, the interface supports the MLRF FRF.16 encapsulation type and Compressed Real-Time Transport Protocol (CRTP).

MLPPP and MLFR

Multilink Point-to-Point Protocol (MLPPP) is a protocol for aggregating multiple constituent links into one larger PPP bundle. Multilink Frame Relay (MLFR) allows you to aggregate multiple Frame Relay links by inverse multiplexing. MLPPP and MLFR provide service options between low-speed T1 and E1 services. In addition to providing additional bandwidth, bundling multiple links can add a level of fault tolerance to your dedicated access service. Because you can implement bundling

across multiple interfaces, you can protect users against loss of access when a single interface fails.

MLFR Frame Relay Forum

JUNOS supports FRF.12 fragmentation header formats for both FRF.15 (MLFR) and FRF.16 (MFR).

MLFR Frame Relay Forum 15 (FRF.15) combines multiple permanent virtual circuits (PVCs) into one aggregated virtual circuit (AVC). This process provides fragmentation over multiple PVCs on one end and reassembly of the AVC on the other end. MLFR FRF.15 is supported on the `ls-0/0/0` interface.

MLFR FRF.16 is supported on the `ls-0/0/0:channel`, which carries a single MLFR FRF.16 bundle. MLFR FRF.16 combines multiple links to form one logical link. Packet fragmentation and reassembly occur on each virtual circuit. Each bundle can support multiple virtual circuits.



NOTE: If you configure a permanent virtual circuit (PVC) between T1, E1, T3, or E3 interfaces in J-series router and another vendor, and the other vendor does not have the same FRF.12 support or supports FRF.12 in a different way, the J-series interface might discard a fragmented packet containing FRF.12 headers and count it as a "Policed Discard." Therefore, when you configure a PVC between T1, E1, T3, or E3 interfaces in J-series router and another vendor, you should configure multilink bundles on both peers and configure fragmentation thresholds on the multilink bundle.

CRTP

Real-Time Transport Protocol (RTP) can help achieve interoperability among different implementations of network audio and video applications. However, the header can be too large a payload for networks using low-speed lines such as dial-up modems. Compressed Real-Time Transport Protocol (CRTP) can reduce network overhead on a low-speed link. On a Services Router, CRTP can operate on a T1 or E1 interface with PPP encapsulation.

Chapter 3

Configuring Ethernet, DS1, DS3, and Serial Interfaces

Each Services Router supports multiple types of interfaces that perform different functions. The router uses DS1, DS3, Fast Ethernet, Gigabit Ethernet, and serial network interfaces to transmit and receive network traffic. For network interfaces to operate, you must configure properties such as logical interfaces, the encapsulation type, and certain settings specific to the interface type.

In most cases, you can use either J-Web Quick Configuration or a configuration editor to configure network interfaces.



NOTE: You cannot configure channelized T1 or E1 interfaces through a J-Web Quick Configuration page. You must use the J-Web or CLI configuration editor. Even after configuration, channelized interfaces do not appear on the Quick Configuration Interfaces page.

This chapter includes the following topics. For more information about interfaces, see “Interfaces Overview” on page 41 and the *JUNOS Network Interfaces Configuration Guide*. To configure channelized interfaces, see “Configuring Channelized T1/E1/ISDN PRI Interfaces” on page 139. To configure DSL interfaces, see “Configuring Digital Subscriber Line Interfaces” on page 157. To configure PPPoE interfaces, see “Configuring Point-to-Point Protocol over Ethernet” on page 191. To configure ISDN interfaces, see “Configuring ISDN” on page 211.

- Before You Begin on page 103
- Configuring DS1, DS3, Ethernet, and Serial Interfaces with Quick Configuration on page 104
- Configuring Network Interfaces with a Configuration Editor on page 131
- Verifying Interface Configuration on page 135

Before You Begin

Before you configure network interfaces, you need to perform the following tasks:

- Install Services Router hardware. For more information, see the Getting Started Guide for your router.

- Establish basic connectivity. For more information, see the Getting Started Guide for your router.
- If you do not already have a basic understanding of physical and logical interfaces and Juniper Networks interface conventions, read “Interfaces Overview” on page 41.

Although it is not a requirement, you might also want to plan how you are going to use the various network interfaces before you start configuring them. You can see a list of the physical interfaces installed on the J-series Services Router by displaying the Quick Configuration page, as shown in Figure 22 on page 104.

Configuring DS1, DS3, Ethernet, and Serial Interfaces with Quick Configuration

The Quick Configuration page allows you to configure most network interfaces on a Services Router, as shown in Figure 22 on page 104.

For information about interface names, see “Network Interface Naming” on page 47.

Figure 22: Quick Configuration Interfaces Page

Router - J6300

Monitor Configuration Diagnose Manage Events Alarms Logged in as: regress Help About Logout

Configuration > Quick Configuration > Interfaces

Quick Configuration

Interfaces

Interface Name	Link State	Configured	Description
fe-0/0/0	Up	Yes	Fast Ethernet Interface 'fe-0/0/0'
fe-0/0/0.0	Up	Yes	Logical Unit 0 on Fast Ethernet Interface 'fe-0/0/0'
fe-0/0/1	Down	No	Fast Ethernet Interface 'fe-0/0/1'
se-1/0/0	Down	No	Other Interface 'se-1/0/0'
se-1/0/1	Down	No	Other Interface 'se-1/0/1'
fe-3/0/0	Up	No	Fast Ethernet Interface 'fe-3/0/0'
fe-3/0/1	Up	No	Fast Ethernet Interface 'fe-3/0/1'
t1-4/0/0	Up	Yes	T1 Interface 't1-4/0/0'
t1-4/0/0.0	Up	Yes	Logical Unit 0 on T1 Interface 't1-4/0/0'
t1-4/0/1	Up	No	T1 Interface 't1-4/0/1'
e1-5/0/0	Up	No	E1 Interface 'e1-5/0/0'
e1-5/0/1	Down	No	E1 Interface 'e1-5/0/1'
t1-6/0/0	Up	No	T1 Interface 't1-6/0/0'
t1-6/0/1	Down	No	T1 Interface 't1-6/0/1'
lo0	Up	Yes	Loopback Interface 'lo0'
lo0.0	Up	Yes	Logical Unit 0 on Loopback Interface 'lo0'
pp0	Up	No	Point-to-Point Protocol over Ethernet Interface 'pp0'

OK Cancel Apply

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. Trademark Notice. Privacy. Juniper Your Net.

To configure a network interface with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Quick Configuration > Interfaces**.

A list of the network interfaces present on the Services Router is displayed, as shown in Figure 22 on page 104. (For information about interface names, see “Network Interface Naming” on page 47.) The third column indicates whether the interface has been configured.



NOTE: Channelized T1 and E1 interfaces are not displayed in the list of interfaces on the J-Web Quick Configuration Interfaces page. However, you can configure and view channelized T1/ interfaces with the J-Web configuration editor. For details, see “Configuring Channelized T1/E1/ISDN PRI Interfaces” on page 139.

2. To configure properties for a network interface, select the interface name and proceed with configuration as described in one of the following topics:
 - Configuring an E1 Interface with Quick Configuration on page 105
 - Configuring an E3 Interface with Quick Configuration on page 109
 - Configuring a Fast Ethernet Interface with Quick Configuration on page 112
 - Configuring Gigabit Ethernet Interfaces with Quick Configuration on page 115
 - Configuring T1 Interfaces with Quick Configuration on page 120
 - Configuring T3 Interfaces with Quick Configuration on page 124
 - Configuring Serial Interfaces with Quick Configuration on page 127

Configuring an E1 Interface with Quick Configuration

To configure properties on an E1 interface:

1. From the Quick Configuration page, as shown in Figure 22 on page 104, select the E1 interface you want to configure.

The properties you can configure on an E1 interface are displayed, as shown in Figure 23 on page 106. (For information about interface names, see “Network Interface Naming” on page 47.)

Figure 23: E1 Interfaces Quick Configuration Page

Juniper
NETWORKS

ROUTER - J6300

Monitor Configuration Diagnose Manage Events Alarms Logged in as: regress Help About Logout

Quick Configuration View and Edit History Rescue

Configuration > Quick Configuration > Interfaces

Quick Configuration

Interfaces Physical Interface: 'e1-5/0/0'

Logical Interfaces

No logical interfaces configured.

Add...

Physical Interface Description

Physical Interface Description

MTU (bytes) ?

Clocking (internal) ?

Per Unit Scheduler ☐ ?

Encapsulation

Encapsulation

Enable CHAP ☐

CHAP Local Identity

Use System Host Name ☒

Local Name

• CHAP Peer Identity

• CHAP Secret

E1 Options

Framing Mode (9704) ?

Invert Data ☐ ?

Timeslots (1-24)

Frame Checksum (16) ?

OK Cancel Apply

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. Trademark Notice. Privacy. Juniper Your Net.

2. Enter information into the Quick Configuration page, as described in Table 28 on page 107.
3. Click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration page, click **Apply**.

- To apply the configuration and return to the main configuration page, click **OK**.
 - To cancel your entries and return to the main page, click **Cancel**.
4. To verify that the E1 interface is configured correctly, see “Verifying Interface Configuration” on page 135.

Table 28: E1 Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
Add logical interfaces	Defines one or more logical units that you connect to this physical E1 interface. You must define at least one logical unit for an E1 interface. You can define multiple units if the encapsulation type is Frame Relay.	Click Add .
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	<ol style="list-style-type: none"> 1. Type one or more IPv4 addresses and prefixes. For example: 10.10.10.10/24 2. Click Add. 3. Click OK.
Physical Interface Description	(Optional) Adds supplementary information about the physical E1 interface.	Type a text description of the E1 interface to more clearly identify it in monitoring displays.
MTU (bytes)	Specifies the maximum transmission unit size for the E1 interface.	Type a value between 256 and 9192 bytes. The default MTU for E1 interfaces is 1504.
Clocking	Specifies the transmit clock source for the E1 line.	From the list, select one of the following: <ul style="list-style-type: none"> ■ internal—Services Router's own system clock (the default) ■ external—Clock received from the E1 interface
Per unit scheduler	Enables scheduling on logical interfaces. Allows you to configure multiple output queues on a logical interface and associate an output scheduler and shaping rate with the queues.	<ul style="list-style-type: none"> ■ To enable scheduling, select the check box. ■ To disable scheduling, clear the check box.
Encapsulation		

Table 28: E1 Quick Configuration Summary *(continued)*

Field	Function	Your Action
Encapsulation	Specifies the encapsulation type for traffic on the interface.	From the list, select the encapsulation for this E1 interface: <ul style="list-style-type: none"> ■ PPP ■ Frame Relay ■ Cisco HDLC
Enable CHAP	Enables or disables CHAP authentication on an E1 interface with PPP encapsulation only.	<ul style="list-style-type: none"> ■ To enable CHAP, select the check box. ■ To disable CHAP, clear the check box.
CHAP Local Identity (available if CHAP is enabled)		
Use System Host Name	Specifies that the E1 interface uses the Services Router's system hostname in CHAP challenge and response packets.	<ul style="list-style-type: none"> ■ To enable, select the check box (the default). ■ To disable, clear the check box.
Local Name	If Use System Host Name is disabled, specifies the local name for CHAP to use.	Type a local name for this E1 interface.
CHAP Peer Identity	Identifies the client or peer with which the Services Router communicates on this E1 interface.	Type the CHAP client name.
CHAP Secret	Specifies the secret password for CHAP authentication, known to both sides of the connection.	Type a password that is known to the other side of the connection. Use a combination of letters and numbers that is difficult for others to guess.
E1 Options		
Framing Mode	Specifies the framing mode for the E1 line.	From the list, select one of the following: <ul style="list-style-type: none"> ■ g704—The default ■ g704-no-crc4—G704 without cyclic redundancy check 4 (CRC4) ■ unframed—Unframed transmission format
Invert Data	Enables or disables data inversion. Data inversion is normally used only in alternate mark inversion (AMI) mode.	<ul style="list-style-type: none"> ■ To enable, select the check box. ■ To disable, clear the check box.
Timeslots	Specifies the number of time slots allocated to a fractional E1 interface. By default, an E1 interface uses all the time slots.	Type numeric values from 2 through 32. Separate discontinuous entries with commas, and use hyphens to indicate ranges. For example: 2,4,7–9
Frame Checksum	Specifies the number of bits in the frame checksum. A 32-bit checksum provides more reliable packet verification, but is not supported by some older equipment.	Select 16 or 32 . The default checksum is 16.

Configuring an E3 Interface with Quick Configuration

To configure properties on an E3 interface:

1. From the Quick Configuration page, as shown in Figure 22 on page 104, select the interface you want to configure.

The properties you can configure on an E3 interface are displayed, as shown in Figure 24 on page 109. (For information about interface names, see “Network Interface Naming” on page 47.)

Figure 24: E3 Interfaces Quick Configuration Page

The screenshot shows the Juniper Networks Router - J6300 Quick Configuration page. The left sidebar contains navigation links: Quick Configuration, View and Edit, History, and Rescue. The main content area is titled 'Quick Configuration' and 'Interfaces'. The physical interface is 'e3-1/0/0'. The configuration options for the E3 interface are as follows:

Option	Value	Help
Bert Algorithm	(pseudo-2e15-o151)	?
Bert Error Rate	(3)	?
Bert Period	(10)	?
Compatibility Mode	Off	
Digital-Link	Subrate	?
Kentrox	Subrate	?
Frame Checksum	(16)	?
Idle Cycle Flag	(flags)	?
Loopback		?
Payload Scrambler	Yes No	?
Start End Flag	(filler)	?
Unframed	Yes No	?

Buttons: OK, Cancel, Apply

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. Trademark Notice. Privacy. Juniper Your Net.

2. Enter information into the Quick Configuration page, as described in Table 29 on page 110.
3. Click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
 - To apply the configuration and return to the main configuration page, click **OK**.

- To cancel your entries and return to the main page, click **Cancel**.
4. To verify that the E3 interface is configured correctly, see “Verifying Interface Configuration” on page 135.

Table 29: E3 Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
Add logical interfaces	Defines one or more logical units that you connect to this physical E3 interface. You must define at least one logical unit for an E3 interface. You can define multiple units if the encapsulation type is Frame Relay.	Click Add .
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	<ol style="list-style-type: none"> 1. Type one or more IPv4 addresses and prefixes. For example: 10.10.10.10/24 2. Click Add. 3. Click OK.
Physical Interface Description	(Optional) Adds supplementary information about the physical E3 interface.	Type a text description of the E3 interface to more clearly identify it in monitoring displays.
MTU (bytes)	Specifies the maximum transmission unit size for the E3 interface.	Type a value between 256 and 9192 bytes. The default MTU for E3 interfaces is 4474 .
Clocking	Specifies the transmit clock source for the E3 line.	From the list, select one of the following: <ul style="list-style-type: none"> ■ internal—Services Router's own system clock (the default) ■ external—Clock received from the E3 interface
Encapsulation		
Encapsulation	Specifies the encapsulation type for traffic on the interface.	From the list, select the encapsulation for this E3 interface: <ul style="list-style-type: none"> ■ PPP ■ Frame Relay ■ Cisco HDLC
Enable CHAP	Enables or disables CHAP authentication on an E3 interface with PPP encapsulation only.	<ul style="list-style-type: none"> ■ To enable CHAP, select the check box. ■ To disable CHAP, clear the check box.
CHAP Local Identity (available if CHAP is enabled)		
Use System Host Name	Specifies that the E3 interface uses the Services Router's system hostname in CHAP challenge and response packets.	<ul style="list-style-type: none"> ■ To enable, select the check box (the default). ■ To disable, clear the check box.

Table 29: E3 Quick Configuration Summary (*continued*)

Field	Function	Your Action
Local Name	If Use System Host Name is disabled, specifies the local name for CHAP to use.	Type a local name for this E3 interface.
CHAP Peer Identity	Identifies the client or peer with which the Services Router communicates on this E3 interface.	Type the CHAP client name.
CHAP Secret	Specifies the secret password for CHAP authentication, known to both sides of the connection.	Type a password that is known to the other side of the connection. Use a combination of letters and numbers that is difficult for others to guess.
E3 Options		
Bert Algorithm	<p>Specifies the bit error rate test (BERT) algorithm to use during a BERT.</p> <p>BERT is supported only when transmission is unframed. (See the Unframed option.)</p>	<p>From the Bert Algorithm list, select the algorithm to use:</p> <ul style="list-style-type: none"> ■ all-ones-repeating ■ alternating-ones-zeros ■ all-zeros-repeating ■ pseudo-2e11-o152 ■ pseudo-2e15-o151 ■ pseudo-2e20-o151 ■ pseudo-2e20-o153 ■ pseudo-2e23-o151 ■ pseudo-2e29 ■ pseudo-2e31 ■ pseudo-2e9-o153 <p>The default is pseudo-2e15-o151.</p>
Bert Error Rate	Specifies the exponent n in the bit error rate 10^{-n} .	Type a value between 3 and 7, or 0. For example, a value of 6 specifies that 1 bit out of 1,000,000 is transmitted in error. The default is 0 (no bits are transmitted in error).
Bert Period	Specifies the length of time—in seconds—of the BERT.	Type a value between 1 and 240. The default is 10.
Compatibility Mode	<p>Defines the transmission mode and subrating to use on the E3 interface. The mode must be set to the type of channel service unit (CSU) connected to the interface. The subrating specified must be the same subrating configured on the CSU.</p> <p>CSU compatibility mode and subrating are supported only when transmission is unframed. (See the Unframed option.)</p>	<p>Select one of the following check boxes:</p> <ul style="list-style-type: none"> ■ Off—CSU compatibility is disabled. ■ Digital-Link—Compatible with a Digital Link CSU. ■ Kentrox—Compatible with a Kentrox CSU. <p>If you select Digital-Link, you can optionally specify a subrate by selecting a value from the Subrate list.</p> <p>If you select Kentrox, you can optionally specify a subrate by typing a value from 1 through 48 in the Subrate box.</p> <p>If you do not specify a subrate, the full E3 rate is used.</p>

Table 29: E3 Quick Configuration Summary (*continued*)

Field	Function	Your Action
Frame Checksum	Specifies the number of bits in the frame checksum. A 32-bit checksum provides more reliable packet verification, but is not supported by some older equipment.	From the Frame Checksum list, select 16 or 32 . The default value is 16 .
Idle Cycle Flag	Specifies the value to transmit during idle cycles.	From the Idle Cycle Flag list, select one of the following: <ul style="list-style-type: none"> ■ flags—Transmits the value 0x7E during idle cycles. This is the default. ■ ones—Transmits the value 0xFF during idle cycles.
Loopback	<p>Configures the E3 interface as a loopback interface for testing purposes.</p> <p>When E3 is configured as a local loopback interface, the router transmits test traffic simultaneously to the CSU and to the receiver at the E3 interface.</p> <p>When E3 is configured as a remote loopback interface, test traffic transmitted by the CSU is simultaneously received at the E3 interface and transmitted back to the CSU.</p>	From the Loopback list, select one of the following: <ul style="list-style-type: none"> ■ local—Traffic loops from the transmitter to the receiver at the E3 interface during tests. ■ remote—Traffic loops from the receiver to the transmitter at the E3 interface during tests.
Payload Scrambler	<p>Specifies whether the payload of the packet is to be scrambled, or randomized, when transmitted. Scrambling eliminates nonvariable bit patterns in the transmission, which can generate link-layer errors across an E3 link.</p> <p>The payload scrambler is supported only when CSU compatibility is enabled and transmission is framed. (See the Compatibility Mode and Unframed options).</p>	<p>Select one of the following check boxes:</p> <ul style="list-style-type: none"> ■ Yes—Transmission is scrambled. ■ No—Transmission is not scrambled.
Start End Flag	Specifies whether the end and start flags are separated.	From the Start End Flag list, select one of the following: <ul style="list-style-type: none"> ■ filler—Flags are separated by idle cycles. ■ shared—Flags overlap (no separation).
Unframed	Specifies whether the transmission is framed (G.751 framing) or unframed.	<p>Select one of the following check boxes:</p> <ul style="list-style-type: none"> ■ Yes—Unframed transmission. ■ No—Framed transmission.

Configuring a Fast Ethernet Interface with Quick Configuration

To configure properties on a Fast Ethernet interface:

1. From the Quick Configuration page, as shown in Figure 22 on page 104, select the interface you want to configure.

The properties you can configure on a Fast Ethernet interface are displayed, as shown in Figure 25 on page 113. (For information about interface names, see “Network Interface Naming” on page 47.)

Figure 25: Fast Ethernet Interfaces Quick Configuration Page

The screenshot shows the Juniper Networks configuration interface for a J6300 router. The 'Configuration' tab is active, and the 'Quick Configuration' page is displayed. The 'Physical Interface' is 'fe-0/0/0'. Below this, the 'Logical Interfaces' table is shown with one entry: 'fe-0/0/0' with a link state of 'Up' and a description of 'Logical Unit 0 on Fast Ethernet Interface 'fe-0/0/0''. At the bottom, there are input fields for 'Physical Interface Description', 'MTU (bytes)', and 'Per Unit Scheduler', along with 'OK', 'Cancel', and 'Apply' buttons.

Logical Interface Name	Link State	Configured	Description
<input type="checkbox"/> fe-0/0/0	Up	Yes	Logical Unit 0 on Fast Ethernet Interface 'fe-0/0/0'

Buttons: Add..., Delete, OK, Cancel, Apply

2. Enter information into the Quick Configuration page, as described in Table 30 on page 113.
3. Click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
 - To apply the configuration and return to the main configuration page, click **OK**.
 - To cancel your entries and return to the main page, click **Cancel**.
4. To verify that the Fast Ethernet interface is configured correctly, see “Verifying Interface Configuration” on page 135.

Table 30: Fast Ethernet Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		

Table 30: Fast Ethernet Quick Configuration Summary *(continued)*

Field	Function	Your Action
Add logical interfaces	Defines one or more logical units that you connect to this physical Fast Ethernet interface. You must define at least one logical unit for a Fast Ethernet interface. You can define multiple units if the encapsulation type is Frame Relay.	Click Add .
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	<ol style="list-style-type: none"> 1. Type one or more IPv4 addresses and prefixes. For example: 10.10.10.10/24 2. Click Add. 3. Click OK.
ARP Address	<p>Enables the router to create a static Address Resolution Protocol (ARP) entry for this interface by specifying the IP address of a node to associate with its media access control (MAC) address. The IP address must be in the same subnet as the IPv4 address or prefix of the interface you are configuring.</p> <p>Static ARP entries associate the IP addresses and MAC addresses of nodes on the same subnet, enabling a Services Router to respond to ARP requests having destination addresses that are not local to the incoming interface.</p>	Type an IPv4 address that you want to associate with the MAC address—for example, 10.10.10.1.
MAC Address	<p>Specifies the hardware media access control (MAC) address associated with the ARP address.</p> <p>The MAC address uniquely identifies the system and is expressed in the following format: mm:mm:mm:ss:ss:ss. The first three octets denote the hardware manufacturer ID, and the last three are serial numbers identifying the router.</p>	Type the MAC address to be mapped to the ARP entry—for example, 00:12:1E:A9:8A:80.
Publish	<p>Enables the router to reply to ARP requests for the specified address.</p> <p>For more information, see “Configuring Static ARP Entries on Ethernet Interfaces” on page 133.</p>	<ul style="list-style-type: none"> ■ To enable publishing, select the check box. ■ To disable publishing, clear the check box.

Table 30: Fast Ethernet Quick Configuration Summary (*continued*)

Field	Function	Your Action
Physical Interface Description	(Optional) Adds supplementary information about the physical Fast Ethernet interface.	Type a text description of the Fast Ethernet interface to more clearly identify it in monitoring displays.
MTU (bytes)	Specifies the maximum transmission unit size for the Fast Ethernet interface.	<p>Type a value between 256 bytes and one of the following values:</p> <ul style="list-style-type: none"> ■ For built-in Fast Ethernet interfaces and Dual-Port Fast Ethernet PIM interfaces, 9192 bytes ■ For 4-Port Fast Ethernet ePIM interfaces, 1514 bytes <p>The default MTU for Fast Ethernet interfaces is 1514.</p>
Per unit scheduler	<p>Enables scheduling on logical interfaces.</p> <p>Allows you to configure multiple output queues on a logical interface and associate an output scheduler and shaping rate with the queues.</p>	<ul style="list-style-type: none"> ■ To enable scheduling, select the check box. ■ To disable scheduling, clear the check box.



NOTE: You can also manually set the speed and link mode for a Fast Ethernet interface using the CLI commands `set interfaces fe-pim/0/port speed 10m | 100m` and `set interfaces fe-pim/0/port link-mode half-duplex | full-duplex`.

Configuring Gigabit Ethernet Interfaces with Quick Configuration

To configure properties on a Gigabit Ethernet interface:

1. From the Quick Configuration page, as shown in Figure 22 on page 104, select the Gigabit Ethernet interface you want to configure.

The properties you can configure on a Gigabit Ethernet interface are displayed, as shown in Figure 26 on page 116. (For information about interface names, see “Network Interface Naming” on page 47.)

Figure 26: Gigabit Ethernet Interfaces Quick Configuration Page

The screenshot displays the Juniper J-series Services Router Quick Configuration page for Gigabit Ethernet interfaces. The page is titled "ROUTER - M5" and includes a navigation bar with links for Monitor, Configuration, Diagnose, Manage, Events, Alarms, and Logout. The left sidebar shows the "Quick Configuration" menu with options for View and Edit, and History. The main content area is titled "Quick Configuration" and "Interfaces". It shows the "Physical Interface: 'ge-0/3/0'". Under "Logical Interfaces", it states "No logical interfaces configured." and provides an "Add..." button. The "Physical Interface Description" field is empty. The "MTU (bytes)" field is set to 1500. The "Per Unit Scheduler" checkbox is unchecked. The "Gigabit Ethernet Options" section includes checkboxes for Flow Control, Loopback, Source Filtering, and Auto Negotiation, all of which are unchecked. The "Auto Negotiation Remote Fault" dropdown is set to "Auto". The "Source MAC Address Filters" section has an "Add" button and a "Delete" button. The "Tag Protocol ID" field is empty. The "MAC Learning" checkbox is unchecked. The "802.3ad" section has an "Add" button and a "Delete" button. At the bottom, there are "OK", "Cancel", and "Apply" buttons.

2. Enter information into the Quick Configuration page, as described in Table 31 on page 117.
3. Click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
 - To apply the configuration and return to the main Configuration page, click **OK**.

- To cancel your entries and return to the main page, click **Cancel**.
4. To verify that the Gigabit Ethernet interface is configured correctly, see “Verifying Interface Configuration” on page 135.

Table 31: Gigabit Ethernet Quick Configuration Page Summary

Field	Function	Your Action
Logical Interfaces		
Add logical interfaces	Defines one or more logical units that you connect to this physical Gigabit Ethernet interface. You must define at least one logical unit for a Gigabit Ethernet interface.	Click Add .
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	<ol style="list-style-type: none"> 1. Type one or more IPv4 addresses and prefixes. For example: 10.10.10.10/24 2. Click Add. 3. Click OK. <p>To delete an IP address and prefix, select them in the Source Addresses and Prefixes box, then click Delete.</p>
ARP Address	<p>Enables the router to create a static Address Resolution Protocol (ARP) entry for this interface by specifying the IP address of a node to associate with its media access control (MAC) address. The IP address must be in the same subnet as the IPv4 address or prefix of the interface you are configuring.</p> <p>Static ARP entries associate the IP addresses and MAC addresses of nodes on the same subnet, enabling a Services Router to respond to ARP requests having destination addresses that are not local to the incoming interface.</p>	Type an IPv4 address that you want to associate with the MAC address—for example, 10.10.10.1.
MAC Address	<p>Specifies the hardware media access control (MAC) address associated with the ARP address.</p> <p>The MAC address uniquely identifies the system and is expressed in the following format: mm:mm:mm:ss:ss:ss. The first three octets denote the hardware manufacturer ID, and the last three are serial numbers identifying the router.</p>	Type the MAC address to be mapped to the ARP entry—for example, 00:12:1E:A9:8A:80.

Table 31: Gigabit Ethernet Quick Configuration Page Summary (*continued*)

Field	Function	Your Action
Publish	<p>Enables the router to reply to ARP requests for the specified address.</p> <p>For more information, see “Configuring Static ARP Entries on Ethernet Interfaces” on page 133.</p>	<ul style="list-style-type: none"> ■ To enable publishing, select the check box. ■ To disable publishing, clear the check box.
Physical Interface Description	(Optional) Adds supplementary information about the physical Gigabit Ethernet interface.	Type a text description of the Gigabit Ethernet interface to more clearly identify it in monitoring displays.
MTU (bytes)	Specifies the maximum transmission unit size for the Gigabit Ethernet interface.	Type a value between 256 and 9014 bytes. The default MTU for Gigabit Ethernet interfaces is 1514 .
Per unit scheduler	<p>Enables scheduling on logical interfaces.</p> <p>Allows you to configure multiple output queues on a logical interface and associate an output scheduler and shaping rate with the queues.</p>	<ul style="list-style-type: none"> ■ To enable scheduling, select the check box. ■ To disable scheduling, clear the check box.
Gigabit Ethernet Options		
Flow Control	Enables or disables flow control.	Select Yes to enable flow control to regulate the amount of traffic sent out of the interface, or select No to disable flow control and permit unrestricted traffic. Flow control is disabled by default.
Loopback	Enables or disables the loopback option.	Select Yes to enable the loopback diagnostic option, or select No to disable the loopback option. By default, loopback is disabled.
Source Filtering	<p>Enables or disables source filtering. Enabling source filtering blocks all incoming packets to the interface.</p> <p>For more information, see the <i>JUNOS Network Interfaces Configuration Guide</i></p>	<p>Select Yes to enable source filtering, or select No to disable the source filtering. By default, source filtering is disabled.</p> <p>NOTE: If you disable source filtering, the source MAC address filtering is also disabled.</p>

Table 31: Gigabit Ethernet Quick Configuration Page Summary (*continued*)

Field	Function	Your Action
Auto Negotiation	<p>Enables or disables autonegotiation.</p> <p>By default, Gigabit Ethernet interfaces autonegotiate the link mode and speed settings. If you disable autonegotiation and do not manually configure link mode and speed, the link is negotiated at 1000 Mbps, full duplex.</p> <p>When you configure both the link mode and the speed, the link negotiates with the manually configured settings whether autonegotiation is enabled or disabled.</p>	<p>Select Yes to enable autonegotiation, or select No to disable it. By default, autonegotiation is enabled.</p>
Auto Negotiation Remote Fault	<p>Indicates the autonegotiation remote fault value.</p>	<p>Select the autonegotiation remote fault value from the list of options given. This field is enabled only if autonegotiation is enabled.</p>
Source MAC Address Filters	<p>Displays the list of media access control (MAC) addresses from which you want to receive packets on this interface.</p> <p>NOTE: To enable source MAC address filtering, first enable source filtering by selecting the Yes check box next to Source Filtering.</p>	<p>To add MAC addresses, type them in the boxes above the Add button, then click Add.</p> <p>To delete a MAC address, select it in the Source Addresses box, then click Delete.</p>
802.3ad	<p>Specifies a “bundle” of Gigabit Ethernet interfaces on this router with which to share traffic.</p> <p>To use this feature, you must already have configured an aggregate Ethernet interface, 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.</p>	<p>Type an aggregated Ethernet interface value—for example, ae0. Aggregated Ethernet interface names range from ae0 through ae15.</p> <p>NOTE: The J-Web interface displays error messages if you enter an incorrect value.</p> <p>For more information, see the <i>JUNOS Network Interfaces Configuration Guide</i>.</p>
Tag Protocol ID	<p>Displays a list of IEEE 802.1Q Tag Protocol Identifier (TPID) values. The TPIDs identify frames that are to be processed as virtual LAN (VLAN)-tagged frames.</p> <p>To use this feature, you must already have enabled VLAN tagging. VLAN tags enable you to channelize an Ethernet interface into multiple logical interfaces, allowing many hosts to be connected to the same Gigabit Ethernet switch.</p>	<p>To add a TPID, type it in the boxes above the Add button, then click Add.</p> <p>To delete a TPID, select it in the box, then click Delete.</p> <p>Each Gigabit Ethernet port can have up to eight TPIDs.</p> <p>For more information about TPIDs and VLAN tagging, see the <i>JUNOS Network Interfaces Configuration Guide</i>.</p>

Table 31: Gigabit Ethernet Quick Configuration Page Summary (*continued*)

Field	Function	Your Action
MAC Learning	Enables or disables source and destination MAC address learning dynamically.	Select Yes to enable dynamic MAC address learning, or select No to disable it. By default, dynamic MAC address learning is disabled.



NOTE: You can also manually set the speed and link mode for built-in and copper PIM Gigabit Ethernet interfaces on J4350 and J6350 routers using the CLI commands `set interfaces ge-pim/0/port speed 10m | 100m | 1000m` and `set interfaces ge-pim/0/port link-mode half-duplex | full-duplex`. (You cannot manually configure speed and link mode on SFP Gigabit Ethernet PIMs.) You must configure both link mode and speed—if you configure only one or the other, the system ignores the configuration and generates a system log message.

Configuring T1 Interfaces with Quick Configuration

To configure properties on a T1 interface:

1. From the Quick Configuration page, as shown in Figure 22 on page 104, select the interface you want to configure.

The properties you can configure on a T1 interface are displayed, as shown in Figure 27 on page 121. (For information about interface names, see “Network Interface Naming” on page 47.)

Figure 27: T1 Interfaces Quick Configuration Page

Juniper NETWORKS ROUTER - J6300

Monitor Configuration Diagnose Manage Events Alarms Logged in as: regress Help About Logout

Configuration > Quick Configuration > Interfaces

Quick Configuration

Interfaces Physical Interface: 't1-4/0/0'

Logical Interfaces

	Logical Interface Name	Link State	Configured	Description
<input type="checkbox"/>	t1-4/0/0	Up	Yes	Logical Unit 0 on T1 Interface 't1-4/0/0'

Add... Delete

Physical Interface Description

MTU (bytes) ?

Clocking (internal) ?

Per Unit Scheduler ☐ ?

Encapsulation

Encapsulation

Enable CHAP ☐

CHAP Local Identity

Use System Host Name ☒

Local Name

* CHAP Peer Identity

* CHAP Secret

T1 Options

Framing Mode (esf) ?

Line Encoding (b8zs) ?

Byte Encoding (nx64) ?

Invert Data ☐ ?

Timeslots ? (1-24)

Frame Checksum (16) ?

Line Buildout (0-132) ?

OK Cancel Apply

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. Trademark Notice. Privacy. Juniper Your Net.

2. Enter information into the Quick Configuration page, as described in Table 32 on page 122.

3. Click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
 - To apply the configuration and return to the main configuration page, click **OK**.
 - To cancel your entries and return to the main page, click **Cancel**.
4. To verify that the T1 interface is configured correctly, see “Verifying Interface Configuration” on page 135.

Table 32: T1 Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
Add logical interfaces	Defines one or more logical units that you connect to this physical T1 interface. You must define at least one logical unit for a T1 interface. You can define multiple units if the encapsulation type is Frame Relay.	Click Add .
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	<ol style="list-style-type: none"> 1. Type one or more IPv4 addresses and prefixes. For example: 10.10.10.10/24 2. Click Add. 3. Click OK.
Physical Interface Description	(Optional) Adds supplementary information about the physical T1 interface.	Type a text description of the T1 interface to more clearly identify it in monitoring displays.
MTU (bytes)	Specifies the maximum transmission unit size for the T1 interface.	Type a value between 256 and 9192 bytes. The default MTU for T1 interfaces is 1504.
Clocking	Specifies the transmit clock source for the T1 line.	From the list, select one of the following: <ul style="list-style-type: none"> ■ internal—Services Router's own system clock (the default) ■ external—Clock received from the T1 interface
Per unit scheduler	Enables scheduling on logical interfaces. Allows you to configure multiple output queues on a logical interface and associate an output scheduler and shaping rate with the queues.	<ul style="list-style-type: none"> ■ To enable scheduling, select the check box. ■ To disable scheduling, clear the check box.
Encapsulation		

Table 32: T1 Quick Configuration Summary (continued)

Field	Function	Your Action
Encapsulation	Specifies the encapsulation type for traffic on the interface.	From the list, select the encapsulation for this T1 interface: <ul style="list-style-type: none"> ■ PPP ■ Frame Relay ■ Cisco HDLC
Enable CHAP	Enables or disables CHAP authentication on a T1 interface with PPP encapsulation only.	<ul style="list-style-type: none"> ■ To enable CHAP, select the check box. ■ To disable CHAP, clear the check box.
CHAP Local Identity (available if CHAP is enabled)		
Use System Host Name	Specifies that the T1 interface uses the Services Router's system hostname in CHAP challenge and response packets.	<ul style="list-style-type: none"> ■ To enable, select the check box (the default). ■ To disable, clear the check box.
Local Name	If Use System Host Name is disabled, specifies the local name for CHAP to use.	Type a local name for this T1 interface.
CHAP Peer Identity	Identifies the client or peer with which the Services Router communicates on this T1 interface.	Type the CHAP client name.
CHAP Secret	Specifies the secret password for CHAP authentication, known to both sides of the connection.	Type a password that is known to the other side of the connection. Use a combination of letters and numbers that is difficult for others to guess.
T1 Options		
Framing Mode	Specifies the framing mode for the T1 line.	From the list, select one of the following: <ul style="list-style-type: none"> ■ esf—Extended superframe (the default) ■ sf—Superframe
Line Encoding	Specifies the line encoding method.	From the list, select one of the following: <ul style="list-style-type: none"> ■ ami—Alternate mark inversion ■ b8zs—Binary 8 zero substitution (the default)
Byte Encoding	Specifies the byte encoding method.	From the list, select one of the following: <ul style="list-style-type: none"> ■ nx56—7 bits per byte ■ nx64—8 bits per byte (the default)
Invert Data	Enables or disables data inversion. Data inversion is normally used only in alternate mark inversion (AMI) mode.	<ul style="list-style-type: none"> ■ To enable, select the check box. ■ To disable, clear the check box.
Timeslots	Specifies the number of time slots allocated to a fractional T1 interface. By default, a T1 interface uses all the time slots.	Type numeric values from 1 through 24 . You can use any combination of time slots. To configure ranges, use hyphens. To configure discontinuous slots, use commas. For example: 1-5,10,24

Table 32: T1 Quick Configuration Summary *(continued)*

Field	Function	Your Action
Frame Checksum	Specifies the number of bits in the frame checksum. A 32-bit checksum provides more reliable packet verification, but is not supported by some older equipment.	Select 16 or 32 . The default value is 16 .
Line Buildout	<p>Specifies the T1 line buildout in feet for cables 655 feet (200 m) or shorter, or in decibels for longer cables.</p> <p>Line buildout compensates for the loss in decibels based on the distance from the device to the first repeater in the circuit.</p>	<p>From the list, select one of the following line buildouts:</p> <ul style="list-style-type: none"> ■ 0–132 (0 m–40 m) (the default) ■ 133–265 (40 m–81 m) ■ 266–398 (81 m–121 m) ■ 399–531 (121 m–162 m) ■ 532–655 (162 m–200 m) ■ long-0db ■ long-7.5db ■ long-15db ■ long-22.5db

Configuring T3 Interfaces with Quick Configuration

To configure properties on a T3 (DS3) interface:

1. From the Quick Configuration page, as shown in Figure 22 on page 104, select the interface you want to configure.

The properties you can configure on a T3 interface are displayed, as shown in Figure 28 on page 125. (For information about interface names, see “Network Interface Naming” on page 47.)

Figure 28: T3 Interfaces Quick Configuration Page

Juniper
NETWORKS

ROUTER - J6300

Monitor **Configuration** Diagnose Manage Events Alarms Logged in as: regress Help About Logout

[Configuration](#) > [Quick Configuration](#) > [Interfaces](#)

Quick Configuration

Interfaces Physical Interface: 't3-3/0/0'

Logical Interfaces

No logical interfaces configured.

[Add...](#)

Physical Interface Description

MTU (bytes) ?

Clocking (internal) ?

Per Unit Scheduler ☐ ?

Encapsulation

Encapsulation

Enable CHAP ☐

CHAP Local Identity

Use System Host Name ☒

Local Name

CHAP Peer Identity

CHAP Secret

T3 Options

Frame Checksum (16) ?

Enable Long Buildout ☐ ?

Disable C-bit parity mode ☐ ?

[OK](#) [Cancel](#) [Apply](#)

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. Trademark Notice. Privacy. Juniper your Net.

2. Enter information into the Quick Configuration page, as described in Table 33 on page 126.
3. Click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration page, click **Apply**.

- To apply the configuration and return to the main configuration page, click **OK**.
 - To cancel your entries and return to the main page, click **Cancel**.
4. To verify that the T3 interface is configured correctly, see “Verifying Interface Configuration” on page 135.

Table 33: T3 Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
Add logical interfaces	Defines one or more logical units that you connect to this physical T3 interface. You must define at least one logical unit for a T3 interface. You can define multiple units if the encapsulation type is Frame Relay.	Click Add .
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	<ol style="list-style-type: none"> 1. Type one or more IPv4 addresses and prefixes. For example: 10.10.10.10/24 2. Click Add. 3. Click OK.
Physical Interface Description	(Optional) Adds supplementary information about the physical T3 interface.	Type a text description of the T3 interface to more clearly identify it in monitoring displays.
MTU (bytes)	Specifies the maximum transmission unit size for the T3 interface.	Type a value between 256 and 9192 bytes. The default MTU for T3 interfaces is 4474 .
Clocking	Specifies the transmit clock source for the T3 line.	From the list, select one of the following: <ul style="list-style-type: none"> ■ internal—Services Router's own system clock (the default) ■ external—Clock received from the T3 interface
Per unit scheduler	Enables scheduling on logical interfaces. Allows you to configure multiple output queues on a logical interface and associate an output scheduler and shaping rate with the queues.	<ul style="list-style-type: none"> ■ To enable scheduling, select the check box. ■ To disable scheduling, clear the check box.
Encapsulation		

Table 33: T3 Quick Configuration Summary (continued)

Field	Function	Your Action
Encapsulation	Specifies the encapsulation type for traffic on the interface.	From the list, select the encapsulation for this T3 interface: <ul style="list-style-type: none"> ■ PPP ■ Frame Relay ■ Cisco HDLC
Enable CHAP	Enables or disables CHAP authentication on a T3 interface with PPP encapsulation only.	<ul style="list-style-type: none"> ■ To enable CHAP, select the check box. ■ To disable CHAP, clear the check box.
CHAP Local Identity (available if CHAP is enabled)		
Use System Host Name	Specifies that the T3 interface uses the Services Router's system hostname in CHAP challenge and response packets.	<ul style="list-style-type: none"> ■ To enable, select the check box (the default). ■ To disable, clear the check box.
Local Name	If Use System Host Name is disabled, specifies the local name for CHAP to use.	Type a local name for this T3 interface.
CHAP Peer Identity	Identifies the client or peer with which the Services Router communicates on this T3 interface.	Type the CHAP client name.
CHAP Secret	Specifies the secret password for CHAP authentication, known to both sides of the connection.	Type a password that is known to the other side of the connection. Use a combination of letters and numbers that is difficult for others to guess.
T3 Options		
Frame Checksum	Specifies the number of bits in the frame checksum. A 32-bit checksum provides more reliable packet verification, but is not supported by some older equipment.	Select 16 or 32 . The default value is 16 .
Enable Long Buildout	Specifies a short or long cable length for copper-cable-based T3 interfaces. A long cable is longer than 225 feet (68.6m).	<ul style="list-style-type: none"> ■ To enable long buildout, select the check box. ■ To disable long buildout, clear the check box.
Disable C-Bit Parity Mode	Enables or disables C-bit parity mode, which controls the type of framing that is present on the transmitted T3 signal.	<ul style="list-style-type: none"> ■ To disable, select the check box. ■ To enable, clear the check box.

Configuring Serial Interfaces with Quick Configuration

A serial interface uses a serial line protocol—such as EIA-530, X.21, RS-449/422, RS-232, or V.35—to control the transmission of signals across the interface. You do not need to explicitly configure the serial line protocol, because it is automatically detected by a Services Router based on the cable plugged into the serial interface.

To configure properties on a serial interface:

1. From the Quick Configuration page, as shown in Figure 22 on page 104, select the interface you want to configure.

The properties you can configure on a serial interface are displayed, as shown in Figure 29 on page 128. (For information about interface names, see “Network Interface Naming” on page 47.)

Figure 29: Serial Interfaces Quick Configuration Page

The screenshot shows the Juniper J6300 Router's Quick Configuration page for Serial Interfaces. The page has a blue header with the Juniper logo and "ROUTER - J6300". Below the header is a navigation bar with tabs: Monitor, Configuration, Diagnose, Manage, Events, Alarms, and a status bar showing "Logged in as: regress". The left sidebar contains links: Quick Configuration, View and Edit, History, and Rescue. The main content area is titled "Quick Configuration" and "Interfaces". It shows the "Physical Interface: 'se-1/0/0'". Under "Logical Interfaces", it says "No logical interfaces configured." and has an "Add..." button. The "Physical Interface Description" field is empty. The "MTU (bytes)" field is empty with a help icon. The "Per Unit Scheduler" checkbox is unchecked. The "Encapsulation" dropdown is empty. The "Enable CHAP" checkbox is unchecked. Under "CHAP Local Identity", the "Use System Host Name" checkbox is checked. The "Local Name" field is empty. The "CHAP Peer Identity" and "CHAP Secret" fields are empty. Under "Serial Options", the "Clock Rate" dropdown is empty with a help icon. At the bottom are "OK", "Cancel", and "Apply" buttons. The footer contains copyright information and the Juniper logo.

2. Enter information into the Quick Configuration page, as described in Table 34 on page 129.
3. Click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
 - To apply the configuration and return to the main configuration page, click **OK**.

- To cancel your entries and return to the main page, click **Cancel**.
4. To verify that the serial interface is configured correctly, see “Verifying Interface Configuration” on page 135.

Table 34: Serial Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
Add logical interfaces	Defines one or more logical units that you connect to this physical serial interface. You must define at least one logical unit for a serial interface. You can define multiple units if the encapsulation type is Frame Relay.	Click Add .
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	<ol style="list-style-type: none"> 1. Type one or more IPv4 addresses and prefixes. For example: 10.10.10.10/24 2. Click Add. 3. Click OK.
Physical Interface Description	(Optional) Adds supplementary information about the physical serial interface.	Type a text description of the serial interface to more clearly identify it in monitoring displays.
MTU (bytes)	Specifies the maximum transmission unit size for a serial interface.	Type a value between 256 and 9192 bytes. The default MTU for serial interfaces is 1504.
Per unit scheduler	<p>Enables scheduling on logical interfaces.</p> <p>Allows you to configure multiple output queues on a logical interface and associate an output scheduler and shaping rate with the queues.</p>	<ul style="list-style-type: none"> ■ To enable scheduling, select the check box. ■ To disable scheduling, clear the check box.
Encapsulation		
Encapsulation	Specifies the encapsulation type for traffic on the interface.	<p>From the list, select the encapsulation for this serial interface:</p> <ul style="list-style-type: none"> ■ PPP ■ Frame Relay ■ Cisco HDLC
Enable CHAP	Enables or disables CHAP authentication on a serial interface with PPP encapsulation only.	<ul style="list-style-type: none"> ■ To enable CHAP, select the check box. ■ To disable CHAP, clear the check box.
CHAP Local Identity (available if CHAP is enabled)		

Table 34: Serial Quick Configuration Summary (*continued*)

Field	Function	Your Action
Use System Host Name	Specifies that the serial interface use the Services Router's system hostname in CHAP challenge and response packets.	<ul style="list-style-type: none"> ■ To enable, select the check box (the default). ■ To disable, clear the check box.
Local Name	If Use System Host Name is disabled, specifies the local name for CHAP to use.	Type a local name for this serial interface.
CHAP Peer Identity	Identifies the client or peer with which the Services Router communicates on this serial interface.	Type the CHAP client name.
CHAP Secret	Specifies the secret password for CHAP authentication, known to both sides of the connection.	Type a password that is known to the other side of the connection. Use a combination of letters and numbers that is difficult for others to guess.
Serial Options		
Clocking Mode	<p>Specifies the clock source to determine the timing on serial interfaces.</p> <p>If you use an externally timed clocking mode—dce or loop—long cables might introduce a phase shift of DTE-transmitted clock and data. At high speeds, this phase shift might cause errors.</p> <p>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, do either of the following:</p> <ul style="list-style-type: none"> ■ In the J-Web configuration editor, set the Transmit clock value to invert on the Interfaces > interface-name > Serial options page. ■ In the CLI configuration editor, include the transmit-clock invert statement at the <code>[edit interfaces se-pim/0/port serial-options]</code> hierarchy level. 	<p>From the list, select one of the following timing sources:</p> <ul style="list-style-type: none"> ■ dce—Uses a transmit clock generated by the data circuit-terminating equipment (DCE) for the Services Router's DTE. ■ internal—Uses the Services Router's internal clock. ■ loop—Uses the DCE's or DTE's receive clock (the default). <p>For X.21 serial interfaces, you must use the loop clocking mode.</p> <p>When the Services Router is functioning as DTE, you must use the dce clocking mode for all interfaces except X.21 serial interfaces.</p> <p>When the Services Router is functioning as DCE, we recommend using the internal clocking mode for all interfaces.</p>

Table 34: Serial Quick Configuration Summary *(continued)*

Field	Function	Your Action
Clock Rate NOTE: RS-232 serial interfaces cannot function error-free with a clock rate greater than 200 KHz.	Specifies the line speed in kilohertz or megahertz for serial interfaces that use the DTE clocking mode.	From the list, select one of the following clock rates: <ul style="list-style-type: none"> ■ 1.2 KHz ■ 2.4 KHz ■ 9.6 KHz ■ 19.2 KHz ■ 38.4 KHz ■ 56.0 KHz ■ 64.0 KHz ■ 72.0 KHz ■ 125.0 KHz ■ 148.0 KHz ■ 250.0 KHz ■ 500.0 KHz ■ 800.0 KHz ■ 1.0 MHz ■ 1.3 MHz ■ 2.0 MHz ■ 4.0 MHz ■ 8.0 MHz

Configuring Network Interfaces with a Configuration Editor

To enable the interfaces installed on your Services Router to work properly, you must configure their properties. You can perform basic interface configuration using the J-Web Quick Configuration pages, as described in “Configuring DS1, DS3, Ethernet, and Serial Interfaces with Quick Configuration” on page 104. You can perform the same configuration tasks using the J-Web or CLI configuration editor. In addition, you can configure a wider variety of options that are encountered less frequently.

You can perform the following tasks to configure interfaces:

- Adding a Network Interface with a Configuration Editor on page 131
- Configuring Static ARP Entries on Ethernet Interfaces on page 133
- Deleting a Network Interface with a Configuration Editor on page 134

For information about using the J-Web and CLI configuration editors, see “Using Services Router Configuration Tools” on page 3.

Adding a Network Interface with a Configuration Editor

After you install a PIM, connect the interface cables to the ports, and power on the router, you must complete initial configuration of each network interface, as described in the following procedure:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 35 on page 132.
3. When you are finished configuring the router, commit the configuration.
4. To verify that the network interface is configured correctly, see “Verifying Interface Configuration” on page 135.

Table 35: Adding an Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Configure or Edit. 	From the [edit] hierarchy level, enter <code>edit interfaces interface-name</code>
Create the new interface.	<ol style="list-style-type: none"> 1. Next to Interface, click Add new entry. 2. Enter the name of the new interface in the Interface name box. Make sure the name conforms to the interface naming rules. For more information, see “Network Interface Naming” on page 47. 3. Click OK. 	For information about interface names, see “Network Interface Naming” on page 47.
Create the basic configuration for the new interface.	<ol style="list-style-type: none"> 1. Under Interface Name in the table, click the name of the new interface. 2. Enter values in the other fields on this page if warranted. All these entries are optional, but you need to set values for Clocking and Encapsulation in particular if the default values are not suitable. 	Enter values for physical interface properties as needed. Examples include changes to the default values for physical encapsulation or MTU. For example: <code>set interface-name encapsulation ppp</code>
Add values for interface-specific options. Most interface types have optional parameters that are specific to the interface type.	<ol style="list-style-type: none"> 1. Under Nested configuration, click Configure for the appropriate interface type. 2. In the interface-specific page that appears, enter the values you need to supply or change the default values. 3. When you are finished, click OK to confirm your changes or Cancel to cancel them and return to the previous page. 	<ol style="list-style-type: none"> 1. From the [edit interfaces <i>interface-name</i>] hierarchy level, type <code>edit interface-options</code> 2. Enter the statement for each interface-specific property for which you need to change the default value.

Table 35: Adding an Interface (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
Add logical interfaces.	<ol style="list-style-type: none"> 1. In the main Interface page for this interface, next to Unit, click Add new entry. 2. On the Unit page for logical interfaces that appears, type a number from 0 through 16384 in the Interface unit number box. 3. Enter values in other fields as required for your network. 4. To configure protocol family values if needed, under Family, click Configure next to the appropriate protocol. 5. To access additional subordinate hierarchies under Nested configuration, click Configure next to any parameter you want to configure. 6. When you are finished, click OK. 	<ol style="list-style-type: none"> 1. From the [edit interfaces <i>interface-name</i>] hierarchy level, type <code>set unitlogical-unit-number</code> Replace <i>logical-unit-number</i> with a value from 0 through 16384. 2. Enter additional values for properties you need to configure on the logical interface, such as logical encapsulation or protocol family.

Configuring Static ARP Entries on Ethernet Interfaces

By default, the Services Router responds to an Address Resolution Protocol (ARP) request only if the destination address of the ARP request is on the local network of the incoming interface. For Fast Ethernet or Gigabit Ethernet interfaces, you can configure static ARP entries that associate the IP addresses of nodes on the same Ethernet subnet with their media access control (MAC) addresses. These static ARP entries enable the Services Router to respond to ARP requests even if the destination address of the ARP request is not local to the incoming Ethernet interface.

In this example, you configure a static ARP entry on Gigabit Ethernet interface **ge-0/0/3** of the Services Router consisting of the IP address and corresponding MAC address of a node on the same Ethernet subnet. The **ge-0/0/3** interface has the IP address **10.1.1.1/24**. The node has the IP address **10.1.1.3** and the MAC address **00:ff:85:7f:78:03**. If the node on your network is another router running the JUNOS Internet software, you can enter the **show interfaces *interface-name*** command to learn the IP and MAC (hardware) address of the node.

For more information about configuring static ARP entries, see the *JUNOS Network Interfaces Configuration Guide*.

To configure a static ARP entry on the **ge-0/0/3** interface:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 36 on page 134.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying Interface Configuration” on page 135.

Table 36: Configuring Static ARP Entries

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	From the [edit] hierarchy level, enter edit interfaces ge-0/0/3
Select the Gigabit Ethernet interface ge-0/0/3 .	In the Interface name column, click ge-0/0/3 .	
Configure a static ARP entry on logical unit 0 with the source address 10.1.1.1/24 on the ge-0/0/3 interface.	<ol style="list-style-type: none"> 1. Under Unit, next to 0, click Edit. 2. Under Family, next to Inet, click Edit. 	<ol style="list-style-type: none"> 1. Enter edit unit 0
Set the IP address of the subnet node to 10.1.1.3 and the corresponding MAC address to 00:ff:85:7f:78:03.	<ol style="list-style-type: none"> 3. Under Address, next to 10.1.1.1/24, click Edit. 4. Next to Arp, click Add new entry. 	<ol style="list-style-type: none"> 2. Enter edit family inet address 10.1.1.1/24
To have the router reply to ARP requests from the node, use the publish option.	<ol style="list-style-type: none"> 5. In the Address box, type the IP address of the node—10.1.1.3. 6. Select the Publish check box. 7. From the Mac address type list, select Mac. 8. In the Mac box, type the MAC address 00:ff:85:7f:78:03 of node. 9. Click OK until you return to the Interfaces page. 	<ol style="list-style-type: none"> 3. Enter set arp 10.1.1.3 mac 00:ff:85:7f:78:03 publish

Deleting a Network Interface with a Configuration Editor

To delete an interface on a Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 37 on page 135.



NOTE: Performing this action removes the interface from the software configuration and disables it. Network interfaces remain physically present, and their identifiers continue to appear on the J-Web Monitor and Quick Configuration pages.

Table 37: Deleting an Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	<p>From the [edit] hierarchy level, enter</p> <p>edit interfaces</p>
Select the interface you want to delete.	<p>In the Interface table, under Interface name, select the name of the interface you want to delete.</p> <p>For information about interface names, see “Network Interface Naming” on page 47.</p>	<p>Enter</p> <p>delete interface-name</p>
Execute the selection.	<ol style="list-style-type: none"> 1. Click Discard. 2. In the page that appears, select the appropriate option button. <p>If you have not made any previous changes, the only selection available is Delete Configuration Below This Point.</p>	<p>Commit the configuration change:</p> <p>commit</p>

Verifying Interface Configuration

To verify an interface configuration, perform these tasks:

- Verifying the Link State of All Interfaces on page 136
- Verifying Interface Properties on page 137

Verifying the Link State of All Interfaces

Purpose By using the ping tool on each peer address in the network, verify that all interfaces on the Services Router are operational.

Action For each interface on the Services Router:

1. In the J-Web interface, select **Diagnose > Ping Host**.
2. In the Remote Host box, type the address of the interface for which you want to verify the link state.
3. Click **Start**. Output appears on a separate page.

```
PING 10.10.10.10 : 56 data bytes
64 bytes from 10.10.10.10: icmp_seq=0 ttl=255 time=0.382 ms
64 bytes from 10.10.10.10: icmp_seq=1 ttl=255 time=0.266 ms
```

What It Means If the interface is operational, it generates an ICMP response. If this response is received, the round-trip time in milliseconds is listed in the **time** field. For more information about the output, see the *J-series Services Router Administration Guide*.

See Also For more information about using the J-Web interface to ping a host, see the *J-series Services Router Administration Guide*.
For information about the **ping** command, see the *J-series Services Router Administration Guide* or the *JUNOS System Basics and Services Command Reference*

Verifying Interface Properties

Purpose Verify that the interface properties are correct.

Action From the CLI, enter the `show interfaces detail` command.

```
user@host> show interfaces detail
Physical interface: ge-1/0/0, Enabled, Physical link is Up
Interface index: 134, SNMP ifIndex: 27, Generation: 17
Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
Source filtering: Disabled, Flow control: Enabled
Device flags   : Present Running
Interface flags: SNMP-Traps 16384
Link flags     : None
CoS queues     : 4 supported
Hold-times     : Up 0 ms, Down 0 ms
Current address: 00:90:69:87:44:9d, Hardware address: 00:90:69:87:44:9d
Last flapped   : 2004-08-25 15:42:30 PDT (4w5d 22:49 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes :                0                0 bps
Output bytes :                0                0 bps
Input packets:                0                0 pps
Output packets:              0                0 pps
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 best-effort                0                0                0
  1 expedited-fo              0                0                0
  2 assured-forw              0                0                0
  3 network-cont              0                0                0

Active alarms : None
Active defects : None
```

What It Means The output shows a summary of interface information. Verify the following information:

- The physical interface is **Enabled**. If the interface is shown as **Disabled**, do either of the following:
 - In the CLI configuration editor, delete the **disable** statement at the `[edit interfaces interface-name]` level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the **Interfaces > *interface-name*** page.
- The physical link is **Up**. A link state of **Down** indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The **Last Flapped** time is an expected value. The **Last Flapped** time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected

throughput for the physical interface. To clear the statistics and see only new changes, use the `clear interfaces statistics interface-name` command.

See Also For a complete description of `show interfaces detail` output, see the *JUNOS Interfaces Command Reference*.

Chapter 4

Configuring Channelized T1/E1/ISDN PRI Interfaces

The Services Router supports the software-configurable interfaces on the Dual-Port Channelized T1/E1/ISDN PRI PIM. Each interface can be partitioned into T1 or E1 DS0 channels, or into a combination of T1 or E1 and ISDN Primary Rate Interface (PRI) B-channels and a D-channel.



NOTE: You cannot configure channelized T1/E1/ISDN/PRI interfaces through a J-Web Quick Configuration page. You must use the J-Web or CLI configuration editor. Even after configuration, channelized interfaces do not appear on the Quick Configuration Interfaces page.

This chapter includes the following topics. For more information about interfaces, see “Interfaces Overview” on page 41 and the *JUNOS Network Interfaces Configuration Guide*. For ISDN information, see “Configuring ISDN” on page 211.

- Channelized T1/E1/ISDN PRI Terms on page 139
- Channelized T1/E1/ISDN PRI Overview on page 140
- Before You Begin on page 142
- Configuring Channelized T1/E1/ISDN PRI interfaces with a Configuration Editor on page 142
- Verifying Channelized T1/E1/ISDN PRI Interfaces on page 150
- Frequently Asked Questions About Channelized T1/E1/ISDN PRI Interfaces on page 153

Channelized T1/E1/ISDN PRI Terms

Before configuring channelized T1/E1/ISDN PRI interfaces on a Services Router, become familiar with the terms defined in Table 38 on page 140.

Table 38: Channelized T1/E1/ISDN PRI Terms

Term	Definition
channel group	Combination of DS0 or ISDN PRI B-channels interfaces partitioned from a channelized interface into a single logical bundle.
channelized E1	2.048-Mbps interface that can be configured as a single clear-channel E1 interface or channelized into as many as 31 discrete DS0 interfaces, or up to 30 ISDN PRI B-channels and 1 D-channel. On J-series channelized T1/E1/ISDN PRI interfaces, time slots are numbered from 1 through 31, and time slot 1 is reserved for framing. When the interface is configured for ISDN PRI service, time slot 16 is reserved for the D-channel.
channelized interface	Interface that is a subdivision of a larger interface, minimizing the number of Physical Interface Modules (PIMs) that an installation requires. On a channelized PIM, each port can be configured as a single T1 or E1 clear channel or partitioned into multiple discrete DS0 interfaces or ISDN PRI channels.
channelized T1	1.544-Mbps interface that can be configured as a single clear-channel T1 interface or channelized into as many as 24 discrete DS0 interfaces, or up to 23 ISDN PRI B-channels and 1 D-channel. When the interface is configured for ISDN PRI service, time slot 24 is reserved for the D-channel.
E1 interface	Physical WAN interface for transmitting signals in European digital transmission (E1) format. The E1 signal format transmits information at a rate of 2.048 Mbps and can carry 32 channels of 64 Kbps each.
Primary Rate Interface (PRI)	ISDN service intended for higher-bandwidth applications than ISDN BRI. ISDN PRI consists of a single D-channel for control and signaling, plus a number of 64-Kbps B-channels—either 23 B-channels on a T1 line or 30 B-channels on an E1 line—to carry network traffic.
T1 interface	Physical WAN interface for transmitting digital signals in the T-carrier system used in the United States, Japan, and Canada. The T1 signal format carries 24 pulse code modulation (PCM) signals using time-division multiplexing (TDM) at an overall rate of 1.544 Mbps.

Channelized T1/E1/ISDN PRI Overview

You can configure a channelized T1/E1/ISDN PRI interface for T1 or E1 or ISDN PRI service.

On a channelized T1/E1/ISDN PRI PIM configured for channelized operation, you can use the "drop-and-insert" feature to integrate voice and data on a single T1 or E1 link, and save the cost of two lines.

This overview contains the following topics:

- Channelized T1/E1/ISDN PRI Interfaces on page 140
- Drop and Insert on page 141
- ISDN PRI Transmission on Channelized Interfaces on page 141

Channelized T1/E1/ISDN PRI Interfaces

Each port on a channelized T1/E1/ISDN PRI PIM is software configurable for T1, E1, or ISDN PRI service. Each channelized T1 or E1 interface can be configured as a

single clear channel, or for fractional (NxDs0) or channelized operation, where *N* is channels 1 to 31 for an E1 interface and channels 1 to 24 for a T1 interface.

Each channelized interface can be configured as ISDN PRI B-channels and one D-channel or as a combination of T1 or E1 DS0 channels and ISDN PRI channels.

J-series ISDN PRI interfaces support the following switch types:

- ATT5E—AT&T 5ESS
- ETSI—NET3 for the United Kingdom and Europe
- NI2—National ISDN-2
- NTDMS100—Northern Telecom DMS-100
- NTT—NTT Group switch for Japan

For more information, see “ISDN PRI Transmission on Channelized Interfaces” on page 141.

Channelized T1/E1/ISDN PRI interfaces are configured through a configuration editor only.

A channelized T1/E1/ISDN PRI interface supports CoS configuration. For information about CoS features, see the *J-series Services Router Advanced WAN Access Configuration Guide*.

Drop and Insert

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. You need not use the same time slots on both interfaces, but the time slots count must be the same.

The channels that are not configured for the drop-and-insert feature are used for normal traffic.

ISDN PRI Transmission on Channelized Interfaces

The Dual-Port Channelized T1/E1/ISDN PRI PIM provides support for ISDN PRI services such as dial-in at the central office, callback from the central office, and primary or backup network connections from branch offices. For more information about the services, see “Configuring ISDN” on page 211.

You can configure up to 23 time slots in a channelized T1 PRI interface and up to 30 time slots in a channelized E1 PRI interface as B-channels. The 24th time slot in a T1 interface and the 16th time slot in an E1 interface are configured as the D-channel interface for signaling purposes. Each B-channel supports 64 Kbps of traffic. The unconfigured time slots can be used as regular DS0 interfaces on top of the T1 or E1 physical layer.

You can install channelized T1/E1/ISDN PRI PIMs and ISDN BRI PIMs and configure both ISDN PRI and ISDN BRI service on the same Services Router.

Before You Begin

Before you configure network interfaces, you need to perform the following tasks:

- Install Services Router hardware. For more information, see the Getting Started Guide for your router.
- Establish basic connectivity. For more information, see the Getting Started Guide for your router.
- If you do not already have a basic understanding of physical and logical interfaces and Juniper Networks interface conventions, read “Interfaces Overview” on page 41.

Although it is not a requirement, you might also want to plan how you are going to use the various network interfaces before you start configuring them.

Configuring Channelized T1/E1/ISDN PRI interfaces with a Configuration Editor

Each port on a channelized T1/E1/ISDN PRI PIM is software configurable as a T1 or E1 clear channel. You can partition each port into up to 24 DS0 channels on a T1 interface or up to 31 DS0 channels on an E1 interface, and can insert channels from one port into another with the drop-and-insert feature.

Channelized T1/E1/ISDN PRI ports can also be partitioned into channels for ISDN PRI service.

Channelized T1/E1/ISDN PRI interfaces are configured through a configuration editor only.

This section includes the following topics:

- Configuring Channelized T1/E1/ISDN PRI Interface as a Clear Channel on page 142
- Configuring Channelized T1/E1/ISDN PRI Interface to Drop and Insert Time Slots on page 145
- Configuring Channelized T1/E1/ISDN PRI Interfaces for ISDN PRI Operation on page 147

Configuring Channelized T1/E1/ISDN PRI Interface as a Clear Channel

To configure or edit a channelized T1/E1/ISDN PRI interface as a clear channel:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 39 on page 143.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying Interface Configuration” on page 135.

Table 39: Configuring a Channelized T1/E1/ISDN PRI interface as a Clear Channel

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	1. In the J-Web interface, select Configuration > Edit Configuration > View and Edit .	From the [edit] hierarchy level, enter one of the following:
For information about interface names, see “Network Interface Naming” on page 47.	2. Next to Interfaces, click Edit .	edit interfaces ct1-3/0/0 edit interfaces ce1-3/0/0
Create the new interface—for example, ct1-3/0/0 or ce1-3/0/0.	1. Next to Interfaces, click Add new entry . 2. In the Interface name box, type one of the following interface names: ■ ct1-3/0/0 ■ ce1-3/0/0 3. Click OK .	
Configure interface options:	1. In the Interface table, under Interface name, click the interface you are configuring: ■ ct1-3/0/0 ■ ce1-3/0/0 2. From the Clocking list, select internal . 3. In the Description box, type one of the following descriptions: ■ clear t1 interface ■ clear e1 interface 4. Under Hold time: Next to Down, type 500. Next to Up, type 500. 5. Under No partition, from the Interface type list, select the type of interface: ■ t1 ■ e1 6. From the Scheduler type list, select Per unit scheduler .	1. Enter set clocking internal 2. Add a description: ■ For T1 interfaces, enter set description clear t1 interface. ■ For E1 interfaces, enter set description clear e1 interface. 3. Enter set hold-time down 500 up 500 4. Specify a clear channel: ■ For T1 interfaces, enter set no-partition interface-type t1. ■ For E1 interfaces, enter set no-partition interface-type e1. 5. Enter set per-unit-scheduler

Table 39: Configuring a Channelized T1/E1/ISDN PRI interface as a Clear Channel *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure T1 or E1 options: <ul style="list-style-type: none"> ■ Bit error rate test (BERT) algorithm—for example, all ones repeating. ■ BERT error rate, a value from 0 through 7—for example, 5. ■ BERT period, in seconds, a value from 1 through 240—for example, 5. ■ (T1 interfaces only) Line buildout, in feet for cables 655 ft (200 m) or shorter—for example, 0-132—or in decibels for longer cables. ■ Framing mode: <ul style="list-style-type: none"> ■ For T1 interfaces, either superframe or extended superframe (ESF)—for example, ESF ■ For E1 interfaces, G704, G704 without cyclic redundancy check 4 (CRC4), or G703 unframed—for example, G704. ■ (T1 interfaces only) Line encoding method—for example, alternate mark inversion (AMI). ■ Loopback mode—for example, local. 	<ol style="list-style-type: none"> Next to T1 options or E1 options, click Configure. From the Bert algorithm list, select all-ones-repeating. In the Bert error rate box, type 5. In the Bert period box, type 5. For T1 interfaces only, from the Buildout list, select 0-132. From the Framing list: <ul style="list-style-type: none"> ■ For T1 interfaces, select esf. ■ For E1 interfaces, select g704. For T1 interfaces only, from the Line encoding list, select ami From the Loopback list, select local. Click OK 	<ol style="list-style-type: none"> Enter set bert-algorithm all-ones-repeating Enter set bert-error-rate 5 Enter set bert-period 5 For T1 interfaces only, enter set buildout 0-132 Set the framing mode: <ul style="list-style-type: none"> ■ For T1 interfaces, enter set framing esf. ■ For E1 interfaces, enter set framing g704. For T1 interfaces only, enter set line encoding ami Enter set loopback local
Configure trace options.	<ol style="list-style-type: none"> Next to Traceoptions, select the check box and click Configure. Next to Flag, click Add new entry. From the Flag name list, select all. Click OK until you return to the Interface page. 	Enter set traceoptions flag all
Configure advanced options. For example, apply configuration settings from one or more groups except the test group.	<ol style="list-style-type: none"> Next to Advanced, click the expand (+) icon. Next to Apply groups except, click Add new entry. In the Value box, type test. Click OK. 	Enter set interfaces apply-groups test

Configuring Channelized T1/E1/ISDN PRI Interface to Drop and Insert Time Slots

On channelized T1/E1/ISDN PRI 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. Although you need not use the same time slots on both interfaces, the time slots count must be the same. The channels that are not configured for the drop-and-insert feature are used for normal traffic.

You must ensure that the signaling channels (port 16 for an E1 interface and port 24 for a T1 interface) are also 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 “Frequently Asked Questions About Channelized T1/E1/ISDN PRI Interfaces” on page 153.

To configure or edit the drop-and-insert feature on a channelized T1/E1/ISDN PRI interface:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 40 on page 146.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying Interface Configuration” on page 135.

Table 40: Configuring a Channelized T1/E1/ISDN PRI Interface to Drop and Insert Time Slots

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Navigate to the Interfaces level in the configuration hierarchy.</p> <p>For information about interface names, see “Network Interface Naming” on page 47.</p>	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > Edit Configuration > View and Edit. 2. Next to Interfaces, click Edit. 	<p>From the [edit] hierarchy level, enter</p> <p>edit interfaces ct1-3/0/0</p>
<p>Create a new interface—for example, ct1-3/0/0.</p>	<ol style="list-style-type: none"> 1. Next to Interfaces, click Add new entry. 2. In the Interface name box, type ct1-3/0/0. 3. Click OK. 	
<p>Configure the clock source and partition on ct1-3/0/0.</p> <p>NOTE: While configuring the drop-and-insert feature, you must ensure that both ports on the channelized T1/E1 PIM run on the same clock.</p> <p>For more details about valid clock combinations, see “Frequently Asked Questions About Channelized T1/E1/ISDN PRI Interfaces” on page 153.</p>	<ol style="list-style-type: none"> 1. In the Interface name column, click ct1-3/0/0. 2. On the Interfaces page, next to Clocking, select the check box and click Configure. 3. From the Clocking choices list, select external. 4. Click OK. 5. On the Interfaces page, next to Partition, click Add new entry. 6. On the Interface Partition page, type 1 in the Partition number box. 7. From the Interface type list, Select ds. 8. In the Timeslots box, type 1-10. 9. Click OK twice. 	<p>From the [edit] hierarchy level, enter</p> <p>set interfaces ct1-3/0/0 clocking external</p> <p>set interfaces ct1-3/0/0 partition 1 timeslots 1-10</p> <p>set interfaces ct1-3/0/0 partition 1 interface-type ds</p>
<p>Create a new interface—for example, ct1-3/0/1.</p>	<ol style="list-style-type: none"> 1. On the Interfaces Configuration page, next to Interface, click Add new entry. 2. In the Interface name box, type ct1-3/0/1. 3. Click OK. 	<p>From the [edit] hierarchy level, enter</p> <p>edit interfaces ct1-3/0/1</p>

Table 40: Configuring a Channelized T1/E1/ISDN PRI Interface to Drop and Insert Time Slots *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Configure the clock source and partition on ct1-3/0/1.</p> <p>NOTE: While configuring the drop-and-insert feature, you must ensure that both ports on the channelized T1/E1 PIM run on the same clock.</p> <p>For more details about valid clock combinations, see “Frequently Asked Questions About Channelized T1/E1/ISDN PRI Interfaces” on page 153.</p>	<ol style="list-style-type: none"> On the Interfaces Configuration page, click ct1-3/0/1 in the Interface name column. Next to Clocking, select the Yes check box, and click Configure. From the Clocking choices list, select external. Next to External, click Configure. In the Interface box, type ct1-3/0/0. Click OK twice. On the Interfaces page, next to Partition, click Add new entry. On the Interface Partition page, type 1 in the Partition number box. From the Interface type list, Select ds. In the Timeslots box, type 1-10. Click OK twice. 	<p>From the [edit] hierarchy level, enter</p> <pre>set interfaces ct1-3/0/1 clocking external interface ct1-3/0/0</pre> <p>set interfaces ct1-3/0/1 partition 1 timeslots 1-10</p> <p>set interfaces ct1-3/0/1 partition 1 interface-type ds</p>
<p>Create new interfaces—for example, ds-3/0/0:1, ds-3/0/1:1 and configure drop-and-insert feature.</p> <p>NOTE: Both interfaces configured for the drop-and-insert feature must exist on the same PIM. For example, you can configure ds-3/0/0:1 as the data input interface for ds-3/0/1:1, but not for ds-4/0/0:1.</p>	<ol style="list-style-type: none"> On the Interfaces Configuration page, next to Interface, click Add new entry. In the Interface name box, type ds-3/0/0:1. Click OK. On the Interfaces Configuration page, click ds-3/0/0:1 in the Interface name column. Next to Data input, click Configure. From the Input choice list, select interface. In the Interface box, type ds-3/0/1:1. Click OK. 	<p>From the [edit] hierarchy level, enter</p> <pre>edit interfaces ds-3/0/0:1</pre> <p>Enter</p> <pre>set interfaces ds-3/0/0:1 data-input interface ds-3/0/1:1</pre>

Configuring Channelized T1/E1/ISDN PRI Interfaces for ISDN PRI Operation

On a J-series Services Router with Dual-Port Channelized T1/E1/ISDN PRI PIMs, you can configure each port for either T1, E1, or ISDN PRI service, or for a combination of ISDN PRI and either channelized T1 or E1 service. For a channelized T1 interface with ISDN PRI service, you can configure 23 B-channels and for a channelized E1 interface with ISDN PRI service, you can configure 30 B-channels.

You must also explicitly configure a D-channel: time slot 24 on a channelized T1 interface and time slot 16 on a channelized E1 interface. In addition, you select a switch type and trace options.

Setting up the router for ISDN PRI operation is a multipart process. First, you add ISDN PRI service on a channelized interface as shown here. Second, you follow the instructions in “Configuring Dialer Interfaces (Required)” on page 226 to configure a dialer interface. You can then configure ISDN services such as dial-in, callback, and backup. For details, see “Configuring ISDN” on page 211.

To configure an ISDN PRI network service on a channelized T1 or E1 interface for the Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 41 on page 148.
3. Go on to “Configuring Dialer Interfaces (Required)” on page 226.

Table 41: Adding an ISDN PRI Service to a Channelized T1/E1/ISDN PRI Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	From the [edit] hierarchy level, enter edit interfaces ct1-2/0/0
Create a new interface—for example, ct1-2/0/0. For information about interface names, see “Network Interface Naming” on page 47.	<ol style="list-style-type: none"> 1. Next to Interfaces, click Add new entry. 2. In the Interface name box, type ct1-2/0/0. 3. Click OK. 	

Table 41: Adding an ISDN PRI Service to a Channelized T1/E1/ISDN PRI Interface (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Configure the partition and interface type. For example, partition the interface into time slots 1 through 23 for B-channels and time slot 24 for the D-channel.</p> <p>For a channelized T1 interface, you can configure 1 through 23 as B-channels and the 24th channel as the signaling channel (D-channel).</p> <p>For a channelized E1 interface, you can configure 1 through 15 and 17 through 31 as B-channels and the 16th channel as the signaling channel (D-channel).</p>	<ol style="list-style-type: none"> 1. In the Interface name column, click ct1-2/0/0. 2. On the Interfaces page, next to Partition, click Add new entry. 3. In the Partition number box, type 1-23. 4. In the Timeslots box, type 1-23. 5. From the Interface type list, Select bc. 6. Click OK. 7. On the Interfaces page, next to Partition, click Add new entry. 8. In the Partition number box, type 24. 9. In the Timeslots box, type 24. 10. From the Interface type list, Select dc. 11. Click OK. 	<p>From the [edit] hierarchy level, enter</p> <p>set interfaces ct1-2/0/0 partition 1-23 timeslots 1-23</p> <p>set interfaces ct1-2/0/0 partition 1-23 interface-type bc</p> <p>set interfaces ct1-2/0/0 partition 24 timeslots 24</p> <p>set interfaces ct1-2/0/0 partition 24 interface-type dc</p>
Configure a trace options flag.	<ol style="list-style-type: none"> 1. Next to Traceoptions, select the check box and click Configure. 2. Next to Flag, click Add new entry. 3. From the Flag name list, select q921. 4. Click OK until you return to the Interface page. 	<p>From the [edit] hierarchy level, enter</p> <p>set interfaces ct1-2/0/0 traceoptions flag q921</p>
Configure B-channel allocation order for allocating a free B-channel for dial-out calls. You can allocate from the lowest-numbered or highest-numbered time slot. The default value is descending .	<ol style="list-style-type: none"> 1. On the Interfaces page, next to Isdn options, click Configure. 2. From the Bchannel allocation list, select ascending. 3. Click OK. 	<p>To set the ISDN options, from the [edit] hierarchy level, enter</p> <p>set interfaces ct1-2/0/0 isdn-options bchannel-allocation ascending</p>
<p>Select the type of ISDN switch—for example, NI2. The following switches are compatible with Services Routers:</p> <ul style="list-style-type: none"> ■ ATT5E—AT&T 5ESS ■ ETSI—NET3 for the UK and Europe ■ NI2—National ISDN-2 ■ NTDMS-100—Northern Telecom DMS-100 ■ NTT—NTT Group switch for Japan 	From the Switch type list, select ni2 .	<p>From the [edit] hierarchy level, enter</p> <p>set interfaces ct1-2/0/0 isdn-options switch-type ni2</p>

Table 41: Adding an ISDN PRI Service to a Channelized T1/E1/ISDN PRI Interface *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure Q.931 timers. Q.931 is a Layer 3 protocol for the setup and termination of connections. The default value for each timer is 10 seconds, but can be configured between 1 and 65536 seconds—for example, 15.	<ol style="list-style-type: none"> 1. In the T310 box, type 15. 2. Click OK. 	<p>From the [edit] hierarchy level, enter</p> <pre>set isdn-options t310 15</pre>
<p>Configure dialer options.</p> <ul style="list-style-type: none"> ■ Name the dialer pool—for example, ISDN-dialer-group. ■ Set the dialer pool priority—for example, 1. <p>Dialer pool priority has a range from 1 to 255, with 1 designating lowest-priority interfaces and 255 designating the highest-priority interfaces.</p>	<ol style="list-style-type: none"> 1. On the Interfaces page, next to Dialer options, select Yes and then click configure. 2. Next to Pool, click Add new entry. 3. In the Pool identifier box, type <code>isdn-dialer-group</code>. 4. In the Priority box, type 1. 5. Click OK. 	<p>From the [edit interfaces ct1-2/0/0] hierarchy level, enter</p> <pre>set dialer-options pool isdn-dialer-group priority 1</pre>

To configure a dialer interface, see “Configuring Dialer Interfaces (Required)” on page 226.

Verifying Channelized T1/E1/ISDN PRI Interfaces

To verify an interface configuration, perform these tasks:

- Verifying Channelized Interfaces on page 151
- Verifying Clear-Channel Interfaces on page 152
- Verifying ISDN PRI Configuration on Channelized T1/E1/ISDN PRI Interfaces on page 153

Verifying Channelized Interfaces

Purpose Verify that your configurations for the channelized interfaces are correct.

Action From the CLI, enter the `show interfaces ct1-3/0/1` command.

```
user@host> show interfaces ct1-3/0/1
```

```
Physical interface: ct1-3/0/1, Enabled, Physical link is Up
  Interface index: 151, SNMP ifIndex: 28
  Link-level type: Controller, Clocking: Internal, Speed: E1, Loopback: None,
  Framing: G704, Parent: None
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : None
  CoS queues     : 8 supported, 8 maximum usable queues
  Last flapped   : 2006-10-05 21:11:48 PDT (06:45:04 ago)
  DS1 alarms     : None
  DS1 defects    : None
  Line encoding  : HDB3
```

What It Means The output shows a summary of information about the physical parent interface—a channelized T1 interface in this example.

See Also For a complete description of `show interfaces` output, see the *JUNOS Interfaces Command Reference*.

Verifying Clear-Channel Interfaces

Purpose Verify that your configurations for the clear-channel interfaces are correct.

Action From the CLI, enter the `show interfaces e1-3/0/1` command.

```
user@host> show interfaces e1-3/0/1
```

```
Physical interface: e1-3/0/1, Enabled, Physical link is Up
  Interface index: 212, SNMP ifIndex: 237
  Link-level type: PPP, MTU: 1504, Speed: E1, Loopback: None, FCS: 16,
  Parent: ce1-3/0/1 Interface index 151
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link flags     : Keepalives
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 1066 (00:00:02 ago), Output: 1066 (00:00:02 ago)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Opened, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Closed
  CoS queues      : 8 supported, 8 maximum usable queues
  Last flapped    : 2006-10-06 01:01:36 PDT (02:57:27 ago)
  Input rate      : 88 bps (0 pps)
  Output rate     : 58144 bps (157 pps)
  DS1 alarms     : None
  DS1 defects    : None

Logical interface e1-3/0/1.0 (Index 66) (SNMP ifIndex 238)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: PPP
  Bandwidth: 1984kbps
  Protocol inet, MTU: 1500
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 47.47.47.0/30, Local: 47.47.47.2, Broadcast: 47.47.47.3
  Protocol inet6, MTU: 1500
    Flags: None
    Addresses, Flags: Is-Preferred Is-Primary
      Destination: 8b8b:8b01::/64, Local: 8b8b:8b01::2
    Addresses, Flags: Is-Preferred
      Destination: fe80::/64, Local: fe80::205:85ff:fec5:d3d0
```

What It Means The output shows a summary of interface information. Although the parent interface is `ce1-3/0/1`, the physical and logical clear-channel interfaces are named `e1-3/0/1` and `e1-3/0/1.0`.

See Also For a complete description of `show interfaces` output, see the *JUNOS Interfaces Command Reference*.

Verifying ISDN PRI Configuration on Channelized T1/E1/ISDN PRI Interfaces

Purpose Verify that your configuration of ISDN PRI service on a channelized interface is correct.

Action From the J-Web interface, select **Configuration > View and Edit > View Configuration Text**. Alternatively, from configuration mode in the CLI, enter the `show interfaces ct1-2/0/0` command.

```
user@host# show interfaces ct1-2/0/0

traceoptions {
    flag q921;
    file {
        isdnback;
    }
}
clocking external;
isdn-options {
    switch-type ni2;
}
dialer-options {
    isdn-dialer-group priority 1;
}
partition 24 timeslots 24 interface-type dc;
partition 1-23 timeslots 1-23 interface-type bc;

[edit]
```

What It Means Verify that the output shows your intended ISDN PRI interface configuration.

See Also For more information about the format of a configuration file, see “Viewing the Configuration Text” on page 9.
To additionally verify ISDN PRI configuration, see “Verifying the ISDN Configuration” on page 245.

Frequently Asked Questions About Channelized T1/E1/ISDN PRI Interfaces

Use answers to the following question to solve configuration problems on a channelized T1/E1/ISDN PRI interface:

- What Clock Combinations Are Possible for Channelized T1/E1/ISDN PRI Drop and Insert? on page 153

What Clock Combinations Are Possible for Channelized T1/E1/ISDN PRI Drop and Insert?

When you configure the drop-and-insert feature on a channelized T1/E1/ISDN PRI PIM, you must ensure that both ports run on the same clock. The following clock combinations are valid:

- When port 0 is configured to use the *internal* clock, port 1 must also be configured to use the *internal* clock.
- When port 0 is configured to use the *external* clock, port 1 must be configured to run on the same clock, the *external clock for port 0*.

- When port 1 is configured to use the *external* clock, port 0 must be configured to run on the same clock, the *external clock for port 1*.

Services Routers connected to one another must have complementary clock sources configured. Consider a scenario where Router R1 is connected to Routers R2 and R3. Port 0 on the channelized T1/E1/ISDN PRI PIM of R1 is connected to R2, and port 1 is connected to R3. The drop-and-insert feature is configured on R1 to insert input coming from R2 on port 0 into port 1 for transmission to R3.

Routers R1, R2, and R3 can be configured in three ways, according to whether the drop-and-insert clock source on R1 is the external clock for port 0, the external clock for port 1, or the router's internal clock.

To configure the drop-and-insert interfaces on Router R1 to use the external clock for port 0:

1. On Router R2, configure:

```
user@hostR2# set interfaces ct1-6/0/0 partition 1 timeslots 1-10
user@hostR2# set interfaces ct1-6/0/0 partition 1 interface-type ds
user@hostR2# set interfaces ds-6/0/0:1 unit 0 family inet address 10.46.46.1/30
```

2. On Router R3, configure:

```
user@hostR3# set interfaces ct1-3/0/0 clocking external
user@hostR3# set interfaces ct1-3/0/0 partition 1 timeslots 1-10
user@hostR3# set interfaces ct1-3/0/0 partition 1 interface-type ds
user@hostR3# set interfaces ds-3/0/0:1 unit 0 family inet address 10.46.46.2/30
```

3. On Router R1, configure:

```
user@hostR1# set interfaces ct1-3/0/0 clocking external
user@hostR1# set interfaces ct1-3/0/0 partition 1 timeslots 1-10
user@hostR1# set interfaces ct1-3/0/0 partition 1 interface-type ds
user@hostR1# set interfaces ds-3/0/0:1 data-input interface ds-3/0/1:1
user@hostR1# set interfaces ct1-3/0/1 clocking external interface ct1-3/0/0
user@hostR1# set interfaces ct1-3/0/1 partition 1 timeslots 1-10
user@hostR1# set interfaces ct1-3/0/1 partition 1 interface-type ds
```

To configure the drop-and-insert interfaces on Router R1 to use the external clock for port 1:

1. On Router R2, configure:

```
user@hostR2# set interfaces ct1-6/0/0 clocking external
user@hostR2# set interfaces ct1-6/0/0 partition 1 timeslots 1-10
user@hostR2# set interfaces ct1-6/0/0 partition 1 interface-type ds
user@hostR2# set interfaces ds-6/0/0:1 unit 0 family inet address 10.46.46.1/30
```

2. On Router R3, configure:

```
user@hostR3# set interfaces ct1-3/0/0 partition 1 timeslots 1-10
user@hostR3# set interfaces ct1-3/0/0 partition 1 interface-type ds
user@hostR3# set interfaces ds-3/0/0:1 unit 0 family inet address 10.46.46.2/30
```

3. On Router R1, configure:

```
user@hostR1# set interfaces ct1-3/0/0 clocking external interface ct1-3/0/1
user@hostR1# set interfaces ct1-3/0/0 partition 1 timeslots 1-10
```

```

user@hostR1# set interfaces ct1-3/0/0 partition 1 interface-type ds
user@hostR1# set interfaces ds-3/0/0:1 data-input interface ds-3/0/1:1
user@hostR1# set interfaces ct1-3/0/1 clocking external
user@hostR1# set interfaces ct1-3/0/1 partition 1 timeslots 1-10
user@hostR1# set interfaces ct1-3/0/1 partition 1 interface-type ds

```

To configure the drop-and-insert interfaces on Router R1 to use the router's internal clock:

1. On Router R2, configure:

```

user@hostR2# set interfaces ct1-6/0/0 clocking external
user@hostR2# set interfaces ct1-6/0/0 partition 1 timeslots 1-10
user@hostR2# set interfaces ct1-6/0/0 partition 1 interface-type ds
user@hostR2# set interfaces ds-6/0/0:1 unit 0 family inet address 10.46.46.1/30

```

2. On Router R3, configure:

```

user@hostR3# set interfaces ct1-3/0/0 clocking external
user@hostR3# set interfaces ct1-3/0/0 partition 1 timeslots 1-10
user@hostR3# set interfaces ct1-3/0/0 partition 1 interface-type ds
user@hostR3# set interfaces ds-3/0/0:1 unit 0 family inet address 10.46.46.2/30

```

3. On Router R1, configure:

```

user@hostR1# set interfaces ct1-3/0/0 clocking internal
user@hostR1# set interfaces ct1-3/0/0 partition 1 timeslots 1-10
user@hostR1# set interfaces ct1-3/0/0 partition 1 interface-type ds
user@hostR1# set interfaces ds-3/0/0:1 data-input interface ds-3/0/1:1
user@hostR1# set interfaces ct1-3/0/1 clocking internal
user@hostR1# set interfaces ct1-3/0/1 partition 1 timeslots 1-10
user@hostR1# set interfaces ct1-3/0/1 partition 1 interface-type ds

```


Chapter 5

Configuring Digital Subscriber Line Interfaces

The Services Router supports DSL features including ATM-over-ADSL and ATM-over-SHDSL interfaces.

You can use either J-Web Quick Configuration or a configuration editor to configure ATM-over-ADSL or ATM-over-SHDSL interfaces.



NOTE: Payload loopback functionality is not supported on ATM-over-SHDSL interfaces.

This chapter contains the following topics.

- DSL Terms on page 157
- Before You Begin on page 158
- Configuring ATM-over-ADSL Interfaces on page 159
- Configuring ATM-over-SHDSL Interfaces on page 169
- Configuring CHAP on DSL Interfaces (Optional) on page 179
- Verifying DSL Interface Configuration on page 180

DSL Terms

Before configuring DSL on a Services Router, become familiar with the terms defined in Table 42 on page 157.

Table 42: DSL Terms

Term	Definition
asymmetric digital subscriber line (ADSL) interface	Physical WAN interface for connecting a Services Router to a digital subscriber line access multiplexer (DSLAM). An ADSL interface allocates line bandwidth asymmetrically with downstream (provider-to-customer) data rates of up to 8 Mbps for ADSL, 12 Mbps for ADSL2, and 25 Mbps for ADSL2 + , and upstream (customer-to-provider) rates of up to 800 Kbps for ADSL and 1 Mbps for ADSL2 and ADSL2 + , depending on the implementation.

Table 42: DSL Terms *(continued)*

Term	Definition
ADSL2 interface	An ADSL interface that supports ITU-T Standards G.992.3 and G.992.4 and allocates downstream (provider-to-customer) data rates of up to 12 Mbps and upstream (customer-to-provider) rates of up to 1 Mbps.
ADSL2+ interface	An ADSL interface that supports ITU-T Standard G.992.5 and allocates downstream (provider-to-customer) data rates of up to 25 Mbps and upstream (customer-to-provider) rates of up to 1 Mbps.
Annex A	ITU-T Standard G.992.1 that defines how ADSL works over plain old telephone service (POTS) lines.
Annex B	ITU-T Standard G.992.1 that defines how ADSL works over Integrated Services Digital Network (ISDN) lines.
ITU-T G.991.2	International Telecommunication Union standard describing a data transmission method for symmetric high-speed digital subscriber line (SHDSL) as a means for data transport in telecommunications access networks. The standard also describes the functionality required for interoperability of equipment from various manufacturers.
ITU-T G.992.1	International Telecommunication Union standard that requires the downstream (provider-to-customer) data transmission to consist of full-duplex low-speed bearer channels and simplex high-speed bearer channels. In the upstream (customer-to-provider) transmissions, only low-speed bearer channels are provided.
ITU-T G.994.1	International Telecommunication Union standard describing the types of signals, messages, and procedures exchanged between digital subscriber line (DSL) equipment when the operational modes of equipment need to be automatically established and selected.
ITU-T G.997.1	International Telecommunication Union standard describing the physical layer management for asymmetric digital subscriber line (ADSL) transmission systems. The standard specifies the means of communication on a transport transmission channel defined in the physical layer recommendations. In addition, the standard describes the content and syntax of network elements for configuration, fault management, and performance management.
symmetric high-speed digital subscriber line (G.SHDSL)	Physical WAN symmetric DSL interface capable of sending and receiving high-speed symmetrical data streams over a single pair of copper wires at rates between 192 Kbps and 2.31 Mbps. G.SHDSL incorporates features of other DSL technologies such as asymmetric DSL and transports T1, E1, ISDN, Asynchronous Transfer Mode (ATM), and IP signals.
symmetric high-speed digital subscriber line (SHDSL) transceiver unit-remote (STU-R)	Equipment that provides symmetric high-speed digital subscriber line (SHDSL) connections to remote user terminals such as data terminals or telecommunications equipment.

Before You Begin

Before you begin configuring DSL interfaces, complete the following tasks:

- Establish basic connectivity. See the Getting Started Guide for your router.
- If you do not already have a basic understanding of physical and logical interfaces and Juniper Networks interface conventions, read “Interfaces Overview” on page 41.

- Configure network interfaces as necessary. See “Configuring Ethernet, DS1, DS3, and Serial Interfaces” on page 103.

Configuring ATM-over-ADSL Interfaces

Services Routers with ADSL Annex A or Annex B PIMs can use an Asynchronous Transfer Mode (ATM) interface to send network traffic through a point-to-point connection to a DSL access multiplexer (DSLAM). ADSL is currently not supported on the J2300 Services Router.



NOTE: You can configure Services Routers with ADSL PIMs for connections through ADSL only, not for direct ATM connections.

To configure Point-to-Point Protocol (PPP), see the *JUNOS Network Interfaces Configuration Guide*.

You configure the underlying ADSL interface as an ATM interface, with an interface name of **at-pim/0/port**. (For information about interface names, see “Network Interface Naming” on page 47.) Multiple encapsulation types are supported on both the physical and logical ATM-over-ADSL interface.

This section contains the following topics:

- Configuring an ATM-over-ADSL Interface with Quick Configuration on page 159
- Adding an ATM-over-ADSL Network Interface with a Configuration Editor on page 164

Configuring an ATM-over-ADSL Interface with Quick Configuration

The Quick Configuration pages allow you to configure ATM-over-ADSL interfaces on a Services Router.

To configure an ATM-over-ADSL interface with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Quick Configuration > Interfaces**.

A list of the network interfaces present on the Services Router is displayed. (For information about interface names, see “Network Interface Naming” on page 47.)

2. Select the **at-pim/0/port** interface name for the ADSL port you want to configure.

The ATM-over-ADSL Quick Configuration page is displayed, as shown in Figure 30 on page 160.

Figure 30: ATM-over-ADSL Interface Quick Configuration Page

Juniper
NETWORKS

ROUTER - J6300

Monitor Configuration Diagnose Manage Events Alarms Logged in as: regress Help About Logout

Configuration > Quick Configuration > Interfaces

Quick Configuration

Interfaces DSL Physical Interface: 'at-5/0/0'

Logical Interfaces

No logical interfaces configured.

Add...

Physical Interface Description

MTU (bytes) ?

Encapsulation

VPI ?

ADSL Options

Operating Mode ?

OK Cancel Apply

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. Trademark Notice. Privacy. Juniper Your Net.

3. Enter information into the ATM-over-ADSL Quick Configuration pages, as described in Table 43 on page 160.
4. From the ATM-over-ADSL Quick Configuration main page, click one of the following buttons:
 - To apply the configuration and stay on the ATM-over-ADSL Quick Configuration main page, click **Apply**.
 - To apply the configuration and return to the Interfaces Quick Configuration page, click **OK**.
 - To cancel your entries and return to the Interfaces Quick Configuration page, click **Cancel**.
5. To verify that the ATM-over-ADSL interface is configured properly, see “Verifying DSL Interface Configuration” on page 180.

Table 43: ATM-over-ADSL Interface Quick Configuration Pages Summary

Field	Function	Your Action
Configuring Logical Interfaces		

Table 43: ATM-over-ADSL Interface Quick Configuration Pages Summary *(continued)*

Field	Function	Your Action
Logical Interfaces	Lists the logical interfaces for this ATM-over-ADSL physical interface.	<ul style="list-style-type: none"> ■ To add a logical interface, click Add. ■ To edit a logical interface, select the interface from the list. ■ To delete a logical interface, select the check box next to the name and click Delete.
Adding or Editing a Logical Interface		
Add logical interfaces	Defines one or more logical units that you connect to this physical ADSL interface.	Click Add .
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.

Table 43: ATM-over-ADSL Interface Quick Configuration Pages Summary (*continued*)

Field	Function	Your Action
Encapsulation	Specifies the type of encapsulation on the DSL logical interface.	<p>From the list, select one of the following types of encapsulations.</p> <p>For ATM-over-ADSL interfaces that use <code>inet</code> (IPv4) protocols only, select one of the following:</p> <ul style="list-style-type: none"> ■ ATM VC multiplexing—Use ATM virtual circuit multiplex encapsulation. ■ ATM NLPID—Use ATM network layer protocol identifier (NLPID) encapsulation. ■ Cisco-compatible ATM NLPID—Use Cisco NLPID encapsulation. ■ Ethernet over ATM (LLC/SNAP)—For interfaces that carry IPv4 traffic, use Ethernet over logical link control (LLC) encapsulation. You cannot configure multipoint interfaces if you use this type of encapsulation. <p>For ATM-over-ADSL for PPP-over-ATM (PPPoA) interfaces only, select one of the following:</p> <ul style="list-style-type: none"> ■ ATM PPP over AAL5/LLC—Use AAL5 logical link control (LLC) encapsulation. ■ ATM PPP over Raw AAL5—Use AAL5 multiplex encapsulation. <p>For other encapsulation types on the ATM-over-ADSL interfaces, select one of the following:</p> <ul style="list-style-type: none"> ■ PPPoE over ATM (LLC/SNAP)—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. ■ Ethernet over ATM (LLC/SNAP)—Use ATM subnetwork attachment point (SNAP) encapsulation.
VCI	Configures the ATM virtual circuit identifier (VCI) for the interface.	In the VCI box, type the number for the VCI.
Add IPv4 address prefixes and destinations	Specifies one or more IPv4 addresses and destination addresses.	Click Add .
IPv4 Address Prefix	Specifies an IPv4 address for the interface.	Type an IPv4 address and prefix. For example: 10.10.10.10/24

Table 43: ATM-over-ADSL Interface Quick Configuration Pages Summary (*continued*)

Field	Function	Your Action
Destination Address	Specifies the destination address.	<ol style="list-style-type: none"> 1. Type an IPv4 address for the destination. 2. Click OK.
Configuring Physical Interface Properties		
Physical Interface Description	(Optional) Adds supplementary information about the physical ATM-over-ADSL interface.	Type a text description of the physical ATM-over-ADSL interface to more clearly identify it in monitoring displays. Specify that it is an ADSL interface.
MTU (bytes)	Specifies the maximum transmit size of a packet for the ATM-over-ADSL interface.	Type a value from 256 to 9192 .
Encapsulation	Selects the type of encapsulation for traffic on this physical interface.	<p>From the list, select the type of encapsulation for this ATM-over-ADSL interface:</p> <ul style="list-style-type: none"> ■ ATM permanent virtual circuits—Use this type of encapsulation for PPP over ATM (PPPoA) over ADSL interfaces. This is the default encapsulation for ATM-over-ADSL interfaces. ■ Ethernet over ATM encapsulation—Use this type of encapsulation for PPP over Ethernet (PPPoE) over ATM-over-ADSL interfaces that carry IPv4 traffic.
VPI	Configures the ATM virtual path identifier for the interface.	Type a VPI value between 0 and 255 .
Configuring ADSL Options		

Table 43: ATM-over-ADSL Interface Quick Configuration Pages Summary (*continued*)

Field	Function	Your Action
Operating Mode	Specifies the type of DSL operating mode for the ATM-over-ADSL interface.	<p>From the list, select one of the following types of DSL operating modes—for example auto.</p> <p>For Annex A or Annex B, select one of the following:</p> <p>For Annex A only, select one of the following:</p> <p>For Annex B only, select one of the following:</p> <ul style="list-style-type: none"> ■ auto—Configure the ADSL interface to autonegotiate settings with the DSLAM located at the central office. For Annex A, the ADSL interface trains in either ANSI T1.413 Issue II mode or ITU G.992.1 mode. For Annex B, the ADSL interface trains in ITU G.992.1 mode. ■ itu-dmt—Configure the ADSL interface to train in ITU G.992.1 mode. ■ adsl2plus—Configure the ADSL interface to train in ITU G.992.5 mode. You can configure this mode only when it is supported on the DSLAM. ■ itu-dmt-bis—Configure the ADSL interface to train in ITU G.992.3 mode. You can configure this mode only when it is supported on the DSLAM. ■ ansi-dmt—Configure the ADSL interface to train in the ANSI T1.413 Issue II mode. ■ etsi—Configure the ADSL line to train in the ETSI TS 101 388 V1.3.1 mode. ■ itu-annexb-ur2—Configure the ADSL line to train in the G.992.1 Deutsche Telekom UR-2 mode. ■ itu-annexb-non-ur2—Configure the ADSL line to train in the G.992.1 Non-UR-2 mode.

Adding an ATM-over-ADSL Network Interface with a Configuration Editor

To configure ATM-over-ADSL network interfaces for the Services Router with a configuration editor:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 44 on page 165.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:

- To enable authentication on the interface, see “Configuring CHAP on DSL Interfaces (Optional)” on page 179.
 - To configure PPP over Ethernet (PPPoE) encapsulation on an Ethernet interface or on an ATM-over-ADSL interface, see “Configuring Point-to-Point Protocol over Ethernet” on page 191.
5. To check the configuration, see “Verifying DSL Interface Configuration” on page 180.

Table 44: Adding an ATM-over-ADSL Network Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	<p>From the [edit] hierarchy level, enter</p> <p>edit interfaces at-2/0/0</p>
Create the new interface—for example, at-2/0/0.	<ol style="list-style-type: none"> 1. Next to Interface, click Add new entry. 2. In the Interface name box, type at-2/0/0. 3. Click OK. 	
Configuring Physical Properties		

Table 44: Adding an ATM-over-ADSL Network Interface *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure ATM virtual path identifier (VPI) options for the interface—for example, at-2/0/0 .	1. In the Interface name box, select at-2/0/0 .	1. To configure the VPI value, enter
■ ATM VPI—A number between 0 and 255—for example, 25.	2. Next to Atm options , click Configure .	<code>set atm-options vpi 25</code>
■ Operation, Maintenance, and Administration (OAM) F5 loopback cell thresholds (“liveness”) on ATM virtual circuits. The range is between 1 and 255, and the default is 5 cells.	3. Next to Vpi , click Add new entry .	2. To configure OAM liveness values on a VPI, enter
■ Down count—Number of consecutive OAM loopback cells an ATM virtual circuit must lose to be identified as unavailable—for example, 200.	4. In the Vpi number box, type 25.	<code>set atm-options vpi 25</code> <code>oam-liveness up-count 200</code> <code>down-count 200</code>
■ Up count—Number of consecutive OAM loopback cells an ATM virtual interface must receive to be identified as operational—for example, 200.	5. Click OK .	3. To configure the OAM period, enter
■ OAM period—Interval, in seconds, at which OAM cells are transmitted on ATM virtual circuits—for example, 100. The range is between 1 and 900 seconds.	6. In the Actions box, click Edit .	<code>set atm-options vpi 25</code> <code>oam-period 100</code>
	7. Next to Oam liveness , click Configure .	
	8. In the Down count box, type 200.	
	9. In the Up count box, type 200.	
	10. Click OK .	
	11. Next to Oam period , click Configure .	
	12. From the Oam period choices list, select Oam period .	
	13. In the Oam period box, type 100.	
	14. Click OK until you return to the Interface page.	

Table 44: Adding an ATM-over-ADSL Network Interface (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the type of DSL operating mode for the ATM-over-ADSL interface—for example auto .	1. Next to Dsl options, click Configure .	Enter
Annex A and Annex B support the following operating modes:	2. From the Operating Mode list, select auto .	set dsl-options operating-mode auto
Annex A supports the following operating modes:	3. Click OK .	
Annex B supports the following operating modes:		
<ul style="list-style-type: none"> ■ auto—Configures the ADSL interface to autonegotiate settings with the DSLAM located at the central office. For Annex A, the ADSL interface trains in either ANSI T1.413 Issue II mode or ITU G.992.1 mode. For Annex B, the ADSL interface trains in ITU G.992.1 mode. ■ itu-dmt—Configures the ADSL interface to train in ITU G.992.1 mode. ■ adsl2plus—Configures the ADSL interface to train in ITU G.992.5 mode. You can configure this mode only when it is supported on the DSLAM. ■ itu-dmt-bis—Configures the ADSL interface to train in ITU G.992.3 mode. You can configure this mode only when it is supported on the DSLAM. ■ ansi-dmt—Configures the ADSL interface to train in the ANSI T1.413 Issue II mode. ■ etsi—Configures the ADSL line to train in the ETSI TS 101 388 V1.3.1 mode. ■ itu-annexb-ur2—Configures the ADSL line to train in the G.992.1 Deutsche Telekom UR-2 mode. ■ itu-annexb-non-ur2—Configures the ADSL line to train in the G.992.1 Non-UR-2 mode. 		
Configure the encapsulation type—for example, ethernet-over-atm .	From the Encapsulation list, select ethernet-over-atm .	Enter
<ul style="list-style-type: none"> ■ atm-pvc—ATM permanent virtual circuits is the default encapsulation for ATM-over-ADSL interfaces. For PPP over ATM (PPPoA) over ADSL interfaces, use this type of encapsulation. ■ ethernet-over-atm—Ethernet over ATM encapsulation. For PPP over Ethernet (PPPoE) over ATM-over-ADSL interfaces that carry IPv4 traffic, use this type of encapsulation. 		set encapsulation ethernet-over-atm
Configuring Logical Properties		

Table 44: Adding an ATM-over-ADSL Network Interface *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
Add the logical interface.	1. Scroll down the page to Unit, and click Add new entry .	Enter
Set a value from 0 and 16385—for example, 3.		set unit 3
Add other values if required by your network.	2. In the Interface unit number box, type 3. 3. Enter other values in the fields required by your network.	
Configure encapsulation for the ATM-for-ADSL logical unit—for example, atm-nlpid .	From the Encapsulation list, select atm-nlpid .	Enter
The following encapsulations are supported on the ATM-over-ADSL interfaces that use inet (IP) protocols only:		set unit 3 encapsulation atm-nlpid
<ul style="list-style-type: none"> ■ atm-vc-mux—Use ATM virtual circuit multiplex encapsulation. ■ atm-nlpid—Use ATM network layer protocol identifier (NLPID) encapsulation. ■ atm-cisco-nlpid—Use Cisco NLPID encapsulation. ■ ether-over-atm-llc—For interfaces that carry IPv4 traffic, use Ethernet over LLC encapsulation. You cannot configure multipoint interfaces if you use this type of encapsulation. 		
The following encapsulations are supported on the ATM-over-ADSL for PPP-over-ATM (PPPoA) interfaces only. (For a sample PPPoA configuration, see “Verifying Interface Configuration” on page 135.)		
<ul style="list-style-type: none"> ■ atm-ppp-llc— AAL5 logical link control (LLC) encapsulation. ■ atm-ppp-vc-mux—Use AAL5 multiplex encapsulation. 		
Other encapsulation types supported on the ATM-over-ADSL interfaces:		
<ul style="list-style-type: none"> ■ 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. ■ atm-snap—Use ATM subnetwork attachment point (SNAP) encapsulation. 		

Table 44: Adding an ATM-over-ADSL Network Interface (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure Operation, Maintenance, and Administration (OAM) options for ATM virtual circuits: <ul style="list-style-type: none"> ■ OAM F5 loopback cell thresholds (“liveness”) on ATM virtual circuits. The range is between 1 and 255, and the default is 5 cells. <ul style="list-style-type: none"> ■ Down count—Number of consecutive OAM loopback cells an ATM virtual circuit must lose to be identified as unavailable—for example, 200. ■ Up count—Number of consecutive OAM loopback cells an ATM virtual interface must receive to be identified as operational—for example, 200. ■ OAM period—Interval, in seconds, at which OAM cells are transmitted on ATM virtual circuits—for example, 100. The range is between 1 and 900 seconds. 	1. Next to Oam liveness, click Configure . 2. In the Down count box, type 200. 3. In the Up count box, type 200. 4. Click OK . 5. Next to Oam period, click Configure . 6. From the Oam period choices list, select Oam period . 7. In the Oam period box, type 100. 8. Click OK .	1. To configure OAM liveness values for an ATM virtual circuit, enter set unit 3 oam-liveness up-count 200 down-count 200 2. To configure the OAM period, enter set unit 3 oam-period 100
Add the Family protocol type—for example, <code>inet</code> .	1. In the Inet box, select Yes and click Configure . 2. Enter the values in the fields required by your network. 3. Click OK .	Enter set unit 3 family inet Commands vary depending on the protocol type.
Configure ATM virtual channel identifier (VCI) options for the interface. <ul style="list-style-type: none"> ■ ATM VCI type—<code>vci</code>. ■ ATM VCI value—A number between 0 and 4089—for example, 35— with VCIs 0 through 31 reserved. 	1. From the Vci Type list, select vci . 2. In the Vci box, type 35. 3. Click OK until you return to the Interfaces page.	1. To configure the VCI value, enter set unit 3 vci 35

Configuring ATM-over-SHDSL Interfaces

Services Routers with G.SHDSL interfaces can use an Asynchronous Transfer Mode (ATM) interface to send network traffic through a point-to-point connection to a DSL access multiplexer (DSLAM).



NOTE: You can configure Services Routers with a G.SHDSL interface for connections through SHDSL only, not for direct ATM connections.

J-series Services Routers with a 2-port G.SHDSL interface installed support the following modes. You can configure only one mode on each interface.

- 2-port two-wire mode (Annex A or Annex B)—Supports autodetection of the line rate or fixed line rates and provides network speeds from 192 Kbps to 2.3 Mbps

in 64-Kbps increments. Two-wire mode provides two separate, slower SHDSL interfaces.

- 1-port four-wire mode (Annex A or Annex B)—Supports fixed line rates only and provides network speeds from 384 Kbps to 4.6 Mbps in 128-Kbps increments, doubling the bandwidth. Four-wire mode provides a single, faster SHDSL interface.

To configure Point-to-Point Protocol (PPP), see the *JUNOS Network Interfaces Configuration Guide*.

You configure the underlying G.SHDSL interface as an ATM interface, with an interface name of **at-pim/O/port**. (For information about interface names, see “Network Interface Naming” on page 47.) Multiple encapsulation types are supported on both the physical and logical ATM-over-SHDSL interface.

This section contains the following topics:

- Configuring an ATM-over-SHDSL Interface with Quick Configuration on page 170
- Adding an ATM-over-SHDSL Interface with a Configuration Editor on page 174

Configuring an ATM-over-SHDSL Interface with Quick Configuration

The ATM-over-SHDSL Quick Configuration pages allow you to configure ATM-over-SHDSL interfaces and SHDSL options.

To configure an ATM-over-SHDSL interface with Quick Configuration:

1. In the J-Web interface, select **Configuration > Quick Configuration > Interfaces**.

A list of the network interfaces installed on the Services Router is displayed. (For information about interface names, see “Network Interface Naming” on page 47.)

2. Select an **at-pim/O/port** interface from the list.

The ATM-over-SHDSL Interface Quick Configuration page is displayed, as shown in Figure 31 on page 171.

Figure 31: ATM-over-SHDSL Interfaces Quick Configuration Main Page

Juniper
NETWORKS

ROUTER - J6300

Monitor Configuration Diagnose Manage Events Alarms Logged in as: regress Help About Logout

Configuration > Quick Configuration > Interfaces

Quick Configuration

Interfaces DSL Physical Interface: 'at-1/0/1'

Logical Interfaces

No logical interfaces configured.

Add...

Physical Interface Description

MTU (bytes) ?

Encapsulation

VPI ?

SHDSL Options

PIC Mode ?

Annex ?

Line Rate ?

Loopback ?

Current SNR Margin

Disable ☐ ?

Value ?

SNEXT SNR Margin

Disable ☐ ?

Value ?

OK Cancel Apply

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. Trademark Notice. Privacy. Juniper your Net.

3. Enter information into the ATM-over-SHDSL Quick Configuration page, as described in Table 45 on page 172.
4. From the ATM-over-SHDSL Quick Configuration main page, click one of the following buttons:
 - To apply the configuration and stay in the ATM-over-SHDSL interface Quick Configuration main page, click **Apply**.
 - To apply the configuration and return to the Interfaces Quick Configuration page, click **OK**.
 - To cancel your entries and return to the Interfaces Quick Configuration page, click **Cancel**.
5. To verify the ATM-over-SHDSL interface properties, see “Verifying DSL Interface Configuration” on page 180.

Table 45: ATM-over-SHDSL Interface Quick Configuration Pages Summary

Field	Function	Your Action
Configuring Logical Interfaces		
Logical Interface Name	Lists the logical interfaces for the ATM-over-SHDSL physical interface.	<p>If you have not added an <i>at-pim/O/port</i> interface, click Add and enter the information required in the Interfaces Quick Configuration fields.</p> <p>If you have already configured a logical interface, select the interface name from the Logical Interface Name list.</p> <p>To delete a logical interface, select the interface and click Delete.</p>
Adding or Editing a Logical Interface		
Add logical interfaces	Defines one or more logical units that you connect to this physical ADSL interface.	Click Add .
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to clearly identify it in monitoring displays.
Encapsulation	Specifies the type of encapsulation on the SHDSL logical interface.	<p>From the list, select one of the following types of encapsulations.</p> <p>For ATM-over-SHDSL interfaces that use <i>inet</i> (IPv4) protocols only, select one of the following:</p> <ul style="list-style-type: none"> ■ Cisco-compatible ATM NLPID—Use Cisco NLPID encapsulation. ■ ATM NLPID—Use ATM network layer protocol identifier (NLPID) encapsulation. ■ ATM PPP over AA5/LLC—Use AAL5 logical link control (LLC) encapsulation. ■ ATM PPP over raw AAL5—Use AAL5 multiplex encapsulation. ■ ATM LLC/SNAP—For interfaces that carry IPv4 traffic, use ATM over logical link control (LLC) encapsulation. You cannot configure multipoint interfaces if you use this type of encapsulation. ■ ATM VC multiplexing—Use ATM virtual circuit multiplex encapsulation. <p>For other encapsulation types on the ATM-over-SHDSL interfaces, select one of the following:</p> <ul style="list-style-type: none"> ■ Ethernet over ATM (LLC/SNAP)—Use ATM subnetwork attachment point (SNAP) encapsulation. ■ PPPoE over ATM (LLC/SNAP)—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.

Table 45: ATM-over-SHDSL Interface Quick Configuration Pages Summary (*continued*)

Field	Function	Your Action
VCI	Configures the ATM virtual circuit identifier (VCI) for the interface.	In the VCI box, type the number for the VCI.
Add IPv4 address prefixes and destinations	Specifies one or more IPv4 addresses and destination addresses.	Click Add .
IPv4 Address Prefix	Specifies an IPv4 address for the interface.	Type an IPv4 address and prefix. For example: 10.10.10.10/24
Destination Address	Specifies the destination address.	1. Type an IPv4 address for the destination. 2. Click OK .
Configuring Physical Properties		
Physical Interface Description	Describes the physical interface description information. (Optional)	Type a description of the interface.
MTU (bytes)	Specifies the maximum transmission unit (MTU) size, in bytes, of a packet on the ATM-over-SHDSL interface.	Type a value for the byte size—for example, 1500.
Encapsulation	Selects the type of encapsulation for traffic on the physical interface.	Select one of the following types of encapsulation: <ul style="list-style-type: none"> ■ ATM permanent virtual circuits—Use this type of encapsulation for PPP over ATM (PPPoA) over SHDSL interfaces. This is the default encapsulation for ATM-over-SHDSL interfaces. ■ Ethernet over ATM encapsulation—Use this type of encapsulation for PPP over Ethernet (PPPoE) over ATM-over-SHDSL interfaces that carry IPv4 traffic.
VPI	Configures the ATM virtual path identifier (VPI) for the interface.	In the VPI field, type a number between 0 and 255—for example, 25.
Configuring SHDSL Options		
PIC Mode	Specifies the mode on the ATM-over-SHDSL interface.	Select either of the following: <ul style="list-style-type: none"> ■ 1-port-atm—1-port four-wire mode ■ 2-port-atm—2-port two-wire mode
Annex	Specifies the type of annex for the interface. Annex defines the System Reference Model for connecting DSL networks to the plain old telephone service (POTS).	Select one of the following: <ul style="list-style-type: none"> ■ Annex A—Used in North American network implementations. ■ Annex B—Used in European network implementations.
Line Rate	Specifies the available line rates, in kilobits per second, to use on an G.SHDSL interface.	Select the appropriate value. For 2-port-atm mode only, you can select auto , which automatically selects a line rate.

Table 45: ATM-over-SHDSL Interface Quick Configuration Pages Summary *(continued)*

Field	Function	Your Action
Loopback	Specifies the type of loopback testing for the interface. Loopback testing is a diagnostic procedure in which a signal is transmitted and returned to the sending device after passing through all or a portion of a network or circuit. The returned signal is compared with the transmitted signal in order to evaluate the integrity of the equipment or transmission path.	Select one of the following: <ul style="list-style-type: none"> ■ local—Used for testing the SHDSL equipment with local network devices. ■ payload—Used to command the remote configuration to send back the received payload. ■ remote—Used to test SHDSL with a remote network configuration.
Current SNR Margin	Specifies the signal-to-noise ratio (SNR) margin or disables SNR.	To disable Current SNR Margin, select Disable .
Disable		To configure a specific value, type a number from 0 to 10—for example, 5.
Value		The range is 0 dB to 10 dB with a default value of 0.
SNEXT SNR Margin	Sets a value, from –10 dB to 10 dB, for the self-near-crosstalk (SNEXT) SNR margin, or disables SNEXT.	To disable SNEXT SNR Margin, select Disable .
Disable		To configure a specific value, type a number from –10 to 10—for example, 5.
Value		

Adding an ATM-over-SHDSL Interface with a Configuration Editor

To configure ATM-over-SHDSL network interfaces for the Services Router with a configuration editor:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 46 on page 175.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To enable authentication on the interface, see “Configuring CHAP on DSL Interfaces (Optional)” on page 179.
 - To configure PPP over Ethernet (PPPoE) encapsulation on an Ethernet interface or on an ATM-over-SHDSL interface, see “Configuring Point-to-Point Protocol over Ethernet” on page 191.
5. To check the configuration, see “Verifying DSL Interface Configuration” on page 180.

Table 46: Adding an ATM-over-SHDSL Network Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Chassis level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Chassis, click Configure. 	<p>From the [edit] hierarchy level, enter</p> <pre>set chassis fpc 6 pic 0 shdsl pic-mode 1-port-atm</pre>
Set the ATM-over-SHDSL mode on the G.SHDSL interface, if required. By default, G.SHDSL interfaces are enabled in two-wire Annex B mode. To configure the four-wire mode on the G.SHDSL interface, follow the tasks in this table.	<ol style="list-style-type: none"> 1. Next to Fpc, click Add new entry. 2. In the Slot box, type 6. 3. Next to Pic, click Add new entry. 4. In the Slot box, type 0. 5. Next to Shdsl, click Configure. 6. From the Pic mode menu, select 1-port-atm. 7. Click OK until you return to the main Configuration page. 	
Navigate to the Interfaces level in the configuration hierarchy.	On the main Configuration page next to Interfaces, click Edit .	<p>From the [edit] hierarchy level, enter</p> <pre>edit interfaces at-2/0/0</pre>
Create the new interface—for example, at-2/0/0.	<ol style="list-style-type: none"> 1. Next to Interface, click Add new entry. 2. In the Interface name box, type at-2/0/0. 3. Click OK. 	
Configuring Physical Properties		

Table 46: Adding an ATM-over-SHDSL Network Interface (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Configure ATM virtual path identifier (VPI) options for the interface—for example, at-2/0/0.</p> <ul style="list-style-type: none"> ■ ATM VPI—A number between 0 and 255—for example, 25. ■ Operation, Maintenance, and Administration (OAM) F5 loopback cell thresholds (“liveness”) on ATM virtual circuits. The range is between 1 and 255, and the default is 5 cells. <ul style="list-style-type: none"> ■ Down count—Number of consecutive OAM loopback cells an ATM virtual circuit must lose to be identified as unavailable—for example, 200. ■ Up count—Number of consecutive OAM loopback cells an ATM virtual interface must receive to be identified as operational—for example, 200. ■ OAM period—Interval, in seconds, at which OAM cells are transmitted on ATM virtual circuits—for example, 100. The range is between 1 and 900 seconds. 	<ol style="list-style-type: none"> 1. In the Interface name box, select at-2/0/0. 2. Next to Atm options, click Configure. 3. Next to Vpi, click Add new entry. 4. In the Vpi number box, type 25. 5. Click OK. 6. In the Actions box, click Edit. 7. Next to Oam liveness, click Configure. 8. In the Down count box, type 200. 9. In the Up count box, type 200. 10. Click OK. 11. Next to Oam period, click Configure. 12. From the Oam period choices list, select Oam period. 13. In the Oam period box, type 100. 14. Click OK until you return to the Interface page. 	<ol style="list-style-type: none"> 1. To configure the VPI value, enter set atm-options vpi 25 2. To configure OAM liveness values on a VPI, enter set atm-options vpi 25 oam-liveness up-count 200 down-count 200 3. To configure the OAM period, enter set atm-options vpi 25 oam-period 100
<p>Configure the encapsulation type—for example, ethernet-over-atm.</p> <ul style="list-style-type: none"> ■ atm-pvc—ATM permanent virtual circuits is the default encapsulation for ATM-over-SHDSL interfaces. For PPP over ATM (PPPoA) over SHDSL interfaces, use this type of encapsulation. ■ ethernet-over-atm—Ethernet over ATM encapsulation. For PPP over Ethernet (PPPoE) over ATM-over-SHDSL interfaces that carry IPv4 traffic, use this type of encapsulation. 	<p>From the Encapsulation list, select ethernet-over-atm.</p>	<p>Enter</p> <p>set encapsulation ethernet-over-atm</p>
<p>Set the annex type.</p> <ul style="list-style-type: none"> ■ Annex A—Used in North American network implementations. ■ Annex B—Used in European network implementations. 	<ol style="list-style-type: none"> 1. Next to Shdsl options, click Configure. 2. From the Annex list, select Annex-a. 	<p>Enter</p> <p>set shdsl-options annex annex-a</p>

Table 46: Adding an ATM-over-SHDSL Network Interface (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Configure the SHDSL line rate for the ATM-over-SHDSL interface—for example, automatic selection of the line rate.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> ■ auto—Automatically selects a line rate. This option is available only in two-wire mode and is the default value. ■ 192 Kbps or higher—Speed of transmission of data on the SHDSL connection. <p>In the four-wire mode, the default line rate is 4608 Kbps.</p>	<p>From the Line Rate list, select auto.</p>	<p>Enter</p> <p>set shdsl-options line-rate auto</p>
<p>Configure the loopback option for testing the SHDSL connection integrity—for example, local loopback.</p> <p>The following values are available:</p> <ul style="list-style-type: none"> ■ local—Used for testing the SHDSL equipment with local network devices. ■ payload—Used to command the remote configuration to send back the received payload. ■ remote—Used to test SHDSL with a remote network configuration. 	<p>From the Loopback list, select local.</p>	<p>Enter</p> <p>set shdsl-options loopback local</p>
<p>Configure the signal-to-noise ratio (SNR) margin—for example, 5 dB for either or both of the following thresholds:</p> <ul style="list-style-type: none"> ■ current—Line trains at higher than current noise margin plus SNR threshold. The range is 0 to 10 dB. The default value is 0. ■ snext—Line trains at higher than self-near-end crosstalk (SNEXT) threshold. The default value is disabled. <p>Setting 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. You can also disable the SNR margin thresholds.</p>	<ol style="list-style-type: none"> Next to Snr margin, select Yes, then click Configure. From the Current list, select Enter Specific Value. In the Value box, type 5. From the Snext list, select Enter Specific Value. In the Value box, type 5. Click OK until you return to the Interface page. 	<ol style="list-style-type: none"> Enter set shdsl-options snr-margin current 5 Enter set shdsl-options snr-margin snext 5
Configuring Logical Properties		
<p>Add the logical interface.</p> <p>Set a value from 0 and 16385—for example, 3.</p> <p>Add other values if required by your network.</p>	<ol style="list-style-type: none"> Scroll down the page to Unit, and click Add new entry. In the Interface unit number box, type 3. Enter other values in the fields required by your network. 	<p>Enter</p> <p>set unit 3</p>

Table 46: Adding an ATM-over-SHDSL Network Interface *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Configure encapsulation for the ATM-for-SHDSL logical unit—for example, <code>atm-nlpid</code>.</p> <p>The following encapsulations are supported on the ATM-over-SHDSL interfaces that use <code>inet</code> (IP) protocols only:</p> <ul style="list-style-type: none"> ■ <code>atm-vc-mux</code>—Use ATM virtual circuit multiplex encapsulation. ■ <code>atm-nlpid</code>—Use ATM network layer protocol identifier (NLPID) encapsulation. ■ <code>atm-cisco-nlpid</code>—Use Cisco NLPID encapsulation. ■ <code>ether-over-atm-llc</code>—For interfaces that carry IPv4 traffic, use Ethernet over LLC encapsulation. You cannot configure multipoint interfaces if you use this type of encapsulation. <p>The following encapsulations are supported on the ATM-over-SHDSL for PPP-over-ATM (PPPoA) interfaces only. (For a sample PPPoA configuration, see “Verifying Interface Configuration” on page 135.)</p> <ul style="list-style-type: none"> ■ <code>atm-ppp-llc</code>—AAL5 logical link control (LLC) encapsulation. ■ <code>atm-ppp-vc-mux</code>—Use AAL5 multiplex encapsulation. <p>Other encapsulation types supported on the ATM-over-SHDSL interfaces:</p> <ul style="list-style-type: none"> ■ <code>ppp-over-ether-over-atm-llc</code>—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. ■ <code>atm-snap</code>—Use ATM subnetwork attachment point (SNAP) encapsulation. 	<p>From the Encapsulation list, select atm-nlpid.</p>	<p>Enter</p> <p>set unit 3 encapsulation atm-nlpid</p>
<p>Configure Operation, Maintenance, and Administration (OAM) options for ATM virtual circuits:</p> <ul style="list-style-type: none"> ■ OAM F5 loopback cell thresholds (“liveness”) on ATM virtual circuits. The range is between 1 and 255, and the default is 5 cells. <ul style="list-style-type: none"> ■ Down count—Number of consecutive OAM loopback cells an ATM virtual circuit must lose to be identified as unavailable—for example, 200. ■ Up count—Number of consecutive OAM loopback cells an ATM virtual interface must receive to be identified as operational—for example, 200. ■ OAM period—Interval, in seconds, at which OAM cells are transmitted on ATM virtual circuits—for example, 100. The range is between 1 and 900 seconds. 	<ol style="list-style-type: none"> Next to Oam liveness, click Configure. In the Down count box, type 200. In the Up count box, type 200. Click OK. Next to Oam period, click Configure. From the Oam period choices list, select Oam period. In the Oam period box, type 100. Click OK. 	<ol style="list-style-type: none"> To configure OAM liveness values for an ATM virtual circuit, enter <p>set unit 3 oam-liveness up-count 200 down-count 200</p> To configure the OAM period, enter <p>set unit 3 oam-period 100</p>

Table 46: Adding an ATM-over-SHDSL Network Interface (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
Add the Family protocol type—for example, inet.	<ol style="list-style-type: none"> 1. In the Inet box, select Yes and click Configure. 2. Enter the values in the fields required by your network. 3. Click OK. 	<p>Enter</p> <p>set unit 3 family inet</p> <p>Commands vary depending on the protocol type.</p>
Configure ATM virtual channel identifier (VCI) options for the interface. <ul style="list-style-type: none"> ■ ATM VCI type—vci. ■ ATM VCI value—A number between 0 and 4089—for example, 35—with VCIs 0 through 31 reserved. 	<ol style="list-style-type: none"> 1. From the Vci type list, select vci. 2. In the Vci box, type 35. 3. Click OK until you return to the Interfaces page. 	<ol style="list-style-type: none"> 1. To configure the VCI value, enter set unit 3 vci 35

Configuring CHAP on DSL Interfaces (Optional)

For interfaces with PPPoA encapsulation, you can optionally 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 set the **passive** option to handle incoming CHAP packets only, the interface does not challenge its peer. However, if the interface is challenged, it responds to the challenge. If you do not set the **passive** option, the interface always challenges its peer.

For more information about CHAP, see the *JUNOS Network Interfaces Configuration Guide*.

To configure CHAP on the ATM-over-ADSL or ATM-over-SHDSL interface:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 47 on page 180.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying DSL Interface Configuration” on page 180.

Table 47: Configuring CHAP

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Access level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Access, click Configure or Edit. 	From the [edit] hierarchy level, enter edit access
Define a CHAP access profile—for example, A-ppp-client—with a client named client 1 and the secret (password) my-secret.	<ol style="list-style-type: none"> 1. Next to Profile, click Add new entry. 2. In the Profile name box, type A-ppp-client. 3. Next to Client, click Add new entry. 4. In the Name box, type client1. 5. In the Chap secret box, type my-secret. 6. Click OK until you return to the main Configuration page. 	Enter set profile A-ppp-client client client1 chap-secret my-secret.
Navigate to the appropriate ATM interface level in the configuration hierarchy—for example, at-3/0/0 unit 0.	<ol style="list-style-type: none"> 1. On the main Configuration page next to Interfaces, click Configure or Edit. 2. In the Interface name box, click at-3/0/0. 3. In the Interface unit number box, click 0. 	From the [edit] hierarchy level, enter edit interfaces at-3/0/0 unit 0
Configure CHAP on the ATM-over-ADSL or ATM-over-SHDSL interface and specify a unique profile name containing a client list and access parameters—for example, A-ppp-client.	<ol style="list-style-type: none"> 1. Next to Ppp options, click Configure. 2. Next to Chap, click Configure. 3. In the Access profile box, type A-ppp-client. 	Enter set ppp-options chap access-profile A-ppp-client
Specify a unique hostname to be used in CHAP challenge and response packets—for example, A-at-3/0/0.0.	In the Local name box, type, A-at-3/0/0.0	Enter set ppp-options chap local-name A-at-3/0/0.0.
Set the passive option to handle incoming CHAP packets only.	<ol style="list-style-type: none"> 1. In the Passive box, click Yes. 2. Click OK. 	Enter set ppp-options chap passive

Verifying DSL Interface Configuration

To verify ATM-over-ADSL or ATM-over-SHDSL, perform these tasks:

- Verifying ADSL Interface Properties on page 181
- Displaying a PPPoA Configuration for an ATM-over-ADSL Interface on page 185
- Verifying an ATM-over-SHDSL Configuration on page 186

Verifying ADSL Interface Properties

Purpose Verify that the interface properties are correct.

Action From the CLI, enter the show interfaces *interface-name* extensive command.

```

user@host> show interfaces at-3/0/0 extensive
Physical interface: at-3/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 23, Generation: 48
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, ADSL mode, Speed: ADSL,

  Loopback: None
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 8 supported
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:c7:44:3c
  Last flapped   : 2005-05-16 05:54:41 PDT (00:41:42 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :                4520                0 bps
    Output bytes  :               39250                0 bps
    Input packets :                 71                0 pps
    Output packets:               1309                0 pps
  Input errors:
    Errors: 0, Drops: 0, Invalid VCs: 0, Framing errors: 0, Policed discards: 0,

    L3 incompletes: 0, L2 channel errors: 1, L2 mismatch timeouts: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,

    Resource errors: 0
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets

    0 best-effort           4                4                0
    1 expedited-fo          0                0                0
    2 assured-forw          0                0                0
    3 network-cont         2340            2340                0

  ADSL alarms   : LOS, LOM, LOCDNI, FAR_LOF, FAR_LOS, FAR_LOCDNI
  ADSL defects  : LOF, LOS, LOCDNI, FAR_LOF, FAR_LOS, FAR_LOCDNI
  ADSL media:
    Seconds      Count  State
    LOF          239206    2 OK
    LOS          239208    1 OK
    LOM           3         1 OK
    LOP           0         0 OK
    LOCDI         3         1 OK
    LOCDNI       239205    1 OK
  ADSL status:
    Modem status : Showtime
    DSL mode      : Auto    Annex A
    Last fail code: ATU-C not detected
  ADSL Statistics:
    ATU-R          ATU-C
    Attenuation (dB) :          0.5          0.0
    Capacity used (%) :           81          72
    Noise margin (dB) :           9.0          9.5

```

```

Output power (dBm)      :              7.5              8.5

                                Interleave      Fast  Interleave      Fast
Bit rate (kbps)   :              0      8128              0      896
CRC               :              0        3              0        0
FEC               :              0        0              0        0
HEC               :              0        3              0        0
Received cells    :              0      287
Transmitted cells :              0     4900
Bit error rate    :              0        0

ATM status:
HCS state:      Hunt
LOC           :      OK
ATM Statistics:
Uncorrectable HCS errors: 0, Correctable HCS errors: 0, Tx cell FIFO overruns:
0,
Rx cell FIFO overruns: 0, Rx cell FIFO underruns: 0, Input cell count: 0,
Output cell count: 0, Output idle cell count: 0, Output VC queue drops: 0,
Input no buffers: 0, Input length errors: 0, Input timeouts: 0, Input invalid
VCs: 0,
Input bad CRCs: 0, Input OAM cell no buffers: 0
Packet Forwarding Engine configuration:
Destination slot: 3
CoS transmit queue      Bandwidth      Buffer Priority
Limit
                                %      bps      %      bytes
0 best-effort           95      7600000   95      0      low
none
3 network-control       5       400000    5      0      low
none

Logical interface at-3/0/0.0 (Index 66) (SNMP ifIndex 28) (Generation 23)
Flags: Point-To-Point SNMP-Traps 16384 Encapsulation: ATM-PPP-LLC
Traffic statistics:
Input bytes   :      2432
Output bytes  :         0
Input packets :      116
Output packets:         0
Local statistics:
Input bytes   :      1810
Output bytes  :         0
Input packets :       78
Output packets:         0
Transit statistics:
Input bytes   :       622      0 bps
Output bytes  :         0      0 bps
Input packets :       38      0 pps
Output packets:         0      0 pps
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive statistics:
Input : 33 (last seen 00:00:03 ago)
Output: 34 (last sent 00:00:03 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
CHAP state: Success
Protocol inet, MTU: 4470, Generation: 24, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 155.55.5.1, Local: 155.55.5.2, Broadcast: Unspecified,
Generation: 45

```


VCI 0.35

Flags: Active, 1024

Total down time: 0 sec, Last down: Never

ATM per-VC transmit statistics:

Tail queue packet drops: 0

Traffic statistics:

Input bytes : 2432

Output bytes : 0

Input packets: 116

Output packets: 0

Logical interface at-3/0/0.32767 (Index 69) (SNMP ifIndex 25) (Generation 21)

Flags: Point-To-Multipoint No-Multicast SNMP-Traps 16384 Encapsulation:

ATM-VCMUX

Traffic statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

Local statistics:

Input bytes : 0

Output bytes : 0

Input packets: 0

Output packets: 0

VCI 0.4

Flags: Active, 1024

Total down time: 0 sec, Last down: Never

ATM per-VC transmit statistics:

Tail queue packet drops: 0

Traffic statistics:

Input bytes : 208

Output bytes : 208

Input packets: 4

Output packets: 4

What It Means The output shows a summary of interface information. Verify the following information:

- The physical interface is **Enabled**. If the interface is shown as **Disabled**, do either of the following:
 - In the CLI configuration editor, delete the **disable** statement at the [edit interfaces *interface-name*] level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the **Interfaces > interface-name** page.
- The physical link is **Up**. A link state of **Down** indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The **Last Flapped** time is an expected value. The **Last Flapped** time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics interface-name** command.
- No ADSL alarms and defects appear that can render the interface unable to pass packets. When a defect persists for a certain amount of time, it is promoted to an alarm. The following are ADSL-specific alarms:
 - LOCDI—Loss of cell delineation for interleaved channel
 - LOCDNI—Loss of cell delineation for non-interleaved channel
 - LOF—Loss of frame
 - LOM—Loss of multiframe
 - LOP—Loss of power
 - LOS—Loss of signal
 - FAR_LOF—Loss of frame in ADSL transceiver unit-central office (ATU-C)
 - FAR_LOS—Loss of signal in ATU-C
 - FAR_LOCDI—Loss of cell delineation for interleaved channel in ATU-C
 - FAR_LOCDNI—Loss of cell delineation for non-interleaved channel in ATU-C

Examine the operational statistics for an ADSL interface. Statistics in the **ATU-R** (ADSL transceiver unit-remote) column are for the near end. Statistics in the **ATU-C** (ADSL transceiver unit-central office) column are for the far end.

- **Attenuation (dB)**—Reduction in signal strength measured in decibels.
- **Capacity used (%)**—Amount of ADSL usage in %.
- **Noise Margin (dB)**—Maximum extraneous signal allowed without causing the output to deviate from an acceptable level.
- **Output Power (dBm)**—Amount of power used by the ADSL interface.

- Bit Rate (kbps)—Data transfer speed on the ADSL interface.

See Also For a complete description of show interfaces extensive output, see the *JUNOS Interfaces Command Reference*.

Displaying a PPPoA Configuration for an ATM-over-ADSL Interface

Purpose Verify the PPPoA configuration for an ATM-over-ADSL interface.

Action From the J-Web interface, select **Configuration > View and Edit > View Configuration Text**. Alternatively, from configuration mode in the CLI, enter the show interfaces *interface-name* and the show access commands from the top level.

```
[edit]
user@host# show interfaces at-3/0/0
at-3/0/0 {
  encapsulation atm-pvc;
  atm-options {
    vpi 0;
  }
  dsl-options {
    operating-mode auto;
  }
  unit 0 {
    encapsulation atm-ppp-llc;
    vci 0.100;
    ppp-options {
      chap {
        access-profile A-ppp-client;
        local-name A-at-3/0/0.0;
        passive;
      }
    }
    family inet {
      negotiate address;
    }
  }
}
user@host# show access
profile A-ppp-client {
  client A-ppp-server chap-secret "$9$G4ikPu0ISyKP5clKv7Nik.PT3"; ## SECRET-DATA
}
```

What It Means Verify that the output shows the intended configuration of PPPoA.

See Also For more information about the format of a configuration file, see “Viewing the Configuration Text” on page 9.

Verifying an ATM-over-SHDSL Configuration

Purpose Verify that the interface properties are correct.

Action From the CLI, enter the show interfaces *interface-name* extensive command.

```

user@host> show interfaces at-6/0/0 extensive
Physical interface: at-6/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 23, Generation: 48
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, ADSL mode, Speed: ADSL,

  Loopback: None
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 8 supported
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:05:85:c7:44:3c
  Last flapped   : 2005-05-16 05:54:41 PDT (00:41:42 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :                4520                0 bps
    Output bytes  :               39250                0 bps
    Input packets :                 71                0 pps
    Output packets:               1309                0 pps
  Input errors:
    Errors: 0, Drops: 0, Invalid VCs: 0, Framing errors: 0, Policed discards: 0,

    L3 incompletes: 0, L2 channel errors: 1, L2 mismatch timeouts: 0, Resource
errors: 0
  Output errors:
    Carrier transitions: 3, Errors: 0, Drops: 0, Aged packets: 0, MTU errors: 0,

    Resource errors: 0
  Queue counters:
    Queued packets  Transmitted packets  Dropped packets

    0 best-effort           4                4                0
    1 expedited-fo         0                0                0
    2 assured-forw         0                0                0
    3 network-cont        2340            2340                0

  SHDSL alarms   : None
  SHDSL defects  : None
  SHDSL media:
    Seconds      Count  State
    LOSD         239206    2  OK
    LOSW         239208    1  OK
    ES            3         1  OK
    SES           0         0  OK
    UAS           3         1  OK

  SHDSL status:
    Line termination :STU-R
  Annex          :Annex B
  Line Mode       :2-wire
  Modem Status    :Data
  Last fail code  :0
  Frammer mode    :ATM
  Dying Gasp      :Enabled

```

```
Chipset version :1
Firmware version :R3.0
SHDSL Statistics:
  Loop Attenuation (dB) :0.600
  Transmit power (dB) :8.5
  Receiver gain (dB) :21.420
  SNR sampling (dB) :39.3690
  Bit rate (kbps) :2304
  Bit error rate :0
  CRC errors :0
  SEGA errors :1
  LOSW errors :0
  Received cells :1155429
  Transmitted cells :1891375
  HEC errors :0
  Cell drop :0
```

What It Means The output shows a summary of interface information. Verify the following information:

- The physical interface is **Enabled**. If the interface is shown as **Disabled**, do either of the following:
 - In the CLI configuration editor, delete the **disable** statement at the [edit interfaces *interface-name*] level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the **Interfaces > interface-name** page.
- The physical link is **Up**. A link state of **Down** indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The **Last Flapped** time is an expected value. The **Last Flapped** time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics interface-name** command.
- No SHDSL alarms and defects appear that can render the interface unable to pass packets. When a defect persists for a certain amount of time, it is promoted to an alarm.
 - **LOS**—Loss of signal. No signal was detected on the line.
 - **LOSW**—Loss of sync word. A message ID was sent.
 - **Power status**—A power failure has occurred.
 - **LOSD**—Loss of signal was detected at the remote application interface.
 - **ES**—Errored seconds. One or more cyclic redundancy check (CRC) anomalies were detected.
 - **SES**—Severely errored seconds. At least 50 CRC anomalies were detected.
 - **UAS**—Unavailable seconds. An interval has occurred during which one or more **LOSW** defects were detected.

Examine the SHDSL interface status:

- **Line termination**—SHDSL transceiver unit-remote (STU-R). (Only customer premises equipment is supported.)
- **Annex**—Either Annex A or Annex B. Annex A is supported in North America, and Annex B is supported in Europe.
- **Line Mode**—SHDSL mode configured on the G.SHDSL interface pair, either 2-wire or 4-wire.
- **Modem Status**—Data. Sending or receiving data.
- **Last fail code**—Code for the last interface failure.
- **Framer mode**—Framer mode of the underlying interface: ATM.

- **Dying Gasp**—Ability of a J-series router that has lost power to send a message informing the attached DSL access multiplexer (DSLAM) that it is about to go offline.
- **Chipset version**—Version number of the chipset on the interface
- **Firmware version**—Version number of the firmware on the interface.

Examine the operational statistics for a SHDSL interface.

- **Loop Attenuation (dB)**—Reduction in signal strength measured in decibels.
- **Transmit power (dB)**—Amount of SHDSL usage in %.
- **Receiver gain (dB)**—Maximum extraneous signal allowed without causing the output to deviate from an acceptable level.
- **SNR sampling (dB)**—Signal-to-noise ratio at a receiver point, in decibels.
- **Bit Rate (kbps)**—Data transfer speed on the SHDSL interface.
- **CRC errors**—Number of cyclic redundancy check errors.
- **SEGA errors**—Number of segment anomaly errors. A regenerator operating on a segment received corrupted data.
- **LOSW errors**—Number of loss of signal defect errors. Three or more consecutively received frames contained one or more errors in the framing bits.
- **Received cells**—Number of cells received through the interface.
- **Transmitted cells**—Number of cells sent through the interface.
- **HEC errors**—Number of header error checksum errors.
- **Cell drop**—Number of dropped cells on the interface.

See Also For a complete description of `show interfaces` extensive output, see the *JUNOS Interfaces Command Reference*.

Chapter 6

Configuring Point-to-Point Protocol over Ethernet

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—a J-series Services Router. Hosts share a common digital subscriber line (DSL), a cable modem, or a wireless connection to the Internet. To use PPPoE, you must initiate a PPPoE session, encapsulate Point-to-Point Protocol (PPP) packets over Ethernet, and configure the Services Router as a PPPoE client.



NOTE: Services Routers with asymmetric digital subscriber line (ADSL) or symmetric high-speed DSL (SHDSL) interfaces can use PPPoE over Asynchronous Transfer Mode (ATM) to connect through DSL lines only, not for direct ATM connections.

You can use the J-Web Quick Configuration, J-Web configuration editor, or CLI configuration editor to configure PPPoE.

This chapter contains the following topics:

- PPPoE Terms on page 191
- PPPoE Overview on page 192
- Before You Begin on page 195
- Configuring PPPoE Interfaces with Quick Configuration on page 195
- Configuring PPPoE with a Configuration Editor on page 198
- Verifying a PPPoE Configuration on page 205

PPPoE Terms

Before configuring PPPoE on a Services Router, become familiar with the terms defined in Table 48 on page 192.

Table 48: PPPoE Terms

Term	Definition
access concentrator	Router that acts as a server in a PPPoE session—for example, an E-series router.
customer premises equipment (CPE)	Router that acts as a PPPoE client in a PPPoE session—for example, a Services Router.
Logical Link Control (LLC)	Encapsulation protocol that allows transport of multiple protocols over a single ATM virtual connection.
Point-to-Point Protocol (PPP)	Encapsulation protocol for transporting IP traffic over point-to-point links.
PPP over Ethernet (PPPoE)	Network protocol that encapsulates PPP frames in Ethernet frames and connects multiple hosts over a simple bridging access device to a remote access concentrator.
PPPoE Active Discovery Initiation (PADI) packet	Initiation packet that is broadcast by the client to start the discovery process.
PPPoE Active Discovery Offer (PADO) packet	Offer packets sent to the client by one or more access concentrators in reply to a PADI packet.
PPPoE Active Discovery Request (PADR) packet	Packet sent by the client to one selected access concentrator to request a session.
PPPoE Active Discovery Session-Confirmation (PADS) packet	Packet sent by the selected access concentrator to confirm the session.
PPPoE Active Discovery Termination (PADT) packet	Packet sent by either the client or the access concentrator to terminate a session.
PPPoE over ATM	Network protocol that encapsulates PPPoE frames in Asynchronous Transfer Mode (ATM) frames for asymmetric digital subscriber line (ADSL) or symmetric high-speed DSL (SHDSL) transmission, and connects multiple hosts over a simple bridging access device to a remote access concentrator.
virtual path identifier (VPI)	An identifier of the virtual path that establishes a route between two devices in a network.
virtual channel identifier (VCI)	An identifier of the virtual channel inside a virtual path. Each virtual path identifier (VPI) can contain multiple virtual channels, each represented by a VCI.

PPPoE Overview

On the Services Router, PPPoE establishes a point-to-point connection between the client (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, Gigabit Ethernet, ATM-over-ADSL, or ATM-over-SHDSL.

interface. PPPoE is easy to configure and allows 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 193
- PPPoE Stages on page 194
- Optional CHAP Authentication on page 194

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 Ethernet, use a PPPoE encapsulation.
- If the interface is ATM-over-ADSL or ATM-over-SHDSL, use a PPPoE over ATM encapsulation.

Ethernet Interface

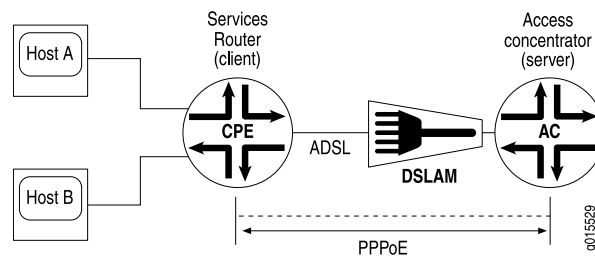
The Services Router encapsulates each PPP frame in an Ethernet frame and transports the frames over an Ethernet loop. Figure 32 on page 193 shows a typical PPPoE session between a Services Router and an access concentrator on the Ethernet loop.

Figure 32: PPPoE Session on the Ethernet Loop



ATM-over-ADSL or ATM-over-SHDSL Interface

When an ATM network is configured with a point-to-point connection, PPPoE can use ATM Adaptation Layer 5 (AAL5) for framing PPPoE-encapsulated packets. The AAL5 protocol provides a virtual connection between the client and the server within the same network. The Services Router encapsulates each PPPoE frame in an ATM frame and transports each frame over an ADSL or SHDSL loop and a digital subscriber line access multiplexer (DSLAM). For example, Figure 33 on page 194 shows a typical PPPoE over ATM session between a Services Router and an access concentrator on an ADSL loop.

Figure 33: PPPoE Session on an ADSL 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.

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 a PPPoE Active Discovery Session-Confirmation (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 per Services Router.

Each PPPoE session is uniquely identified by the Ethernet address of the peer and the session ID.

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 set the **passive** option to handle incoming CHAP packets only, the interface does not challenge its peer. However, if the interface is challenged, it responds to the challenge. If you do not set the **passive** option, the interface always challenges its peer.

You can configure Remote Authentication Dial-In User Service (RADIUS) authentication of PPP sessions using CHAP. CHAP enables you to send RADIUS messages through a routing instance to customer RADIUS servers in a private network. For more information, see the *JUNOS System Basics Configuration Guide*.

For more information about CHAP, see the *JUNOS Network Interfaces Configuration Guide*.

Before You Begin

Before you begin configuring PPPoE, complete the following tasks:

- Establish basic connectivity. See the Getting Started Guide for your router.
- Configure network interfaces. See “Configuring a Fast Ethernet Interface with Quick Configuration” on page 112, “Configuring Gigabit Ethernet Interfaces with Quick Configuration” on page 115, or “Configuring Digital Subscriber Line Interfaces” on page 157.

Configuring PPPoE Interfaces with Quick Configuration

To configure properties on a PPPoE interface:

1. In the J-Web user interface, select **Configuration > Quick Configuration > Interfaces**.

A list of the network interfaces present on the Services Router is displayed, as shown in Figure 22 on page 104. (For information about interface names, see “Network Interface Naming” on page 47.) The third column indicates whether the interface has been configured.

2. Select **pp0**.

The PPPoE Interfaces Quick Configuration main page is displayed, as shown in Figure 34 on page 196.

Figure 34: PPPoE Interfaces Quick Configuration Main Page

Juniper
NETWORKS

ROUTER - J6300

Monitor Configuration Diagnose Manage Events Alarms Logged in as: regress Help About Logout

Configuration > Quick Configuration > Interfaces

Quick Configuration

View and Edit
History
Rescue

Quick Configuration

Interfaces [Add a PPPoE Logical Interface](#)

Interface Information

Logical Interface Description

IPv4 Addresses and Prefixes

PPP Options

Enable CHAP ☐

CHAP Local Identity

Use System Host Name ☐

Local Name

* CHAP Peer Identity

* CHAP Secret

PPPoE Options

Access Concentrator

Auto Reconnect Time

Idle Timeout

Service Name

Underlying Interface

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. Trademark Notice. Privacy. **Juniper Your Net.**

3. Enter information into the Quick Configuration pages, as described in Table 49 on page 197.
4. From the PPPoE Interfaces Quick Configuration main page, click one of the following buttons:
 - To apply the configuration and stay on the PPPoE Quick Configuration main page, click **Apply**.
 - To apply the configuration and return to the Interfaces Quick Configuration page, click **OK**.
 - To cancel your entries and return to the Interfaces Quick Configuration page, click **Cancel**.
5. To verify that the PPPoE interface is configured correctly, see “Verifying a PPPoE Configuration” on page 205.

Table 49: PPPoE Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
Logical Interfaces	Lists the logical interfaces for the PPPoE physical interface.	<ul style="list-style-type: none"> ■ To add a logical interface, click Add. ■ To edit a logical interface, select the interface from the list. ■ To delete a logical interface, select the check box next to the name and click Delete.
Add logical interfaces	Defines one or more logical units that you connect to this physical PPPoE interface. You must define at least one logical unit for a PPPoE interface.	Click Add .
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	<ol style="list-style-type: none"> 1. Type one or more IPv4 addresses and prefixes. For example: 10.10.10.10/24 2. Click Add. 3. Click OK.
Physical Interface Description	(Optional) Adds supplementary information about the physical PPPoE interface.	Type a text description of the PPPoE interface to more clearly identify it in monitoring displays.
PPP Options		
Enable CHAP	Enables or disables CHAP authentication on a PPPoE interface.	<ul style="list-style-type: none"> ■ To enable CHAP, select the check box. ■ To disable CHAP, clear the check box.
CHAP Local Identity (available if CHAP is enabled)		
Use System Host Name	Specifies that the PPPoE interface uses the Services Router's system hostname in CHAP challenge and response packets.	<ul style="list-style-type: none"> ■ To enable, select the check box (the default). ■ To disable, clear the check box.
Local Name	If Use System Host Name is disabled, specifies the local name for CHAP to use.	Type a local name for this PPPoE interface.
CHAP Peer Identity (required if CHAP is enabled)	Identifies the client or peer with which the Services Router communicates on this PPPoE interface.	Type the CHAP client name.
CHAP Secret (required if CHAP is enabled)	Specifies the secret password for CHAP authentication, known to both sides of the connection.	Type a password that is known to the other side of the connection. Use a combination of letters and numbers that is difficult for others to guess.

Table 49: PPPoE Quick Configuration Summary *(continued)*

Field	Function	Your Action
PPPoE Options		
Access Concentrator	Identifies the access concentrator by a unique name.	Type a name for the access concentrator—for example, <code>ispl.com</code> .
Auto Reconnect Time	Specifies the number of seconds to wait before reconnecting after a PPPoE session is terminated.	Type a value from 1 through 4294947295 for automatic reconnection—for example, 100 seconds. Type 0 (the default) for immediate reconnection.
Idle Timeout	Specifies the number of seconds a session can be idle without disconnecting.	Type a value for the timeout. Type 0 if you do not want the session to time out.
Service Name	Identifies the type of service provided by the access concentrator, such as the name of the Internet service provider (ISP), class, or quality of service.	Type the type of service provided by the access concentrator. For example, <code>video@ispl.com</code> .
Underlying Interface	Specifies the logical Ethernet interface or the logical ATM interface as the underlying interface for the PPPoE session.	From the list, select the underlying interface for the PPPoE session—for example, <code>ge-0/0/1.0</code> or <code>at-2/0/0.0</code> . For information about interface names, see “Network Interface Naming” on page 47.

Configuring PPPoE with a Configuration Editor

To configure PPPoE on a Services Router, you must perform the following tasks marked *(Required)*:

- Setting the Appropriate Encapsulation on the Interface (Required) on page 198
- Configuring PPPoE Interfaces (Required) on page 201
- Configuring CHAP on a PPPoE Interface (Optional) on page 204

Setting the Appropriate Encapsulation on the Interface (Required)

For PPPoE on an Ethernet interface, you must configure encapsulation on the logical interface. To configure encapsulation on an Ethernet logical interface, use PPP over Ethernet encapsulation.

For PPPoE on an ATM-over-ADSL or ATM-over-SHDSL interface, you must configure encapsulation on both the physical and logical interfaces. To configure encapsulation on an ATM-over-ADSL or ATM-over-SHDSL physical interface, use Ethernet over ATM encapsulation. To configure encapsulation on an ATM-over-ADSL or ATM-over-SHDSL logical interface, use PPPoE over AAL5 logical link control (LLC) encapsulation. LLC encapsulation allows a single ATM virtual connection to transport multiple protocols.

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.

Perform the task appropriate for the interface on which you are using PPPoE:

- Configuring PPPoE Encapsulation on an Ethernet Interface on page 199
- Configuring PPPoE Encapsulation on an ATM-over-ADSL or ATM-over-SHDSL Interface on page 200

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:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 50 on page 199.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To configure the PPPoE interface, see “Configuring PPPoE Interfaces (Required)” on page 201.
 - To enable authentication on the interface, see “Configuring CHAP on a PPPoE Interface (Optional)” on page 204.
 - To check the configuration, see “Verifying a PPPoE Configuration” on page 205.

Table 50: Configuring PPPoE Encapsulation on an Ethernet Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none">1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration.2. Next to Interfaces, click Configure or Edit.	From the [edit] hierarchy level, enter edit interfaces
Configure encapsulation on a logical Ethernet interface—for example, ge-0/0/1.0 . For information about interface names, see “Network Interface Naming” on page 47.	<ol style="list-style-type: none">1. In the Interface name box, click ge-0/0/1.2. In the Interface unit number box, click 0.3. From the Encapsulation list, select ppp-over-ether.4. Click OK.	Set PPP encapsulation on unit 0 of the Ethernet interface: set ge-0/0/1 unit 0 encapsulation ppp-over-ether

Configuring PPPoE Encapsulation on an ATM-over-ADSL or ATM-over-SHDSL Interface

To configure PPPoE encapsulation on an ATM-over-ADSL or ATM-over-SHDSL interface:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 51 on page 200.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To configure the PPPoE interface, see “Configuring PPPoE Interfaces (Required)” on page 201.
 - To enable authentication on the interface, see “Configuring CHAP on a PPPoE Interface (Optional)” on page 204.
 - To check the configuration, see “Verifying a PPPoE Configuration” on page 205.

Table 51: Configuring PPPoE Encapsulation on an ATM-over-ADSL or ATM-over-SHDSL Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Configure or Edit. 	From the [edit] hierarchy level, enter edit interfaces
Navigate to the ATM-over-ADSL or ATM-over-SHDSL interface—for example, at-2/0/0 —and set the ATM virtual path identifier (VPI) to 0.	<ol style="list-style-type: none"> 1. In the Interface name box, click at-2/0/0. 2. Next to ATM options, click Configure. 3. Next to Vpi, click Add new entry. 4. In the Vpi number box, type 0. 5. Click OK twice. 	Enter set at-2/0/0 atm-options vpi 0

Table 51: Configuring PPPoE Encapsulation on an ATM-over-ADSL or ATM-over-SHDSL Interface *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
Do one of the following:	To configure the ADSL operating mode on the physical ATM interface:	Enter
■ Configure the ADSL operating mode on the physical ATM interface—for example, autonegotiation.	1. Next to Dsl options, click Configure .	set at-2/0/0 dsl-options operating-mode auto
■ Configure the SHDSL options:	2. From the Operating mode list, select auto .	
■ Annex type—for example, Annex A.	3. Click OK .	
■ SHDSL line rate for SHDSL interface—for example, automatic selection of line rate.	To configure the SHDSL options:	Enter
■ Loopback option for testing the SHDSL connection integrity on the physical ATM interface—for example, local.	1. Next to Shdsl options, click Configure .	set at-2/0/0 shdsl-options annex annex-a
	2. From the Annex list, select Annex-a .	line-rate auto loopback local
	3. From the Line Rate list, select auto .	
	4. From the Loopback list, select local .	
	5. Click OK until you return to the Interfaces page.	
Configure Ethernet over ATM encapsulation on the physical ATM-over-ADSL or ATM-over-SHDSL interface.	From the Encapsulation list, select ethernet-over-atm .	Enter
		set at-2/0/0 encapsulation ethernet-over-atm
Create an ATM-over-ADSL or ATM-over-SHDSL logical interface, configure LLC encapsulation, and specify a VCI number.	1. Next to Unit, click Add new entry .	Enter
	2. In the Interface unit number box, type 0.	set at-2/0/0 unit 0 encapsulation
	3. From the Encapsulation list, select ppp-over-ether-over-atm-llc .	ppp-over-ether-over-atm-llc vci 0.120
	4. In the Multicast vci box, type 0.120 and click OK .	

Configuring PPPoE Interfaces (Required)

To create and configure a PPPoE interface over the underlying Ethernet and ATM interface:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 52 on page 202.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To enable authentication on the PPPoE interface, see “Configuring CHAP on a PPPoE Interface (Optional)” on page 204.
 - To check the configuration, see “Verifying a PPPoE Configuration” on page 205.

Table 52: Configuring a PPPoE Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Configure or Edit. 	<p>From the [edit] hierarchy level, enter</p> <p>edit interfaces</p>
Create a PPPoE interface with a logical interface unit 0.	<ol style="list-style-type: none"> 1. Next to Interface, click Add new entry. 2. In the Interface name box, type pp0 and click OK. 3. Under Interface name, click pp0. 4. Next to Unit, click Add new entry. 5. In the Interface unit number box, type 0. 	<p>Enter</p> <p>edit pp0 unit 0</p>
Configure an ISDN interface as the backup interface for the PPPoE interface—for example, dl0.0.	<ol style="list-style-type: none"> 1. Next to Backup options, click Configure. 2. In the Interface box, type dl0.0. 3. Click OK. 	<p>Enter</p> <p>set backup-options interface dl0.0</p>
Specify the logical Ethernet interface or the logical ATM interface as the underlying interface for the PPPoE session—for example, ge-0/0/1.0 or at-2/0/0.0.	<ol style="list-style-type: none"> 1. Next to Pppoe options, click Edit. 2. In the Underlying Interface box, type one of the following interface names: <ul style="list-style-type: none"> ■ For a logical Ethernet interface, type ge-0/0/1.0. ■ For a logical ATM interface type, at-2/0/0.0. 	<p>Enter one of the following commands:</p> <ul style="list-style-type: none"> ■ set pppoe-options underlying-interface ge-0/0/1.0. ■ set pppoe-options underlying-interface at-2/0/0.0.
Identify the access concentrator by a unique name—for example, ispl.com.	In the Access concentrator box type ispl.com.	<p>Enter</p> <p>set pppoe-options access-concentrator ispl.com</p>
Specify the number of seconds (from 1 through 4294967295) to wait before reconnecting after a PPPoE session is terminated—for example, 100. A 0 value (the default) specifies immediate reconnection.	In the Auto reconnect box, type 100.	<p>Enter</p> <p>set pppoe-options auto-reconnect 100</p>
Specify the number of seconds a session can be idle—for example, 100. A 0 value prevents the session from timing out.	In the Idle timeout box, type 100.	<p>Enter</p> <p>set pppoe-options idle-timeout 100.</p>
Specify the J-Series Services Router as the client for the PPPoE interface.	In the Client box, Yes .	<p>Enter</p> <p>set pppoe-options client.</p>

Table 52: Configuring a PPPoE Interface *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
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—for example, <code>video@ispl.com</code> .	<ol style="list-style-type: none"> In the Service name box, type <code>video@ispl.com</code>. Click OK. 	<p>Enter</p> <pre>set pppoe-options service-name video@ispl.com</pre>
Configure the maximum transmission unit (MTU) of the IPv4, IPv6, or Multiprotocol Label Switching (MPLS) protocol families—for example, <code>1492</code> .	<ol style="list-style-type: none"> Select one of the following protocol families: <ul style="list-style-type: none"> For the IPv4 family, in the Inet box, select Yes and click Configure. For the IPv6 family, in the Inet6 box, select Yes and click Configure. For the MPLS family, in the Mpls box, select Yes and click Configure. In the Mtu box, type <code>1492</code>. Click OK until you return to the Unit page. 	<p>Enter one of the following:</p> <pre>set family inet mtu 1492 set family inet6 mtu 1492 set family mpls mtu 1492</pre>
Configure the PPPoE logical interface address in one of the following ways: Do one of the following:	<p>Select one of the following IP address configurations:</p> <p>To assign the source and destination addresses:</p> <ol style="list-style-type: none"> Next to Address, click Add new entry. In the Inet Source box, type <code>192.168.1.1/32</code>, or in the Inet6 Source box, type <code>2004::1/128</code>. In the Inet Destination box, type <code>192.168.1.2</code>, or in the Inet6 Destination box, type <code>2004::2</code>. Click OK until you return to the Unit page. <p>To derive the IPv4 source address and assign the destination address:</p> <ol style="list-style-type: none"> Next to Inet, click Edit. Next to Unnumbered address, select the Yes check box and click Configure. In the Destination box, type <code>192.168.1.2</code>. In the Source box, type <code>lo0.0</code>. Click OK until you return to the Unit page. <p>To obtain an IP address from the remote end:</p> <ol style="list-style-type: none"> Next to Negotiate address, select the Yes check box. Click OK until you return to the Unit page. 	<p>Do one of the following:</p> <ul style="list-style-type: none"> To assign source and destination addresses enter one of the following sets of commands: <ul style="list-style-type: none"> For IPv4 addresses, <code>set family inet address 192.168.1.1/32 destination 192.168.1.2</code> For IPv6 addresses, <code>set family inet6 address 2004::1/128 destination 2004::2</code> To derive the IPv4 source address and assign the destination address, enter <code>set family inet unnumbered-address lo0.0 destination 192.168.1.2</code>. To obtain an IP address from the remote end, enter <code>set family inet negotiate-address</code>.
Disable the sending of keepalives on a logical interface.	<ol style="list-style-type: none"> From the Keepalive choices list, select no keepalives. Click OK to apply your entries to the configuration. 	<p>Enter</p> <pre>set no-keepalives</pre>

To clear a PPPoE session on the pp0.0 interface, enter the `clear pppoe sessions pp0.0` command. To clear all sessions on the PPPoE interface, enter the `clear pppoe sessions` command.

Configuring CHAP on a PPPoE Interface (Optional)

To configure CHAP on the PPPoE interface:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 53 on page 204.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying a PPPoE Configuration” on page 205.

Table 53: Configuring CHAP on a PPPoE Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Profile level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Access, click Configure or Edit. 	From the [edit] hierarchy level, enter set access profile A-ppp-client client client1 chap-secret my-secret
Define a CHAP access profile—for example, A-ppp-client —with a client named client 1 and the secret (password) my-secret .	<ol style="list-style-type: none"> 1. Next to Profile, click Add new entry. 2. In the Profile name box, type A-ppp-client. 3. Next to Client, click Add new entry. 4. In the Name box, type client1. 5. In the Chap secret box, type my-secret. 6. Click OK until you return to the main Configuration page. 	
Navigate to the pp0 unit 0 interface level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. On the main Configuration page next to Interfaces, click Configure or Edit. 2. In the Interface name box, click pp0. 3. In the Interface unit number box, click 0. 	From the [edit] hierarchy level, enter edit interfaces pp0 unit 0
Configure CHAP on the PPPoE interface, and specify a unique profile name containing a client list and access parameters—for example, A-ppp-client .	<ol style="list-style-type: none"> 1. Next to Ppp options, click Configure. 2. Next to Chap, click Configure. 3. In the Access profile box, type A-ppp-client. 	Enter set ppp-options chap access-profile A-ppp-client

Table 53: Configuring CHAP on a PPPoE Interface *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
Specify a unique hostname to be used in CHAP challenge and response packets—for example, A-ge-0/0/1.0 or A-at-2/0/0.0.	<p>In the Local name box, type one of the following:</p> <ul style="list-style-type: none">■ For an Ethernet interface, type A-ge-0/0/1.0.■ For an ATM interface, type A-at-2/0/0.0.	<p>Do one of the following:</p> <ul style="list-style-type: none">■ For the Ethernet interface, enter <code>set ppp-options chap local-name A-ge-0/0/1.0</code>.■ For the ATM interface, enter <code>set ppp-options chap local-name A-at-2/0/0.0</code>.
Set the passive option to handle incoming CHAP packets only.	<ol style="list-style-type: none">1. In the Passive box, click Yes.2. Click OK.	<p>Enter</p> <p><code>set ppp-options chap passive</code></p>

Verifying a PPPoE Configuration

- To verify PPPoE configuration perform the following tasks:
- Displaying a PPPoE Configuration for an Ethernet Interface on page 206
 - Displaying a PPPoE Configuration for an ATM-over-ADSL or ATM-over-SHDSL Interface on page 207
 - Verifying PPPoE Interfaces on page 208
 - Verifying PPPoE Sessions on page 209
 - Verifying the PPPoE Version on page 209
 - Verifying PPPoE Statistics on page 210

Displaying a PPPoE Configuration for an Ethernet Interface

Purpose Verify the PPPoE configuration for an Ethernet interface.

Action From the J-Web interface, select **Configuration > View and Edit > View Configuration Text**. Alternatively, from configuration mode in the CLI, enter the `show interfaces` command from the top level.

```
[edit]
user@host#show interfaces
ge-3/0/0 {
  unit 1 {
  }
}
pp0 {
  unit 1 {
    pppoe-options {
      underlying-interface ge-3/0/0.0;
      idle-timeout 123;
      access-concentrator myac;
      service-name myserv;
      auto-reconnect 10;
      client;
    }
    family inet {
      address 22.2.2.1/32 {
        destination 22.2.2.2;
      }
    }
    family inet6 {
      address 3004::1/128 {
        destination 3004::2;
      }
    }
  }
}
```

What It Means Verify that the output shows the intended configuration of PPPoE.

See Also For more information about the format of a configuration file, see “Viewing the Configuration Text” on page 9.

Displaying a PPPoE Configuration for an ATM-over-ADSL or ATM-over-SHDSL Interface

Purpose Verify the PPPoE configuration for an ATM-over-ADSL or ATM-over-SHDSL interface.

Action From the J-Web interface, select **Configuration > View and Edit > View Configuration Text**. Alternatively, from configuration mode in the CLI, enter the `show interfaces` command from the top level.

```
[edit]
user@host#show interfaces
at-6/0/0 {
    encapsulation ethernet-over-atm;
    atm-options {
        vpi 0;
    }
    dsl-options {
        operating-mode itu-dmt;
    }
    unit 0 {
        encapsulation ppp-over-ether-over-atm-llc;
        vci 35;
    }
}
pp0 {
    unit 0 {
        pppoe-options {
            underlying-interface at-6/0/0.0;
            idle-timeout 123;
            access-concentrator myac;
            service-name myserv;
            auto-reconnect 10;
            client;
        }
        family inet {
            address 11.1.1.1/32 {
                destination 11.1.1.2;
            }
        }
        family inet6 {
            address 2004::1/128 {
                destination 2004::2;
            }
        }
        family mpls;
    }
}
```

What It Means Verify that the output shows the intended configuration of PPPoE.

See Also For more information about the format of a configuration file, see “Viewing the Configuration Text” on page 9.

Verifying PPPoE Interfaces

Purpose Verify that the PPPoE router interfaces are configured properly.

Action From the CLI, enter the `show interfaces pp0` command.

```
user@host> show interfaces pp0
Physical interface: pp0, Enabled, Physical link is Up
  Interface index: 67, SNMP ifIndex: 317
  Type: PPPoE, Link-level type: PPPoE, MTU: 9192
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Link type      : Full-Duplex
  Link flags     : None
  Last flapped   : Never
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

Logical interface pp0.0 (Index 1) (SNMP ifIndex 330)
  Flags: Point-To-Point SNMP-Traps 16384 Encapsulation: PPPoE
  PPPoE:
    State: SessionUp, Session ID: 3304,
    Session AC name: isp1.com, AC MAC address: 00:90:1a:40:f6:4c,
    Service name: video@isp1.com, Configured AC name: isp1.com,
    Auto-reconnect timeout: 60 seconds
    Underlying interface: ge-5/0/0.0 (Index 71)
  Input packets : 23
  Output packets: 22
  Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
  Keepalive: Input: 16 (00:00:26 ago), Output: 0 (never)
  LCP state: Opened
  NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
  Not-configured
  CHAP state: Success
    Protocol inet, MTU: 1492
    Flags: Negotiate-Address
    Addresses, Flags: Kernel Is-Preferred Is-Primary
    Destination: 211.211.211.2, Local: 211.211.211.1
```

What It Means The output shows information about the physical and the logical interface. Verify the following information:

- The physical interface is enabled and the link is up.
- The PPPoE session is running on the correct logical interface.
- Under **State**, the state is active (**up**).
- Under **Underlying interface**, the physical interface on which the PPPoE session is running is correct:
 - For an Ethernet connection, the underlying interface is Fast Ethernet or Gigabit Ethernet—for example, `ge-5/0/0.0`.
 - For an ATM-over-ADSL or ATM-over-SHDSL connection, the underlying interface is ATM—for example, `at-2/0/0.0`.

See Also For a complete description of `show interfaces pp0` output, see the *JUNOS Interfaces Command Reference*.

Verifying PPPoE Sessions

Purpose Verify that a PPPoE session is running properly on the logical interface.

Action From the CLI, enter the `show pppoe interfaces` command.

```
user@host> show pppoe interfaces
pp0.0 Index 67
  State: Session up, Session ID: 31,
  Service name: video@isp1.com, Configured AC name: isp1.com,
  Session AC name: belur, AC MAC address: 00:90:1a:40:f6:4e,
  Auto-reconnect timeout: 1 seconds,
  Underlying interface: ge-0/0/1.0 Index 69
```

What It Means The output shows information about the PPPoE sessions. Verify the following information:

- The PPPoE session is running on the correct logical interface.
- Under **State**, the session is active (**up**).
- Under **Underlying interface**, the physical interface on which the PPPoE session is running is correct:
 - For an Ethernet connection, the underlying interface is Fast Ethernet or Gigabit Ethernet—for example, `ge-0/0/1.0`.
 - For an ATM-over-ADSL or ATM-over-SHDSL connection, the underlying interface is ATM—for example, `at-2/0/0.0`.

See Also For a complete description of `show pppoe interfaces` output, see the *JUNOS Interfaces Command Reference*.

Verifying the PPPoE Version

Purpose Verify the version information of the PPPoE protocol configured on the Services Router interfaces.

Action From the CLI, enter the `show pppoe version` command.

```
user@host> show pppoe version
Point-to-Point Protocol Over Ethernet, version 1. rfc2516
  PPPoE protocol           = Enabled
  Maximum Sessions         = 256
  PADI resend timeout      = 2 seconds
  PADR resend timeout      = 16 seconds
  Max resend timeout       = 64 seconds
  Max Configured AC timeout = 4 seconds
```

What It Means The output shows PPPoE protocol information. Verify the following information:

- The correct version of the PPPoE protocol is configured on the interface.
- Under **PPPoE protocol**, the PPPoE protocol is enabled.

See Also For a complete description of `show pppoe version` output, see the *JUNOS Interfaces Command Reference*.

Verifying PPPoE Statistics

Purpose Display statistics information about PPPoE interfaces.

Action From the CLI, enter the `show pppoe statistics` command.

```
user@host> show pppoe statistics
Active PPPoE sessions: 4
```

PacketType	Sent	Received
PADI	502	0
PADO	0	219
PADR	219	0
PADS	0	219
PADT	0	161
Service name error	0	0
AC system error	0	13
Generic error	0	0
Malformed packets	0	41
Unknown packets	0	0
Timeout		
PADI	42	
PADO	0	
PADR	0	

What It Means The output shows information about active sessions on PPPoE interfaces. Verify the following information:

- Total number of active PPPoE sessions running on the interface.
- Under **Packet Type**, the number of packets of each type sent and received during the PPPoE session.

See Also For a complete description of `show pppoe statistics` output, see the *JUNOS Interfaces Command Reference*.

Chapter 7

Configuring ISDN

ISDN connectivity is supported on the J-series Services Routers as a backup for a primary Internet connection. The J-series Services Routers can be configured to “fail over” to an ISDN interface when the primary connection experiences interruptions in Internet connectivity.

Use ISDN also at the central office to terminate calls that originate at branch office routers and for central office callback for security, accounting, or cost savings at the branch office.

You can use either J-Web Quick Configuration or a configuration editor to configure ISDN BRI interfaces. To configure ISDN PRI, you use either the J-Web configuration editor or CLI configuration editor.



NOTE: This chapter provides instructions for configuring basic ISDN BRI service and features such as dial backup, dial-in, or callback for both ISDN BRI and ISDN PRI. To configure basic ISDN PRI service, see “Configuring Channelized T1/E1/ISDN PRI Interfaces” on page 139.

This chapter contains the following topics:

- ISDN Terms on page 211
- ISDN Overview on page 214
- Before You Begin on page 215
- Configuring ISDN BRI Interfaces with Quick Configuration on page 216
- Configuring ISDN Interfaces and Features with a Configuration Editor on page 223
- Verifying the ISDN Configuration on page 245

ISDN Terms

Before configuring ISDN, become familiar with the terms defined in Table 54 on page 212.

Table 54: ISDN Terminology

Term	Definition
bandwidth on demand	ISDN cost-control feature defining the bandwidth threshold that must be reached on all links before a Services Router initiates additional ISDN data connections to provide more bandwidth.
Basic Rate Interface (BRI)	ISDN service intended for home and small enterprise applications. ISDN BRI consists of two 64-Kbps B-channels to carry voice or data and one 16-Kbps D-channel for control and signaling.
bearer channel (B-channel)	64-Kbps channel used for voice or data transfer on an ISDN interface.
callback	Alternative feature to dial-in that enables a J-series Services Router to call back the caller from the remote end of a backup ISDN connection. Instead of accepting a call from the remote end of the connection, the router rejects the call, waits a configured period of time, and calls a number configured on the router's dialer interface. See also <i>dial-in</i> .
caller ID	Telephone number of the caller on the remote end of a backup ISDN connection, used to dial in and also to identify the caller. Multiple caller IDs can be configured on an ISDN dialer interface. During dial-in, the router matches the incoming call's caller ID against the caller IDs configured on its dialer interfaces. Each dialer interface accepts calls from only callers whose caller IDs are configured on it.
delta-channel (D-channel)	Circuit-switched channel that carries signaling and control for B-channels. In ISDN Basic Rate Interface (BRI) applications, a D-channel can also support customer packet data traffic at speeds up to 9.6 Kbps.
demand circuit	Network segment whose cost varies with usage, according to a service level agreement with a service provider. Demand circuits limit traffic based on either bandwidth (bytes or packets transmitted) or access time. For example, ISDN interfaces can be configured for dial-on-demand routing backup. In OSPF, the demand circuit reduces the amount of OSPF traffic by removing all OSPF protocols when the routing domain is in a steady state.
dial backup	Feature that reestablishes network connectivity through one or more backup ISDN dialer interfaces after a primary interface fails. When the primary interface is reestablished, the ISDN interface is disconnected.
dialer filter	Stateless firewall filter that enables dial-on-demand routing backup when applied to a physical ISDN interface and its dialer interface configured as a passive static route. The passive static route has a lower priority than dynamic routes. If all dynamic routes to an address are lost from the routing table and the router receives a packet for that address, the dialer interface initiates an ISDN backup connection and sends the packet over it. See also <i>dial-on-demand routing backup; floating static route</i> .
dialer interface (dl)	Logical interface for configuring dialing properties and the control interface for a backup ISDN connection.

Table 54: ISDN Terminology (*continued*)

Term	Definition
dial-in	Feature that enables J-series Services Routers to receive calls from the remote end of a backup ISDN connection. The remote end of the ISDN call might be a service provider, a corporate central location, or a customer premises equipment (CPE) branch office. All incoming calls can be verified against caller IDs configured on the router's dialer interface. See also <i>callback</i> .
dial-on-demand routing (DDR) backup	<p>Feature that provides a J-series Services Router with full-time connectivity across an ISDN line.</p> <p>When routes on a primary serial T1, E1, T3, E3, Fast Ethernet, Gigabit Ethernet, or PPPoE interface are lost, an ISDN dialer interface establishes a backup connection. To save connection time costs, the Services Router drops the ISDN connection after a configured period of inactivity. Services Routers with ISDN interfaces support two types of dial-on-demand routing backup: on-demand routing with a dialer filter and dialer watch. See also <i>dialer filter</i>; <i>dialer watch</i>.</p>
dialer profile	Set of characteristics configured for the ISDN dialer interface. Dialer profiles allow the configuration of physical interfaces to be separated from the logical configuration of dialer interfaces required for ISDN connectivity. This feature also allows physical and logical interfaces to be bound together dynamically on a per-connection basis.
dialer watch	Dial-on-demand routing (DDR) backup feature that provides reliable connectivity without relying on a dialer filter to activate the ISDN interface. The ISDN dialer interface monitors the existence of each route on a watch list. If all routes on the watch list are lost from the routing table, dialer watch initiates the ISDN interface for failover connectivity. See also <i>dial-on-demand routing backup</i> .
floating static route	Route with an administrative distance greater than the administrative distance of the dynamically learned versions of the same route. The static route is used only when the dynamic routes are no longer available. When a floating static route is configured on an interface with a dialer filter, the interface can be used for backup.
Integrated Services Digital Network (ISDN)	Digital communication service provided by telecommunication service providers. It is an all-digital dialup (on-demand) service that carries voice, data, and video transmissions over telephone lines.
Primary Rate Interface (PRI)	ISDN service intended for higher-bandwidth applications than ISDN BRI. ISDN PRI consists of a single D-channel for control and signaling, plus a number of 64-Kbps B-channels—either 23 B-channels on a T1 line or 30 B-channels on an E1 line—to carry network traffic.
service profile identifier (SPID)	Number that specifies the services available to you on the service provider switch and defines the feature set ordered when the ISDN service is provisioned.
terminal endpoint identifier (TEI)	Number that identifies a terminal endpoint, an ISDN-capable device attached to an ISDN network through an ISDN interface on the Services Router. The TEI is a number between 0 and 127. The numbers 0–63 are used for static TEI assignment, 64–126 are used for dynamic assignment, and 127 is used for group assignment.

ISDN Overview

Integrated Services Digital Network (ISDN) is a set of standards for digital transmission over different media created by the Consultative Committee for International Telegraph and Telephone (CCITT) and International Telecommunication Union (ITU). As a dial-on-demand service, it has fast call setup and low latency as well as 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: at least one physical interface and a logical interface called the dialer interface.

ISDN Interfaces

Table 55 on page 214 lists the J-series Services Router interfaces available for ISDN connectivity.

Table 55: ISDN Ports and PIMs

J2300 Model	All Other J-series Models
One of the following built-in ISDN BRI interfaces:	For ISDN BRI, up to six of the following field-replaceable units (FRUs):
<ul style="list-style-type: none"> One S/T port supporting ITU-T I.430, ETSI TS 101080, and GR-1089-Core Type III One U port supporting ANSI T.601 and GR-1089-Core 	<ul style="list-style-type: none"> 4-port S/T PIM supporting ITU-T I.430, ETSI TS 101080, and GR-1089-Core Type III 4-port U PIM supporting ANSI T.601 and GR-1089-Core
	For ISDN PRI, up to six Dual-Port Channelized T1/E1/ISDN PRI PIMs, supporting ITU-T Q.920, Q.921: LAPD, Q.930, and Q.931

ISDN BRI Interface Types

A J-series Services Router with one or more ISDN BRI ports has the following types of ISDN interfaces:

- Physical ISDN BRI interface—*br-pim/0/port*
- Physical B-channel interface—*bc-pim/0/port*
- Physical D-channel interface—*dc-pim/0/port*
- Logical dialer interface—*dln*

For information about interface names, see “Network Interface Naming” on page 47.

To configure ISDN BRI service on a Services Router, you configure the physical ISDN BRI interface and the logical dialer interface.

Each ISDN BRI port has two B-channels for transport, identified as *bc-pim/0/port:1* and *bc-pim/0/port:2*, and one D-channel for control, identified as *dc-pim/0/port*. On ISDN BRI interfaces, the B-channels and D-channel have no configurable settings, but you can monitor them for interface status and statistics.

ISDN PRI Interface Types

On a J-series Services Router with one or more Dual-Port Channelized T1/E1/ISDN PRI PIMs, you can configure each port on the PIM for either T1, E1, or ISDN PRI service, or for a combination of ISDN PRI and either T1 or E1 service. For ISDN PRI service, you configure the following types of ISDN interfaces as channels on the channelized T1 or E1 interface:

- Physical B-channel interface—*bc-pim/0/port:channel*
 - On a channelized T1 interface, up to 23 time slots can be configured as ISDN PRI B-channels.
 - On a channelized E1 interface, up to 30 time slots can be configured as ISDN PRI B-channels.
- Physical D-channel interface—*dc-pim/0/port:channel*
 - On a channelized T1 interface, you configure time slot 24 as the D-channel.
 - On a channelized E1 interface, you configure time slot 16 as the D-channel.
- Logical dialer interface—*dlm*

For information about interface names, see “Network Interface Naming” on page 47.

For more information about channelized T1/E1/ISDN PRI interfaces, see “Configuring Channelized T1/E1/ISDN PRI Interfaces” on page 139.

Dialer Interface

The dialer (*dlm*) interface is a logical interface on which you configure dialing properties for ISDN connections. The interface can be configured in two modes:

- Multilink mode using Multilink PPP encapsulation
- Normal mode using PPP or Cisco High-Level Data Link Control (HDLC) encapsulation

The dialer interface can perform backup, dialer filter, and dialer watch functions, but these operations are mutually exclusive. You can configure a single dialer interface to operate in only one of the following ways:

- As a backup interface—for one primary interface
- As a dialer filter
- As a dialer watch interface

Before You Begin

Before you configure ISDN interfaces, you need to perform the following tasks:

- Install Services Router hardware. For more information, see the Getting Started Guide for your router.
- Establish basic connectivity. For more information, see the Getting Started Guide for your router.
- Order an ISDN line from your telecommunications service provider. Contact your service provider for more information.
- If you do not already have a basic understanding of physical and logical interfaces and Juniper Networks interface conventions, read “Interfaces Overview” on page 41.

Although it is not a requirement, you might also want to plan how you are going to use the ISDN interfaces on your network before you begin configuring them. (To display a list of installed ISDN BRI interfaces, select **Configuration > Quick Configuration > Interfaces**.)

Configuring ISDN BRI Interfaces with Quick Configuration

You can use the ISDN Interfaces Quick Configuration pages to configure ISDN BRI interfaces on a router. The Quick Configuration pages allow you to configure ISDN BRI connectivity on a router to back up a primary Internet connection.



NOTE: To configure an ISDN *PRI* interface, you must use the J-Web or CLI configuration editor.

You configure the physical ISDN BRI interface first and then the backup method on the logical dialer interface.

This section contains the following topics:

- Configuring ISDN BRI Physical Interfaces with Quick Configuration on page 216
- Configuring ISDN BRI Dialer Interfaces with Quick Configuration on page 219

Configuring ISDN BRI Physical Interfaces with Quick Configuration

To configure ISDN BRI physical interfaces with Quick Configuration:

1. In the J-Web interface, select **Configuration > Quick Configuration > Interfaces**.
A list of network interfaces installed on the router is displayed.
2. Click the *br-pim/0/port* interface name for the ISDN BRI port you want to configure.

The ISDN BRI Physical Interface Quick Configuration page is displayed as shown in Figure 35 on page 217.

Figure 35: ISDN BRI Physical Interface Quick Configuration Page

Juniper
NETWORKS

ROUTER - J4300

Monitor Configuration Diagnose Manage Events Logged in as: regress Help About Logout

Configuration > Quick Configuration > Interfaces

Quick Configuration

Interfaces Physical Interface: 'br-5/0/1'

Physical Interface Description

Dialer Pools

No dialer pools are configured.

ISDN Options

Calling Number ?

ISDN Switch Type ?

Service Profile Identifier ?

Service Profile Identifier 2 ?

Static TEI Value ?

TEI Option ?

Timer T310 Value ?

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. Trademark Notice. Privacy. Juniper Your Net.

3. Enter information into the ISDN Quick Configuration pages, as described in Table 56 on page 217.
4. From the ISDN Physical Interfaces Quick Configuration page:
 - To apply the configuration and stay on the ISDN Physical Interfaces Quick Configuration page, click **Apply**.
 - To apply the configuration and return to the Interfaces Quick Configuration page, click **OK**.
 - To cancel your entries and return to the Interfaces Quick Configuration page, click **Cancel**.
5. Go on to “Configuring ISDN BRI Dialer Interfaces with Quick Configuration” on page 219.

Table 56: ISDN BRI Quick Configuration Page Summary

Field	Function	Your Action
Configuring ISDN Interfaces		

Table 56: ISDN BRI Quick Configuration Page Summary *(continued)*

Field	Function	Your Action
Physical Interface Description	(Optional) Adds supplemental information about the ISDN physical interface on the router.	Type a text description of the physical ISDN BRI interface in the box to clearly identify it in monitoring displays.
Clocking	<p>Enables internal or external clocking sources for the interface on the router.</p> <ul style="list-style-type: none"> ■ internal—Services Router's own system clock (the default) ■ external—Clock received from the T1 interface 	Select internal or external from the list.
Dialer Pool Options		
Dialer Pools	Displays the list of configured ISDN dialer pools on the router.	<ul style="list-style-type: none"> ■ To add a dialer pool to the interface, click Add. ■ To edit a dialer pool, select the name from the list. You can change the priority, but not the name. ■ To delete a dialer pool, select the check box and click Delete.
Dialer Pool Name (required)	Specifies the group of physical interfaces to be used by the dialer interface.	Type the dialer pool name—for example, isdn-dialer-pool .
Priority	Specifies the priority of this interface within the dialer pool. Interfaces with a higher priority are the first to interact with the dialer interface.	<ol style="list-style-type: none"> 1. Type a priority value from 0 (lowest) to 255 (highest). The default is 0. 2. Click OK to return to the Quick Configuration page.
ISDN Options		
Calling Number	Configures the dialing number used to connect with the service provider.	Type the outgoing calling number for the service provider.
ISDN Switch Type	Specifies the type of ISDN switch used by the service provider.	<p>Select one of the following switch types:</p> <ul style="list-style-type: none"> ■ att5e—AT&T 5ESS ■ etsi—NET3 for the UK and Europe ■ ni1—National ISDN-1 ■ ntdms-100—Northern Telecom DMS-100 ■ ntt—NTT Group switch for Japan
Service Profile Identifier Service Profile Identifier 2	Configures the service profile identifier (SPID) provided by your ISDN service.	Type the SPID in the box. If you have a NTDMS-100 or NI1 switch, an additional SPID field is provided.

Table 56: ISDN BRI Quick Configuration Page Summary (continued)

Field	Function	Your Action
Static TEI Value	<p>Configures the static terminal endpoint identifier (TEI) value from your service provider.</p> <p>The TEI number identifies a terminal endpoint, an ISDN-capable device attached to an ISDN network through an ISDN interface on the Services Router. The TEI is a number between 0 and 127. The numbers 0–63 are used for static TEI assignment, 64–126 are used for dynamic assignment, and 127 is used for group assignment.</p>	<p>Type a value between 0 and 63. If this value is not supplied, the router dynamically acquires a TEI.</p> <p>If you configured more than one SPID, the TEI must be acquired dynamically.</p>
TEI Option	Configures when the TEI negotiates with the ISDN provider.	<ul style="list-style-type: none"> ■ Select first-call to activate the connection when the call setup is sent to the ISDN provider. ■ Select power-up (the default) to activate the connection when the router is powered on.
Timer T310 Value	Sets the Q.931 timer value in seconds.	Type a value between 1 and 65536. The default value is 10 seconds.

Configuring ISDN BRI Dialer Interfaces with Quick Configuration

When ISDN BRI interfaces are installed on the Services Router, links to ISDN Quick Configuration pages for dialer options are displayed on the Interfaces Quick Configuration page as shown in Figure 36 on page 220.

You can use these Quick Configuration pages to configure an ISDN BRI dialer interface for either dial backup or dialer watch. For dial backup you specify the serial interface to back up. For dialer watch you specify a watch list of one or more routes to monitor.

Figure 36: ISDN BRI Dialer Options Quick Configuration Page

The screenshot shows the Juniper J-Web interface for a ROUTER - J4300. The top navigation bar includes links for Monitor, Configuration, Diagnose, Manage, and Events. The left sidebar shows a tree view with options like Quick Configuration, View and Edit, History, and Rescue. The main content area is titled 'Quick Configuration' and 'Interfaces'. It contains a table of network interfaces with columns for Interface Name, Link State, Configured, and Description.

Interface Name	Link State	Configured	Description
fe-0/0/0	Up	Yes	Fast Ethernet Interface 'fe-0/0/0'
fe-0/0/0.0	Up	Yes	Logical Unit 0 on Fast Ethernet Interface 'fe-0/0/0'
fe-0/0/1	Up	Yes	Fast Ethernet Interface 'fe-0/0/1'
fe-0/0/1.0	Up	Yes	Logical Unit 0 on Fast Ethernet Interface 'fe-0/0/1'
br-5/0/0	Down	Yes	ISDN BRI Interface 'br-5/0/0'
br-5/0/1	Up	No	ISDN BRI Interface 'br-5/0/1'
br-5/0/2	Up	No	ISDN BRI Interface 'br-5/0/2'
br-5/0/3	Up	No	ISDN BRI Interface 'br-5/0/3'
e3-6/0/0	Up	No	E3 Interface 'e3-6/0/0'
lo0	Up	Yes	Loopback Interface 'lo0'
lo0.0	Up	Yes	Logical Unit 0 on Loopback Interface 'lo0'
pp0	Up	No	Point-to-Point Protocol over Ethernet Interface 'pp0'

Below the table, there is a section titled 'ISDN Dialer Options' with the text: 'Configure ISDN Dialer features Dial Backup, Dial Watch, and Dial on Demand.' At the bottom of this section are three buttons: OK, Cancel, and Apply.

To configure ISDN BRI dialer interfaces with Quick Configuration:

1. In the J-Web interface, select **Configuration > Quick Configuration > Interfaces**.

A list of network interfaces installed on the Services Router is displayed.

2. Click **ISDN Dialer Options** under the interfaces list.
3. Select a backup method to configure on the dialer interface:
 - Click **Dial Backup** to allow one or more dialer interfaces to back up the primary interface. The backup interfaces are activated only when the primary interface fails.
 - Click **Dialer Watch** to monitor a specified route and initiate dialing of the backup link if that route is not present.
4. Do one of the following:
 - To edit an existing dialer interface, click the dialer interface name. For example, click **dl0** to edit the dialer physical interface, and then click **dl0.0** to edit the dialer logical interface.
 - To add a dialer interface, click **Add**. In the Interface Name box, type a name for the logical interface—for example, **dl1**—then click **Add** under Logical Interfaces.

Figure 37 on page 221 shows the ISDN Quick Configuration page for dialer logical interfaces.

Figure 37: ISDN BRI Dialer Interface Quick Configuration Page

The screenshot shows the Juniper Router - J6300 web interface. The top navigation bar includes links for Monitor, Configuration, Diagnose, Manage, Events, Alarms, and a user login status. The left sidebar shows a menu with Quick Configuration, View and Edit, History, and Rescue. The main content area is titled 'Quick Configuration' and 'Interfaces'. It displays the 'Dialer Logical Interface: 'dl0.0'' configuration page. The page is divided into several sections: 'Logical Interface Description' with a text input field; 'IPv4 Addresses and Prefixes' with a table for adding and deleting addresses, including 'Add' and 'Delete' buttons; 'Dialer Options' with fields for 'Activation Delay', 'Deactivation Delay', 'Dial String' (with a list of strings and 'Add/Delete' buttons), and 'Pool' (with a dropdown menu); and 'Backup Interface' with a dropdown menu for 'Interface to Backup'. At the bottom, there are 'OK', 'Cancel', and 'Apply' buttons. The footer contains copyright information and the Juniper logo.

5. Enter information into the ISDN Quick Configuration page for dialer logical interfaces, as described in Table 57 on page 222.
6. Click one of the following buttons on the ISDN Quick Configuration page:
 - To apply the configuration and stay on the current Quick Configuration page, click **Apply**.
 - To apply the configuration and return to the previous Quick Configuration page, click **OK**.
 - To cancel your entries and return to the previous Quick Configuration page, click **Cancel**.
7. To verify that the ISDN interface is configured correctly, see “Verifying the ISDN Configuration” on page 245.

Table 57: ISDN BRI Dialer Interface Quick Configuration Page Summary

Field	Function	Your Action
Configuring Dialer Interfaces		
Logical Interface Description	Describes the logical interface.	Type a text description of the interface in the box.
IPv4 Addresses and Prefixes	Displays the IPv4 addresses for the interfaces to which the dialer interface is assigned. 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 can be routed through any of the dialer interfaces with the IP subnet address, instead of being routed through the dialer interface to which the ISDN call is mapped.	Type an IP address and a prefix in the boxes. Click Add . To delete an IP address, highlight it in the list, and click Delete .
Dialer Options		
Activation Delay	Displays the time to wait before activating the backup interface once the primary interface is down.	Type a value, in seconds—for example, 30. The default value is 0 seconds with a maximum value of 60 seconds.
Deactivation Delay	Displays the time to wait before deactivating the backup interface once the primary interface is up.	Type a value, in seconds—for example, 30. The default value is 0 seconds with a maximum value of 60 seconds.
Dial String (required)	Displays the dialing number from your ISDN service provider.	Type the dialing number and click Add . To delete a dial string, highlight it and click Delete .
Pool (required)	Displays a list of dialer pools configured on <i>br-pim/0/port</i> interfaces.	Select a dialer pool from the list.
Multilink Dialer Options		
Load Interval	Defines the interval used to calculate the average load on the dialer interface for bandwidth on demand.	Type a value, in seconds—for example, 30. The default value is 60 seconds with a range of 20–80. The value must be a multiple of 10.
Load Threshold	Defines the threshold at which an additional ISDN interface is activated for bandwidth-on-demand. You specify the threshold as a percentage of the cumulative load of all UP links.	Type a percentage—for example, 80. The default value is 100 with a range of 0–100.
Backup Interface (for dial backup only)		
Interface to Backup	Displays a list of interfaces for ISDN backup.	Select an interface from the list for ISDN backup.
Dialer Watch List (for dialer watch only)		

Table 57: ISDN BRI Dialer Interface Quick Configuration Page Summary *(continued)*

Field	Function	Your Action
IPv4 Addresses and Prefixes	Displays the IPv4 addresses in the list of routes to be monitored by the dialer interface.	Type an IP address and a prefix in the boxes. Click Add . To delete an IP address, highlight it in the list, and click Delete .

Configuring ISDN Interfaces and Features with a Configuration Editor

To configure ISDN interfaces on a J-series Services Router, you first configure the basic ISDN interface—either “Adding an ISDN BRI Interface (Required)” on page 224 or “Configuring Channelized T1/E1/ISDN PRI Interfaces for ISDN PRI Operation” on page 147. Second, configure the dialer interface by performing “Configuring Dialer Interfaces (Required)” on page 226.

To configure ISDN interfaces to back up primary Services Router interfaces, you then configure a backup method—either “Configuring Dial Backup” on page 228, “Configuring Dialer Filters for Dial-on-Demand Routing Backup” on page 229, or “Configuring Dialer Watch” on page 231.

To configure ISDN interfaces for dial-in or callback, configure the basic ISDN BRI or PRI interface and then perform “Configuring Dial-In and Callback (Optional)” on page 238.

Perform other tasks as needed on your network.

This section contains the following topics:

- Adding an ISDN BRI Interface (Required) on page 224
- Configuring Dialer Interfaces (Required) on page 226
- Configuring Dial Backup on page 228
- Configuring Dialer Filters for Dial-on-Demand Routing Backup on page 229
- Configuring Dialer Watch on page 231
- Configuring Dial-on-Demand Routing Backup with OSPF Support (Optional) on page 233
- Configuring Bandwidth on Demand (Optional) on page 234
- Configuring Dial-In and Callback (Optional) on page 238
- Disabling Dialing Out Through Dialer Interfaces on page 243
- Disabling ISDN Signaling on page 244

Adding an ISDN BRI Interface (Required)

To enable ISDN BRI interfaces installed on your Services Router to work properly, you must configure the interface properties.

To configure an ISDN BRI network interface for the Services Router:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 58 on page 224.
3. Go on to “Configuring Dialer Interfaces (Required)” on page 226.

Table 58: Adding an ISDN BRI Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Configure or Edit. 	From the [edit] hierarchy level, enter edit interfaces br-1/0/3
Create the new interface—for example, br-1/0/3.	<ol style="list-style-type: none"> 1. Next to Interface, click Add new entry. 2. In the Interface name box, type the name of the new interface, br-1/0/3. 3. Click OK. 	
Configure dialer options. <ul style="list-style-type: none"> ■ Name the dialer pool—for example, isdn-dialer-group. ■ Set the dialer pool priority—for example, 255. Dialer pool priority has a range from 1 to 255, with 1 designating lowest-priority interfaces and 255 designating the highest-priority interfaces.	<ol style="list-style-type: none"> 1. In the Encapsulation column, next to the new interface, click Edit. 2. Next to Dialer options, select Yes, and then click Configure. 3. Next to Pool, click Add new entry. 4. In the Pool identifier box, type isdn-dialer-group. 5. In the Priority box, type 255. 6. Click OK. 	From the [edit interfaces br-1/0/3] hierarchy, enter set dialer-options pool isdn-dialer-group priority 255

Table 58: Adding an ISDN BRI Interface (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Configure ISDN BRI properties.</p> <ul style="list-style-type: none"> ■ Calling number sent to the ISDN switch during the call setup, which represents the caller's number—for example, 18005555555. ■ Service provider ID (SPID)—for example, 00108005555555. ■ Static TEI between 0 and 63 from your service provider—for example, 23. If the field is left blank, the Services Router dynamically acquires a TEI. Also, if you have configured a second SPID, you cannot set a static TEI value. If you have a NTDMS-100 or NI1 switch, an additional box for a service provider ID is provided. If you are using a service provider that requires SPIDs, you cannot place calls until the interface sends a valid, assigned SPID to the service provider when accessing the ISDN connection. ■ Incoming called number—for example, 18883333456. Configure incoming call properties if you have remote locations dialing into the router through the ISDN interface. 	<ol style="list-style-type: none"> 1. Next to Isdn options, click Configure. 2. In the Calling number box, type 18005555555. 3. In the Spid1 box, type 00108005555555. 4. In the Static tei val box, type 23. 5. Next to Incoming called number, click Add new entry. 6. In the Called number box, type 18883333456. 7. Click OK. 	<ol style="list-style-type: none"> 1. To set the ISDN options, enter <pre>set isdn-options calling-number 18005555555</pre> 2. Enter <pre>set isdn-options spid1 00108005555555</pre> 3. Enter <pre>set isdn-options static-tei-val 23</pre> 4. set isdn-options incoming-called-number 18883333456
<p>Select the type of ISDN switch—for example, ATT5E. The following switches are compatible with Services Routers:</p> <ul style="list-style-type: none"> ■ ATT5E—AT&T 5ESS ■ ETSI—NET3 for the UK and Europe ■ NI1—National ISDN-1 ■ NTDMS-100—Northern Telecom DMS-100 ■ NTT—NTT Group switch for Japan 	<p>From the Switch type list, select att5e.</p>	<p>To select the switch type, enter</p> <pre>set isdn-options switch-type att5e</pre>
<p>Configure the Q.931 timer. Q.931 is a Layer 3 protocol for the setup and termination of connections. The default value for the timer is 10 seconds, but can be configured between 1 and 65536 seconds—for example, 15.</p>	<p>In the T310 box, type 15.</p>	<pre>set isdn-options t310 15</pre>

Table 58: Adding an ISDN BRI Interface *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure when the TEI negotiates with the ISDN provider.	1. From the Tei option list, select power-up .	To initiate activation at power-up, enter
■ first-call —Activation does not occur until a call is sent.	2. Click OK to return to the Interfaces page.	set isdn-options tei-option power-up
■ power-up —Activation occurs when the Services Router is powered on. This is the default value.		

Configuring Dialer Interfaces (Required)

The dialer interface (dl) is a logical interface configured to establish ISDN connectivity. You can configure multiple dialer interfaces for different functions on the Services Router.

After configuring the dialer interface, you must configure a backup method—either dial backup, a dialer filter, or dialer watch.

To configure a logical dialer interface for the Services Router:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 59 on page 226.
3. To configure a backup method, go on to one of the following tasks:
 - “Configuring Dial Backup” on page 228.
 - “Configuring Dialer Filters for Dial-on-Demand Routing Backup” on page 229.
 - “Configuring Dialer Watch” on page 231.

Table 59: Adding a Dialer Interface to a Services Router

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Configure or Edit. 	From the [edit] hierarchy level, enter edit interfaces

Table 59: Adding a Dialer Interface to a Services Router (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Create the new interface—for example, d10.</p> <p>Adding a description can differentiate between different dialer interfaces—for example, T1-backup.</p>	<ol style="list-style-type: none"> Next to Interface, click Add new entry. In the Interface name box, type d10. In the Description box, type T1-backup. Click OK. 	<p>Create and name the interface:</p> <ol style="list-style-type: none"> <code>edit d10</code> <code>set description T1-backup</code>
<p>Configure encapsulation options—for example, Cisco HDLC.</p> <ul style="list-style-type: none"> ■ Cisco HDLC—For normal mode (when the router is using only one B-channel). Cisco-compatible High-Level Data Link Control is a group of protocols for transmitting data between network points. ■ PPP—For normal mode (when the router is using only one ISDN B-channel per call). Point-to-Point Protocol is for communication between two computers using a serial interface. ■ Multilink PPP—For multilink mode, when the router is using multiple B-channels per call. Multilink Point-to-Point Protocol (MLPPP) is a protocol for aggregating multiple constituent links into one larger PPP bundle. You can bundle up to eight B-channels. 	<ol style="list-style-type: none"> In the Encapsulation column, next to the new interface, click Edit. From the Encapsulation list, select cisco-hdlc. 	<p>Enter</p> <p><code>set encapsulation cisco-hdlc</code></p>
<p>Enter a hold-time value in milliseconds—for example, 60. The hold-time value is used to damp interface transitions. When an interface goes from up to down, it is not advertised as down to the rest of the system until it remains unavailable for the hold-time period. Similarly, an interface is not advertised as up until it remains operational for the hold-time period. The hold time is three times the interval at which keepalive messages are sent.</p>	<ol style="list-style-type: none"> In the Hold time section, type 60 in the Down box. In the Up box, type 60. 	<ol style="list-style-type: none"> Enter <code>set hold-time down 60</code> Enter <code>set hold-time up 60</code>
<p>Create the logical unit—for example, 0.</p> <p>NOTE: You can set the logical unit to 0 only, unless you are configuring the dialer interface for Multilink PPP encapsulation.</p>	<ol style="list-style-type: none"> Next to Unit, click Add new entry. In the Interface unit number box, type 0. Next to Dialer options, select Yes, and then click Configure. 	<p>Enter</p> <p><code>set unit 0</code></p>

Table 59: Adding a Dialer Interface to a Services Router (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure dialer options. <ul style="list-style-type: none"> ■ Activation delay—Time to wait before activating the backup interface once the primary interface is down—for example, 30. Default value is 0 seconds with a maximum value of 60 seconds. Use only for dialer backup and dialer watch. ■ Deactivation delay—Time to wait before deactivating the backup interface once the primary interface is up—for example, 30. Default value is 0 seconds with a maximum value of 60 seconds. Use only for dialer backup and dialer watch. ■ Idle timeout—Time a connection is idle before disconnecting—for example, 30. Default value is 120 seconds with a range from 0 to 4294967295. This option is used only to configure a dialer filter. ■ Initial route check—Time to wait before checking if the primary interface is up—for example, 30. Default value is 120 seconds with a range of 1 to 300 seconds. This option is used only to configure dialer watch. ■ Pool—Name of a group of ISDN interfaces configured to use the dialer interface—for example, <code>isdn-dialer-group</code>. ■ Redial delay—Number of seconds to wait before redialing a failed outgoing ISDN call. Default value is 3 seconds with a range from 2 to 255. 	1. In the Activation delay box, type 60. 2. In the Deactivation delay box, type 30. 3. In the Pool box, type <code>isdn-dialer-group</code> . 4. In the Redial delay box, type 5.	1. Enter <code>edit unit 0 dialer-options</code> 2. Enter <code>set activation-delay 60</code> 3. Enter <code>set deactivation-delay 30</code> 4. Enter <code>set pool isdn-dialer-group</code> 5. Enter <code>set redial-delay 5</code>
Configure the remote destination to call—for example, 5551212.	1. Next to Dial string, click Add new entry . 2. In the Dial string box, type 5551212. 3. Click OK .	Enter <code>set dial-string 5551212</code>
Configure source and destination IP addresses for the dialer interface—for example, 172.20.10.2 and 172.20.10.1. (The destination IP address is optional.) <p>NOTE: If you configure multiple dialer interfaces, ensure that the same IP subnet address is not configured on different dialer interfaces. Configuring the same IP subnet address on multiple dialer interfaces can result in inconsistency in the route and packet loss. The router might route packets through another dialer interface with the IP subnet address instead of through the dialer interface to which the ISDN modem call is mapped.</p>	1. Select Inet under Family, and click Edit . 2. Next to Address, click Add new entry . 3. In the Source box, type 172.20.10.2. 4. In the Destination box, type 172.20.10.1. 5. Click OK .	1. From the [edit] hierarchy level, enter <code>edit interfaces d10 unit 0</code> 2. Enter <code>set family inet address 172.20.10.2 destination 172.20.10.1</code>

Configuring Dial Backup

Dial backup allows one or more dialer interfaces to be configured as the backup link for a primary interface. The backup dialer interfaces are activated only when the

primary interface fails. ISDN backup connectivity is supported on all interfaces except `ls-0/0/0`.

To configure a primary interface for backup connectivity:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 60 on page 229.
3. If you are finished configuring the router, commit the configuration.
4. Go on to any of the following optional tasks:
 - “Configuring Dial-on-Demand Routing Backup with OSPF Support (Optional)” on page 233.
 - “Configuring Bandwidth on Demand (Optional)” on page 234.
 - Configuring Dial-In and Callback (Optional) on page 238
5. To verify that the network interface is configured correctly, see “Verifying the ISDN Configuration” on page 245.

Table 60: Configuring an Interface for ISDN Backup

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	From the [edit] hierarchy level, enter edit interfaces ge-0/0/0 unit 0
Select the physical interface for backup ISDN connectivity.	<ol style="list-style-type: none"> 1. In the Interface name column, click the physical interface name. 1. Under Unit, in the Nested Configuration column, click Edit. 	
Configure the backup dialer interface—for instance, <code>dl0.0</code> .	<ol style="list-style-type: none"> 1. Next to Backup options, click Configure. 2. In the Interface box, type <code>dl0.0</code>. 3. Click OK until you return to the Interfaces page. 	Enter set backup-options interface dl0.0

Configuring Dialer Filters for Dial-on-Demand Routing Backup

This dial-on-demand routing backup method allows an ISDN line to be activated only when network traffic configured as an “interesting packet” arrives on the network. Once the network traffic is sent, an inactivity timer is triggered and the connection is closed after the timer expires.

You define an interesting packet using the dialer filter feature of the Services Router. There are two steps to configuring dial-on-demand routing backup using a dialer filter:

- Configuring the Dialer Filter on page 230
- Applying the Dial-on-Demand Dialer Filter to the Dialer Interface on page 231

Configuring the Dialer Filter

To configure the dialer filter:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 61 on page 230.
3. Go on to “Applying the Dial-on-Demand Dialer Filter to the Dialer Interface” on page 231.

Table 61: Configuring a Dialer Filter for Interesting Packets

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Firewall level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Firewall, click Edit. 	<p>From the [edit] hierarchy level, enter</p> <p>edit firewall</p>
Configure the dialer filter name—for example, int-packet.	<ol style="list-style-type: none"> 1. Next to Inet, click Configure or Edit. 2. Next to Dialer filter, click Add new entry. 3. In the Filter name box, type int-packet. 	<ol style="list-style-type: none"> 1. Enter <p>edit family inet</p> 2. Then enter <p>edit dialer-filter int-packet</p>
Configure the dialer filter rule name—for example, term1.	<ol style="list-style-type: none"> 1. Next to Term, click Add new entry. 2. In the Rule name box, type term1. 	Enter
Configure term behavior. For example, you might want to configure your interesting packet as an ICMP packet.	<ol style="list-style-type: none"> 3. Next to From, click Configure. 4. From the Protocol choice list, select Protocol. 5. Next to Protocol, click Add new entry. 	set term term1 from protocol icmp
To configure the term completely, include both from and then statements.	<ol style="list-style-type: none"> 6. From the Value keyword list, select icmp. 7. Click OK twice to return to the Term page. 	
Configure the then part of the dialer filter.	<ol style="list-style-type: none"> 1. Next to Then, click Configure. 2. From the Designation list, select Note. 3. Click OK. 	<p>Enter</p> <p>set term1 then note</p>

Applying the Dial-on-Demand Dialer Filter to the Dialer Interface

To complete dial-on-demand routing with dialer filter configuration:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 62 on page 231.
3. When you are finished configuring the router, commit the configuration.
4. Go on to any of the following optional tasks:
 - “Configuring Dial-on-Demand Routing Backup with OSPF Support (Optional)” on page 233.
 - “Configuring Bandwidth on Demand (Optional)” on page 234.
 - Configuring Dial-In and Callback (Optional) on page 238
5. To verify that the network interface is configured correctly, see “Verifying the ISDN Configuration” on page 245.

Table 62: Applying the Dialer Filter to the Dialer Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	From the [edit] hierarchy level, enter edit interfaces d10 unit 0
Select the dialer interface to apply the filter—for example, d10.	<ol style="list-style-type: none"> 1. In the Interface name column, click d10. 2. Under Unit, in the Mtu column, click Edit. 	
Select the dialer filter and apply it to the dialer interface.	<ol style="list-style-type: none"> 1. In the Family section, next to Inet, click Edit. 2. Next to Filter, click Configure. 3. In the Dialer box, type int-packet, the dialer-filter configured in “Configuring the Dialer Filter” on page 230, as the dialer-filter. 4. Click OK. 	<ol style="list-style-type: none"> 1. Enter edit family inet filter 2. Enter set dialer int-packet

Configuring Dialer Watch

Dialer watch is a backup method that integrates backup dialing with routing capabilities and provides reliable connectivity without relying on a dialer filter to trigger outgoing ISDN connections. With dialer watch, the Services Router monitors the existence of a specified route and if the route disappears, the dialer interface initiates the ISDN connection as a backup connection.

Adding a Dialer Watch Interface on the Services Router

To configure dialer watch:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 63 on page 232.
3. Go on to “Configuring the ISDN Interface for Dialer Watch” on page 232.

Table 63: Adding a Dialer Watch Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	From the [edit] hierarchy level, enter edit interfaces
Select a dialer interface—for example, d10 . Adding a description, such as dialer-watch , can help you identify one dialer interface from another.	<ol style="list-style-type: none"> 1. Under Interface name, select d10. 2. In the Description box, type dialer-watch. 	<ol style="list-style-type: none"> 1. Enter edit d10 2. Enter set description dialer-watch
On a logical interface—for example, 0 —specify a dial pool—for example, dw-group —to link the dialer interface to at least one ISDN physical interface. Then configure the list of routes for dialer watch—for example, 172.27.27.0/24 .	<ol style="list-style-type: none"> 1. Under Unit, click the logical unit number 0. 2. Next to Dialer options, click Edit. 3. In the Pool box, type dw-group. 4. Next to Watch list, click Add new entry. 5. In the Prefix box, type 172.27.27.0/24. 6. Click OK. 	<ol style="list-style-type: none"> 1. Enter edit unit 0 dialer-options 2. Enter set pool dw-group 3. Enter set watch-list 172.27.27.0/24

Configuring the ISDN Interface for Dialer Watch

To configure the ISDN interface to participate as a dialer watch interface:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 64 on page 233.
3. If you are finished configuring the router, commit the configuration.
4. Go on to any of the following optional tasks:
 - “Configuring Dial-on-Demand Routing Backup with OSPF Support (Optional)” on page 233.

- “Configuring Bandwidth on Demand (Optional)” on page 234.
 - Configuring Dial-In and Callback (Optional) on page 238
5. To verify that the network interface is configured correctly, see “Verifying the ISDN Configuration” on page 245.

Table 64: Configuring an ISDN Interface for Dialer Watch

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy, and select an ISDN physical interface—for example, <code>br-1/0/3</code> for ISDN BRI.	1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration .	From the [edit] hierarchy level:
For ISDN PRI, select a channelized T1/E1/ISDN PRI interface—for example, <code>ct1-1/0/1</code> .	2. Next to Interfaces, click Edit .	■ For ISDN BRI, enter edit interfaces br-1/0/3 dialer-options pool isdn-dialer-group
	3. Under Interface name: <ul style="list-style-type: none"> ■ For ISDN BRI, click br-1/0/3. ■ For ISDN PRI, click ct1-1/0/1. 	■ For ISDN PRI, enter edit interfaces ct1-1/0/1 dialer-options isdn-dialer-group
Configure dialer watch options for each ISDN interface participating in the dialer watch feature.	1. Next to Dialer options, click Edit .	
	2. Next to Pool, click Add new entry .	
Each ISDN interface must have the same pool identifier to participate in dialer watch. Therefore, the dialer pool name <code>isdn-dialer-group</code> , for the dialer watch interface configured in Table 63 on page 232, is used when configuring the ISDN interface.	3. In the Pool identifier box, type <code>isdn-dialer-group</code> .	
	4. Click OK .	

Configuring Dial-on-Demand Routing Backup with OSPF Support (Optional)

Two types of routing protocol traffic are used by OSPF to establish and maintain network structure. First, periodic hello packets are sent over each link for neighbor discovery and maintenance. Second, OSPF protocol traffic achieves and maintains link-state database synchronization between routers. The OSPF demand circuit feature removes the periodic nature of both traffic types and reduces the amount of OSPF traffic by removing all OSPF protocol traffic from a demand circuit when the routing domain is in a steady state. This feature allows the underlying data-link connections to be closed when no application traffic is on the network.

You must configure OSPF on the Services Router before configuring on-demand routing backup with OSPF support. For information on configuring OSPF, see “Configuring an OSPF Network” on page 411.

To configure OSPF demand circuits:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 65 on page 234.

3. If you are finished configuring the router, commit the configuration.
4. To verify that the network interface is configured correctly, see “Verifying the ISDN Configuration” on page 245.

Table 65: Configuring OSPF Demand Circuits

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Protocols level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Configure or Edit. 3. Next to Ospf, click Configure. 4. Next to Area, click Add new entry. 5. In the Area id box, type 0.0.0.0. 	<p>From the [edit] hierarchy level, enter</p> <p>edit protocols ospf area 0.0.0.0</p>
Configure OSPF on-demand circuits for each ISDN dialer interface participating as an on-demand routing interface—for example, dl0.	<ol style="list-style-type: none"> 1. Next to Interface, click Add new entry. 2. In the Interface name box, type dl0.0. 3. Select Demand circuit. 4. Click OK. 	<ol style="list-style-type: none"> 1. Enter <p>edit interface dl0</p> 2. Enter <p>set demand-circuit</p>

Configuring Bandwidth on Demand (Optional)

You can define a threshold for network traffic on the Services Router using the dialer interface and ISDN interfaces. A number of ISDN interfaces are aggregated together into a bundle and assigned a single dialer profile. Initially, only one ISDN link is active and all packets are sent through this interface. When a configured threshold is exceeded, the dialer interface activates another ISDN link and initiates a data connection. The threshold is specified as a percentage of the cumulative load of all **UP** links that are part of the bundle. When the cumulative load of all **UP** links, not counting the most recently activated link, is at or below the threshold, the most recently activated link is deactivated.

Configuring Dialer Interfaces for Bandwidth on Demand

To configure a dialer interface for bandwidth-on-demand:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 66 on page 235.
3. Go on to “Configuring an ISDN Interface for Bandwidth on Demand” on page 238.

Table 66: Configuring a Dialer Interface for Bandwidth on Demand

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy, and select a dialer interface—for example, d10 .	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 3. Next to d10, click Edit. 	From the [edit] hierarchy level, enter edit interfaces d10
Configure multilink properties on the dialer interface.	<ol style="list-style-type: none"> 1. Select multilink-ppp as the encapsulation type. 	Enter set encapsulation multilink-ppp
Configure the dialer options. <ul style="list-style-type: none"> ■ Dial string—Telephone number for the interface to dial that establishes ISDN connectivity—for example, 4085550115. You can configure a maximum of 15 dial strings per dialer interface. ■ Load interval—Interval of time used to calculate the average load on the dialer interface—for example, 90. Default value is 60 seconds with a range of 20-180 seconds. The value must be a multiple of 10. ■ Load threshold—Threshold above which an additional ISDN interface is activated, specified as a percentage of the cumulative load of all UP links—for example 95. Default value is 100 with a range of 0-100. ■ Pool—Name of a group of ISDN interfaces configured to use the dialer interface—for example, bw-pool. 	<ol style="list-style-type: none"> 1. In the Unit section, click Dialer options under Encapsulation. 2. Next to Dial string, click Add new entry. 3. In the Value box, type 4085550115 and click OK. 4. In the Load interval box, type 90. 5. In the Load threshold box, type 95. 6. In the Pool box, type bw-pool. 7. Click OK. 	<ol style="list-style-type: none"> 1. Enter edit unit 0 2. Enter edit dialer-options 3. Enter set dial-string 4085550115 4. Enter set load-interval 90 5. Enter set load-threshold 95 6. Enter set pool bw-pool
Configure unit properties. To configure multiple dialer interfaces for bandwidth-on-demand, increment the unit number—for example, d10.1, d10.2, and so on. F max period —Maximum number of compressed packets allowed between the transmission of full packets—for example, 100. The value can be between 1 and 65535.	<ol style="list-style-type: none"> 1. Next to Compression, select Yes, and then click Configure. 2. Select Rtp, and then click Configure. 3. In the F max period box, type 100. 4. Next to Queues, click Add new entry. 5. From the Value list, select q3. 6. Click OK until you return to the Unit page. 	<ol style="list-style-type: none"> 1. From the [edit] hierarchy level, enter edit interfaces d10 unit 0 2. Enter set compression rtp f-max-period 500 queues q3

Table 66: Configuring a Dialer Interface for Bandwidth on Demand (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure logical properties. ■ Fragment threshold —Maximum size, in bytes, for multilink packet fragments. The value can be between 128 and 16320 bytes, for example, 1024. The default is 0 bytes (no fragmentation). Any nonzero value must be a multiple of 64 bytes. ■ Maximum received reconstructed unit (MRRU) —This value is expressed as a number between 1500 and 4500 bytes—for example, 1500.	1. In the Fragment threshold box, type 1024. 2. In the Mrru box, type 1500. 3. Click OK until you return to the main Configuration page.	1. Enter set fragment-threshold 1024 2. Enter set mrru 1500
Define a CHAP access profile with a client and a secret password. For example, define bw-profile with client 1 and password my-secret.	1. On the main Configuration page next to Access, click Configure or Edit . 2. Next to Profile, click Add new entry . 3. In the Profile name box, type bw-profile. 4. Next to Client, click Add new entry . 5. In the Name box, type client1. 6. In the Chap secret box, type my-secret. 7. Click OK until you return to the main Configuration page.	From the [edit] hierarchy level, enter set access profile bw-profile client client1 chap-secret my-secret
Navigate to the appropriate dialer interface level in the configuration hierarchy—for example, d10 unit 0.	1. On the main Configuration page next to Interfaces, click Configure or Edit . 2. In the interface name box, click d10 . 3. In the Interface unit number box, click 0 .	From the [edit] hierarchy level, enter edit interfaces d10 unit 0
Configure CHAP on the dialer interface and specify a unique profile name containing a client list and access parameters—for example, bw-profile.	1. Next to Ppp options, click Configure . 2. Next to Chap, click Configure . 3. In the Access profile box, type bw-profile. 4. Click OK .	Enter set ppp-options chap access-profile bw-profile

Table 66: Configuring a Dialer Interface for Bandwidth on Demand (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure packet compression.	1. Under Compression, select Acfc .	Enter
You can configure the following compression types:	2. Click OK until you return to the Unit page.	set ppp-options compression acfc
<ul style="list-style-type: none"> ■ ACFC (address and control field compression)—Conserves bandwidth by compressing the address and control fields of PPP-encapsulated packets. ■ PFC (protocol field compression)—Conserves bandwidth by compressing the protocol field of a PPP-encapsulated packet. 		
Configure the dialer interface to be assigned an IP address in one of the following ways:	<p>Next to Inet, select Yes and click Configure.</p> <p>Select one of the following IP address configurations:</p> <p>To assign source and destination IP addresses:</p> <ol style="list-style-type: none"> 1. Next to Address, click Add new entry. 2. In the Source box, type 172.20.10.2. 3. In the Destination box, type 172.20.10.1. 4. Click OK. <p>To obtain an IP address from the remote end:</p> <ol style="list-style-type: none"> 1. Next to Negotiate address, select the Yes check box. 2. Click OK. <p>To derive the source address and assign the destination address:</p> <ol style="list-style-type: none"> 1. Next to Unnumbered address, select the Yes check box and click Configure. 2. In the Destination box, type 192.168.1.2. 3. In the Source box, type lo0.0. 4. Click OK. 	<p>Do one of the following:</p> <ul style="list-style-type: none"> ■ To assign source and destination IP addresses, enter set family inet address 172.20.10.2 destination 172.20.10.1 ■ To obtain an IP address from the remote end, enter set family inet negotiate-address ■ To derive the source address and assign the destination address, enter set family inet unnumbered-address lo0.0 destination 192.168.1.2

Configuring an ISDN Interface for Bandwidth on Demand

To configure bandwidth on demand on the ISDN interface:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 67 on page 238. Repeat these tasks for each ISDN interface participating in the aggregated link.
3. If you are finished configuring the router, commit the configuration.
4. To verify that the network interface is configured correctly, see “Verifying the ISDN Configuration” on page 245.

Table 67: Configuring an ISDN Interface for Bandwidth on Demand

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy, and select an ISDN BRI physical interface—for example, br-1/0/3 .	1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration .	From the [edit] hierarchy level:
For ISDN PRI, select a channelized T1/E1/ISDN PRI interface—for example, ct1-1/0/1 .	2. Next to Interfaces, click Edit . 3. Under Interface name: ■ For ISDN BRI, click br-1/0/3 . ■ For ISDN PRI, click ct1-1/0/1 .	■ For ISDN BRI, enter edit interfaces br-1/0/3 ■ For ISDN PRI, enter edit interfaces ct1-1/0/1
Because each ISDN interface must have the same pool identifier to participate in bandwidth on demand, use the dialer pool name bw-pool , the dialer interface configured in Table 66 on page 235, to configure the ISDN interfaces participating in the pool.	1. Next to Dialer options, click Dialer options . 2. Next to Pool, click Add new entry . 3. In the Pool identifier box, type the name of the dialer pool—for example, bw-pool .	Enter set dialer-options pool bw-pool
For ISDN BRI, you can group up to four ISDN interfaces together when configuring bandwidth on demand, for a total of eight B-channels (two channels per interface) providing connectivity.	4. Click OK .	
For ISDN PRI, the pool limit is eight B-channels per channelized T1/E1/ISDN PRI port.		

Configuring Dial-In and Callback (Optional)

If you are a service provider or a corporate data center into which a remote location dials in during an emergency, you can configure the Services Router to accept incoming ISDN calls originating from the remote location, or reject the incoming calls and call back the remote location. The callback feature lets you control access by allowing only specific remote locations to connect to the Services Router. You can also configure the Services Router to reject all incoming ISDN calls.



NOTE: Incoming voice calls are currently not supported.

When it receives an incoming ISDN call, the Services Router matches the incoming call's caller ID against the caller IDs configured on its dialer interfaces. If an exact match is not found and the incoming call's caller ID has more digits than the configured caller IDs, the Services Router performs a right-to-left match of the incoming call's caller ID with the configured caller IDs and accepts the incoming call if a match is found. For example, if the incoming call's caller ID is **4085550115** and the caller ID configured on a dialer interface is **5550115**, the incoming call is accepted. Each dialer interface accepts calls from only callers whose caller IDs are configured on it.

The dialer interface of the Services Router and the dialer interface of the remote router must have the same encapsulation—PPP, Multilink PPP, or Cisco HDLC. If the encapsulation is different, the ISDN call is dropped. Table 68 on page 239 describes how the Services Router performs encapsulation monitoring.

Table 68: Encapsulation Monitoring by Services Routers

Encapsulation on Services Router's Dialer Interface	Encapsulation on Remote Router's Dialer Interface	Possible Action on Services Router's Dialer Interface	Encapsulation Monitoring and Call Status
PPP	PPP	■ Accepts an incoming call	Services Router performs encapsulation monitoring.
Multilink PPP	Multilink PPP	■ Rejects an incoming call and calls back the incoming number when callback is enabled on the dialer interface	
PPP	Multilink PPP or Cisco HDLC		Services Router performs encapsulation monitoring.
Multilink PPP	PPP or Cisco HDLC		ISDN call is <i>dropped</i> because of encapsulation mismatch.
PPP or Multilink PPP	PPP, Multilink PPP, or Cisco HDLC	■ Dials out ■ Accepts an incoming call as a result of having originally dialed out, because the dialer interface of the remote router has callback enabled	Services Router does not perform encapsulation monitoring. Success of the ISDN call depends on the encapsulation monitoring capability of the remote router.
Cisco HDLC	PPP, Multilink PPP, or Cisco HDLC	■ Dials out ■ Accepts an incoming call ■ Accepts an incoming call as a result of having originally dialed out, because the dialer interface of the remote router has callback enabled ■ Rejects an incoming call and calls back the incoming number when callback is enabled on the dialer interface	

This section contains the following topics:

- Configuring Dialer Interfaces for Dial-In and Callback on page 240
- Configuring an ISDN Interface to Screen Incoming Calls on page 242
- Configuring the Services Router to Reject Incoming ISDN Calls on page 243

Configuring Dialer Interfaces for Dial-In and Callback

To configure a dialer interface for dial-in and callback:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 69 on page 240.
3. If you are finished configuring the router, commit the configuration.
4. To verify that the network interface is configured correctly, see “Verifying the ISDN Configuration” on page 245.

Table 69: Configuring the Dialer Interface for Dial-In and Callback

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy, and select a dialer interface—for example, d10.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 3. Next to d10, click Edit. 	From the [edit] hierarchy level, enter edit interfaces d10

Table 69: Configuring the Dialer Interface for Dial-In and Callback *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>On a logical interface—for example, 0—configure the incoming map options for the dialer interface. For Services Routers to use dial-in, you must configure an incoming map on the dialer interface.</p> <ul style="list-style-type: none"> Accept all—Dialer interface accepts all incoming calls. You can configure this option for only one of the dialer interfaces associated with an ISDN physical interface. The dialer interface configured to accept all calls is used only if the incoming call's caller ID does not match the caller IDs configured on other dialer interfaces. Caller—Dialer interface accepts calls from a specific caller ID—for example, 4085550115. You can configure a maximum of 15 caller IDs per dialer interface. The same caller ID must not be configured on different dialer interfaces. However, you can configure caller IDs with more or fewer digits on different dialer interfaces. For example, you can configure the caller IDs 14085550115, 4085550115, and 5550115 on different dialer interfaces. 	<ol style="list-style-type: none"> In the Unit section, for logical unit number 0, click Dialer options under Encapsulation. Next to Incoming map, click Configure. From the Caller type menu, select Caller. Next to Caller, click Add new entry. In the Caller id box, type 4085550115. 	<ol style="list-style-type: none"> Enter edit unit 0 Enter edit dialer-options Enter set incoming-map caller 4085550115
<p>Configure callback options for the dialer interface</p> <ul style="list-style-type: none"> Callback—Enable this feature to allow the ISDN interface to reject incoming calls, wait for 5 seconds (the default callback wait period), and then call back the incoming number. Before configuring callback on a dialer interface, ensure that the following conditions exist: <ul style="list-style-type: none"> 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. Callback wait period—Number of seconds to wait before redialing an incoming ISDN call. 	<ol style="list-style-type: none"> Select Callback. In the Callback wait period box, type 5. 	<ol style="list-style-type: none"> Enter set callback Enter set callback-wait-period 5

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 configure the incoming called numbers that you want an ISDN interface to accept. You can also use the reject option to configure a called number that you want an ISDN interface to ignore because the number belongs to another device connected to the same ISDN line. For example, if another device on the same ISDN line has the called number 4085551091, you can configure the called number 4085551091 with the reject option on the ISDN interface so that it does not accept calls with that number.

When it receives an incoming ISDN call, 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 numbers are configured on it.

To configure an ISDN interface to screen incoming ISDN calls:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 70 on page 242.
3. If you are finished configuring the router, commit the configuration.
4. To verify that the network interface is configured correctly, see “Verifying the ISDN Configuration” on page 245.

Table 70: Configuring an ISDN Interface to Screen Incoming ISDN Calls

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy, and select an ISDN physical interface—for example, br-1/0/3 .	1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration .	From the [edit] hierarchy level:
For ISDN BRI, select a channelized T1/E1/ISDN PRI interface—for example, ct1-1/0/1 .	2. Next to Interfaces, click Edit .	■ For ISDN BRI, enter edit interfaces br-1/0/3
	3. Under Interface name:	■ For ISDN PRI, enter edit interfaces ct1-1/0/1
	■ For ISDN BRI, click br-1/0/3 .	
	■ For ISDN PRI, click ct1-1/0/1 .	
Configure the incoming called number—for example, 4085550115 —for the ISDN interface.	1. Next to Isdn options, click Edit .	Enter
	2. Next to Incoming called number, click Add new entry .	set isdn-options incoming-called-number
To configure the ISDN interface to ignore the incoming called number, use the reject option.	3. In the Called number box, type 4085550115 .	4085550115
	4. Click OK .	

Configuring the Services Router to Reject Incoming ISDN Calls

By default, the Services Router is configured to accept incoming ISDN calls. The incoming calls are accepted if dial-in is configured on the Services Router. You can configure the Services Router to reject all incoming ISDN calls.

To configure the Services Router to reject incoming ISDN calls:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 71 on page 243.
3. If you are finished configuring the router, commit the configuration.
4. To verify that the network interface is configured correctly, see “Verifying the ISDN Configuration” on page 245.

Table 71: Configuring the Services Router to Reject Incoming ISDN Calls

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Processes level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to System, click Edit. 3. Next to Processes, click Configure. 4. Next to Isdn signaling, click Configure. 	<p>From the [edit] hierarchy level, enter</p> <p>set system processes isdn-signaling reject-incoming</p>
Configure the Services Router to reject incoming calls.	<ol style="list-style-type: none"> 1. Select the Reject Incoming check box. 2. Click OK. 	

Disabling Dialing Out Through Dialer Interfaces

The JUNOS ISDN dialer services process manages dialing out through dialer interfaces. You can disable dialing out through all dialer interfaces by disabling the dialer services process.



CAUTION: Never disable a software processes unless instructed to do so by a Customer Support engineer.

To disable dialing out through dialer interfaces:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 72 on page 244.
3. If you are finished configuring the router, commit the configuration.

Table 72: Disabling Dialing Out Through Dialer Interfaces

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Processes level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to System, click Edit. 3. Next to Processes, click Configure. 4. Next to Dialer services, click Configure. 	<p>From the [edit] hierarchy level, enter</p> <p>set system processes dialer-services disable</p>
Disable dialing out through dialer interfaces.	<ol style="list-style-type: none"> 1. Select the Disable check box. 2. Click OK. 	

Disabling ISDN Signaling

The JUNOS ISDN signaling process manages ISDN signaling by initializing ISDN connections. You can disable ISDN signaling by disabling the ISDN signaling process.



CAUTION: Never disable a software processes unless instructed to do so by a Customer Support engineer.

To disable ISDN signaling:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 73 on page 244.
3. If you are finished configuring the router, commit the configuration.

Table 73: Disabling ISDN Signaling

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Processes level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to System, click Edit. 3. Next to Processes, click Configure. 4. Next to Isdn signaling, click Configure. 	<p>From the [edit] hierarchy level, enter</p> <p>set system processes isdn-signaling disable</p>
Disable ISDN signaling on the Services Router.	<ol style="list-style-type: none"> 1. Select the Disable check box. 2. Click OK. 	

Verifying the ISDN Configuration

To verify an ISDN configuration, perform the following tasks:

- Displaying the ISDN Status on page 246
- Verifying an ISDN BRI Interface on page 247
- Verifying an ISDN PRI Interface and Checking B-Channel Interface Statistics on page 249
- Checking D-Channel Interface Statistics on page 251
- Displaying the Status of ISDN Calls on page 253
- Verifying Dialer Interface Configuration on page 255

Displaying the ISDN Status

Purpose Display the status of ISDN service on the ISDN interface. For example, you can display ISDN BRI status on the **br-6/0/0** interface and ISDN PRI status on the **ct1-2/0/0** interface.

Action From the operational mode in the CLI, enter **show isdn status**.

```
user@host> show isdn status
Interface: br-6/0/0
Layer 1 status: active
Layer 2 status:
  CES: 0, Q.921: up, TEI: 12
Layer 3 status: 1 Active calls
Switch Type      : ETSI
Interface Type   : USER
T310             : 10 seconds
Tei Option       : Power Up
```

```
user@host> show isdn status
Interface: ct1-2/0/0
Layer 1 status: active
Layer 2 status:
  CES: 0, Q.921: up, TEI: 0
Layer 3 status: 8 Active calls
Switch Type      : NI2
Interface Type   : USER
T310             : 10 seconds
Tei Option       : Power Up
```

What It Means The output shows a summary of interface information. Verify the following information:

- **Interface**—ISDN interface currently active on the Services Router. For ISDN BRI, the interface is a **br-pim/0/port** interface, as shown in the first example for **br-6/0/0**. For ISDN PRI, the interface displayed is a channelized T1 or channelized E1 interface, as shown in the second example for **ct1-2/0/0**.
- **Layer 1 status**—Indicates whether Layer 1 is active or inactive.
- **Layer 2 status**—Indicates whether Q.921 (the D-channel protocol) is up or down.
- **TEI**—Assigned terminal endpoint identifier (TEI) number.
- **Layer 3 status**—Number of active calls.
- **Switch Type**—Type of ISDN switch connected to the Services Router interface.
- **Interface Type**—Default value for the local interface.
- **Calling number**—Telephone number configured for dial-out.
- **T310**—Q.931-specific timer.
- **TEI Option**—Indicates when TEI negotiations occur on the interface.

See Also For a complete description of **show isdn status** output, see the *JUNOS Interfaces Command Reference*.

Verifying an ISDN BRI Interface

Purpose Verify that the ISDN BRI interface is correctly configured.

Action From the CLI, enter the `show interfaces extensive` command. Alternatively, from the J-Web interface select **Monitor > Interfaces > br-6/0/0**.

```
user@host> show interfaces br-6/0/0 extensive
Physical interface: br-6/0/0, Enabled, Physical link is Up
  Interface index: 143, SNMP ifIndex: 59, Generation: 24
  Type: BRI, Link-level type: Controller, MTU: 4092, Clocking: 1, Speed: 144kbps,

  Parent: None
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps Internal: 0x4000
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : S/T
  Hold-times    : Up 0 ms, Down 0 ms
  Last flapped   : 2005-12-07 12:21:11 UTC (04:07:26 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :                0                0 bps
    Output bytes  :                0                0 bps
    Input packets :                0                0 pps
    Output packets:                0                0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Giants: 0, Policed discards:
0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
```

What It Means The output shows a summary of interface information. Verify the following information:

- The physical interface is **Enabled**. If the interface is shown as **Disabled**, do either of the following:
 - In the CLI configuration editor, delete the `disable` statement at the [edit *interfaces interface-name*] level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the **Interfaces > interface-name** page.
- The physical link is **Up**. A link state of **Down** indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The **Last Flapped** time is an expected value. The **Last Flapped** time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the `clear interfaces statistics interface-name` command.

See Also For a complete description of `show interfaces` (ISDN BRI) output, see the *JUNOS Interfaces Command Reference*.

Verifying an ISDN PRI Interface and Checking B-Channel Interface Statistics

- Purpose** Verify that an ISDN B-channel interface is operating properly. For ISDN PRI, verify that a B-channel interface is configured correctly. (To display a list of B-channel interfaces, enter the `show isdn calls` command.)
- Action** From the CLI, enter the `show interfaces extensive` command. Alternatively, from the J-Web interface select **Monitor > Interfaces > bc-0/0/4:1**.

```

user@host> show interfaces bc-0/0/4:1 extensive
Physical interface: bc-0/0/4:1, Administratively down, Physical link is Up
Interface index: 145, SNMP ifIndex: 75, Generation: 26
Type: Serial, Link-level type: Multilink-PPP, MTU: 1510, Clocking: Internal,
Speed: 64kbps,
Parent: br-0/0/4 Interface index 143
Device flags : Present Running
Interface flags: Admin-Test SNMP-Traps 16384
Link type : Full-Duplex
Link flags : None
Physical info : Unspecified
Hold-times : Up 0 ms, Down 0 ms
Current address: Unspecified, Hardware address: Unspecified
Alternate link address: Unspecified
CoS queues : 8 supported, 8 maximum usable queues
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes : 5787 0 bps
Output bytes : 3816 0 bps
Input packets: 326 0 pps
Output packets: 264 0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
6,
Resource errors: 0
Output errors:
Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
Queue counters Queued packets Transmitted Packets Dropped packets
0 best-effort 314335 0 0
1 best-effort 0 0 0
2 best-effort 5 0 0
3 best-effort 5624 5624 0
Packet Forwarding Engine configuration:
Destination slot: 5, PLP byte: 1 (0x00)
CoS transmit queue Bandwidth Buffer Priority
Limit % bps % usec
0 best-effort 95 60800 95 0 low
none
3 network-control 5 3200 5 0 low
none

Logical interface bc-0/0/4:1.0 (Index 71) (SNMP ifIndex 61) (Generation 33)
Flags: Device Down Point-To-Point SNMP-Traps Encapsulation: PPP
Protocol mlppp, Multilink bundle: dl0.0, MTU: 1506, Generation: 18, Route
table: 0

```

What It Means The output shows a summary of B-channel interface information. Verify the following information:

- The physical interface is **Enabled**. If the interface is shown as **Disabled**, do either of the following:
 - In the CLI configuration editor, delete the **disable** statement at the **[edit interfaces *interface-name*]** level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the **Interfaces > *interface-name*** page.
- The physical link is **Up**. A link state of **Down** indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- For ISDN BRI, the **Parent** interface is a **br-*pim*/0/*port*** interface—**br-0/0/4** in this example. For ISDN PRI, the **Parent** interface is a channelized T1 or channelized E1 interface—**ct1-*pim*/0/*port*** or **ce1-*pim*/0/*port***.
- The **Last Flapped** time is an expected value. The **Last Flapped** time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates possible link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics *interface-name*** command.

See Also For a complete description of **show interfaces** (ISDN B-channel) output, see the *JUNOS Interfaces Command Reference*.

Checking D-Channel Interface Statistics

Purpose Verify that the ISDN D-channel interface is operating properly. For ISDN PRI, verify that the D-channel interface is configured correctly.

Action From the CLI, enter the `show interfaces extensive` command. Alternatively, from the J-Web interface select **Monitor > Interfaces > dc-0/0/4**.

```
user@host> show interfaces dc-0/0/4 extensive
Physical interface: dc-0/0/4, Enabled, Physical link is Up
  Interface index: 144, SNMP ifIndex: 60, Generation: 25
  Type: Serial, Link-level type: 55, MTU: 4092, Clocking: Internal, Speed: 16kbps,

  Parent: br-0/0/4 Interface index 143
  Device flags   : Present Running
  Interface flags: SNMP-Traps Internal: 0x4000
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped   : 2005-12-07 12:21:12 UTC (05:46:00 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :                13407                0 bps
    Output bytes  :                16889                0 bps
    Input packets :                3262                0 pps
    Output packets:                3262                0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 1, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
  ISDN alarms    : None
  ISDN media:
    Seconds      Count  State
    LOF          0      1  OK
    LOS          0      0  OK

  Logical interface dc-0/0/4.32767 (Index 70) (SNMP ifIndex 72) (Generation 8)
  Flags: Point-To-Point SNMP-Traps Encapsulation: 60
  Traffic statistics:
    Input bytes   :                13407
    Output bytes  :                82129
    Input packets :                3262
    Output packets:                3262
  Local statistics:
    Input bytes   :                13407
    Output bytes  :                82129
    Input packets :                3262
    Output packets:                3262
```

What It Means The output shows a summary of D-channel interface information. Verify the following information:

- The physical interface is **Enabled**. If the interface is shown as **Disabled**, do either of the following:

- In the CLI configuration editor, delete the **disable** statement at the [edit **interfaces** *interface-name*] level of the configuration hierarchy.
- In the J-Web configuration editor, clear the **Disable** check box on the **Interfaces** > *interface-name* page.
- The physical link is **Up**. A link state of **Down** indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- For ISDN BRI, the **Parent** interface is a *br-pim/0/port* interface—*br-0/0/4* in this example. For ISDN PRI, the **Parent** interface is a channelized T1 or channelized E1 interface—*ct1-pim/0/port* or *ce1-pim/0/port*.
- The **Last Flapped** time is an expected value. The **Last Flapped** time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates possible link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics** *interface-name* command.

See Also For a complete description of **show interfaces** (ISDN D-channel) output, see the *JUNOS Interfaces Command Reference*.

Displaying the Status of ISDN Calls

Purpose Display the status of ISDN calls. This information helps you to verify the dialer interface configuration as described in Verifying Dialer Interface Configuration on page 255. The command also provides a list of the B-channels configured on an ISDN BRI or ISDN PRI interface.

Action From the CLI, enter the `show isdn calls` command.

```
user@host> show isdn calls
Interface: bc-6/0/0:1
  Status: No call in progress
  Most recent error code: No error
Interface: bc-6/0/0:2
  Status: Connected to 384070
  Call Duration: 43 seconds
  Call Direction: Dialout
  Most recent error code: No error

user@host> show isdn calls
Interface: bc-2/0/0:1
  Status: Connected to 384010
  Call Duration: 49782 seconds
  Call Direction: Dialin
  Most recent error code: destination out of order
Interface: bc-2/0/0:2
  Status: Connected to 384011
  Call Duration: 49782 seconds
  Call Direction: Dialin
  Most recent error code: destination out of order
Interface: bc-2/0/0:3
  Status: Connected to 384020
  Call Duration: 49782 seconds
  Call Direction: Dialin
  Most recent error code: destination out of order
...
Interface: bc-2/0/0:20
  Status: No call in progress
  Most recent error code: No error
Interface: bc-2/0/0:21
  Status: No call in progress
  Most recent error code: No error
Interface: bc-2/0/0:22
  Status: No call in progress
  Most recent error code: No error
Interface: bc-2/0/0:23
  Status: No call in progress
  Most recent error code: No error
```

What It Means The output shows a summary of B-channel interfaces and the active ISDN calls on the interfaces. The first example shows the two B-channels on an ISDN BRI interface—`bc-2/0/0:1` and `bc-2/0/0:2`. The second example indicates B-channels `bc-2/0/0:1` through `bc-2/0/0:23`, the 23 B-channels on an ISDN PRI interface. Determine the following information:

- The interfaces on which ISDN calls are in progress
- Whether the call is a dial-in call, dial-out call, or a callback call

See Also For a complete description of `show isdn calls` output, see the *JUNOS Interfaces Command Reference*.

Verifying Dialer Interface Configuration

Purpose Verify that the dialer interface is correctly configured. To determine the ISDN interfaces on which calls are taking place, see *Displaying the Status of ISDN Calls* on page 253.

Action From the CLI, enter the `show interfaces d10 extensive` command. Alternatively, from the J-Web interface select **Monitor > Interfaces > d10**.

```
user@host> show interfaces d10 extensive
Physical interface: d10, Enabled, Physical link is Up
  Interface index: 173, SNMP ifIndex: 26, Generation: 77
  Type: 27, Link-level type: PPP, MTU: 1504, Clocking: Unspecified, Speed:
Unspecified
  Device flags      : Present Running
  Interface flags:  SNMP-Traps
  Link type        : Full-Duplex
  Link flags       : Keepalives
  Physical info    : Unspecified
  Hold-times       : Up 0 ms, Down 0 ms
  Current address:  Unspecified, Hardware address: Unspecified
  Alternate link address: Unspecified
  Last flapped     : Never
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :           13859           0 bps
    Output bytes  :              0           0 bps
    Input packets :            317           0 pps
    Output packets:              0           0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runt: 0, Giants: 0, Policed discards:
0,
    Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0

Logical interface d10.0 (Index 76) (SNMP ifIndex 28) (Generation 148)
  Flags: Point-To-Point SNMP-Traps 0x4000 LinkAddress 23-0 Encapsulation: PPP
  Dialer:
    State: Active, Dial pool: 1
    Dial strings: 384070
    Subordinate interfaces: bc-6/0/0:2 (Index 172)
    Watch list: 11.12.13.14/32
    Activation delay: 0, Deactivation delay: 0
    Initial route check delay: 120
    Redial delay: 3
    Callback wait period: 5
    Load threshold: 0, Load interval: 60
  Bandwidth: 64kbps
  Traffic statistics:
    Input bytes   :           24839
    Output bytes  :           17792
    Input packets :            489
    Output packets:            340
  Local statistics:
    Input bytes   :           10980
    Output bytes  :           17792
    Input packets :             172
```

```
Output packets:          340
Transit statistics:
Input bytes  :          13859          0 bps
Output bytes :           0          0 bps
Input packets:          317          0 pps
Output packets:           0          0 pps
Keepalive settings: Interval 10 seconds, Up-count 1, Down-count 3
Keepalive statistics:
Input : 0 (last seen: never)
Output: 36 (last sent 00:00:09 ago)
LCP state: Opened
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured,
mpls: Not-configured
CHAP state: Success
Protocol inet, MTU: 1500, Generation: 74, Route table: 0
Flags: Negotiate-Address
Addresses, Flags: Kernel Is-Preferred Is-Primary
Destination: 43.1.1.2, Local: 43.1.1.19, Broadcast: Unspecified,
Generation: 37
```

```

user@host> show interfaces d10 extensive
Physical interface: d10, Enabled, Physical link is Up
  Interface index: 140, SNMP ifIndex: 35, Generation: 141
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2007-02-27 01:50:38 PST (1d 15:48 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes :          42980144          0 bps
    Output bytes :           504          0 bps
    Input packets:         934346          0 pps
    Output packets:           6          0 pps
  Frame exceptions:
    Oversized frames          0
    Errored input frames      0
    Input on disabled link/bundle 0
    Output for disabled link/bundle 0
    Queuing drops            0
  Buffering exceptions:
    Packet data buffer overflow 0
    Fragment data buffer overflow 0
  Assembly exceptions:
    Fragment timeout          0
    Missing sequence number    0
    Out-of-order sequence number 0
    Out-of-range sequence number 0
  Hardware errors (sticky):
    Data memory error          0
    Control memory error        0
  Egress queues: 8 supported, 8 in use
  Queue counters:

```

	Queued packets	Transmitted packets	Dropped packets
0 q1	6	6	0
1 q2	0	0	0
2 assured-forw	0	0	0
3 q3	0	0	0

```

Logical interface d10.0 (Index 66) (SNMP ifIndex 36) (Generation 133)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
  Dialer:
    State: Active, Dial pool: 1
    Dial strings: 384010
    Subordinate interfaces: bc-2/0/0:8 (Index 161), bc-2/0/0:7 (Index 160),
      bc-2/0/0:6 (Index 159), bc-2/0/0:5 (Index 158), bc-2/0/0:4 (Index 157),
      bc-2/0/0:3 (Index 156), bc-2/0/0:2 (Index 155), bc-2/0/0:1 (Index 154)
    Activation delay: 0, Deactivation delay: 0
    Initial route check delay: 120
    Redial delay: 3
    Callback wait period: 5
    Load threshold: 100, Load interval: 60
  Bandwidth: 512kbps
  Bundle options:
    MRRU          1504
    Remote MRRU    1504
    Drop timer period 0
    Inner PPP Protocol field compression enabled

```

```

Sequence number format          long (24 bits)
Fragmentation threshold         0
Links needed to sustain bundle  1
Interleave fragments            Disabled
Bundle errors:
  Packet drops                  0 (0 bytes)
  Fragment drops                15827 (759696 bytes)
  MRRU exceeded                 0
  Exception events              0
Statistics      Frames      fps      Bytes      bps
Bundle:
  Fragments:
    Input :      963116      0      50963104      0
    Output:       6         0         540         0
  Packets:
    Input :     934346      0      42980144      0
    Output:       6         0         504         0
Link:
  bc-2/0/0:1.0
    Input :     119656      0      6341806      0
    Output:       1         0         90         0
  bc-2/0/0:2.0
    Input :     120176      0      6369366      0
    Output:       1         0         90         0
  bc-2/0/0:3.0
    Input :     119856      0      6352368      0
    Output:       1         0         90         0
  bc-2/0/0:4.0
    Input :     120315      0      6376695      0
    Output:       0         0         0         0
  bc-2/0/0:5.0
    Input :     120181      0      6369593      0
    Output:       0         0         0         0
  bc-2/0/0:6.0
    Input :     121154      0      6421200      0
    Output:       0         0         0         0
  bc-2/0/0:7.0
    Input :     121181      0      6340321      0
    Output:       0         0         0         0
  bc-2/0/0:8.0
    Input :     120594      0      6391482      0
    Output:       0         0         0         0
NCP state: inet: Opened, inet6: Not-configured, iso: Not-configured, mpls:
Not-configured
Protocol inet, MTU: 1500, Generation: 138, Route table: 0
Flags: None
Addresses, Flags: Is-Preferred Is-Primary
Destination: 1.1.1.0/30, Local: 1.1.1.2, Broadcast: Unspecified,
Generation: 134

```

What It Means The output shows a summary of dialer interface information. The first example is for ISDN BRI service, and the second example is for ISDN PRI service. Verify the following information:

- The physical interface is **Enabled**. If the interface is shown as **Disabled**, do either of the following:
 - In the CLI configuration editor, delete the **disable** statement at the **[edit interfaces *interface-name*]** level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the **Interfaces > *interface-name*** page.
- The physical link is **Up**. A link state of **Down** indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The **Last Flapped** time is an expected value. The **Last Flapped** time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates possible link-layer errors.
- **Subordinate interfaces** correctly lists the B-channel interface or interfaces associated with this dialer interface. The ISDN BRI output in the first example shows that **dl0** supports **bc-6/0/0:2**.

The ISDN PRI output in the second example shows that **dl0** supports **bc-2/0/0:1** through **bc-2/0/0:8**.

- The traffic statistics reflect expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics *interface-name*** command.
- The dialer state is **Active** when an ISDN call is in progress.
- The LCP state is **Opened** when an ISDN call is in progress. An LCP state of **Closed** or **Not Configured** indicates a problem with the dialer configuration that needs to be debugged with the **monitor traffic interface *interface-name*** command. For information about the **monitor traffic** command, see the *J-series Services Router Administration Guide*.

See Also For a complete description of **show interfaces** (ISDN dialer) output, see the *JUNOS Interfaces Command Reference*.

Chapter 8

Configuring Link Services Interfaces

Link services include the multilink services Multilink Point-to-Point Protocol (MLPPP), Multilink Frame Relay (MLFR), and Compressed Real-Time Transport Protocol (CRTP). J-series Services Routers support link services on the `ls-0/0/0` link services interface.

You can use either J-Web Quick Configuration or a configuration editor to configure the link services interface.

This chapter contains the following topics:

- Link Services Terms on page 261
- Link Services Interfaces Overview on page 262
- Before You Begin on page 270
- Configuring the Link Services Interface with Quick Configuration on page 271
- Configuring the Link Services Interface with a Configuration Editor on page 273
- Verifying the Link Services Interface Configuration on page 291
- Frequently Asked Questions About the Link Services Interface on page 301

Link Services Terms

Before configuring a link services interface, become familiar with the terms defined in Table 74 on page 261.

Table 74: Link Services Terminology

Term	Definition
Compressed Real-Time Transport Protocol (CRTP)	Protocol defined in RFC 2508 that compresses the size of IP, UDP, and Real-Time Transport Protocol (RTP) headers and works with reliable and fast point-to-point links for voice over IP (VoIP) traffic.
data-link connection identifier (DLCI)	Identifier for a Frame Relay virtual connection, also called a logical interface.

Table 74: Link Services Terminology (continued)

Term	Definition
link fragmentation and interleaving (LFI)	For MLFR with Frame Relay traffic or MLPPP with PPP traffic, a method of reducing excessive delays by fragmenting long packets into smaller packets and interleaving them with real-time frames. For example, short delay-sensitive packets, such as those of packetized voice, can race ahead of larger delay-insensitive packets, such as common data packets.
link services	Capabilities on an interface that use Multilink Frame Relay (MLFR) and Multilink Point-to-Point Protocol (MLPPP), link fragmentation and interleaving (LFI), Compressed Real-Time Transport Protocol (CRTP), and certain class-of-service (CoS) components to improve packet transmission, especially for time-sensitive voice packets.
Multilink Frame Relay (MLFR)	Protocol that allows multiple Frame Relay links to be aggregated by inverse multiplexing.
Multilink Point-to-Point Protocol (MLPPP)	Protocol that allows you to bundle multiple Point-to-Point Protocol (PPP) links into a single logical unit. MLPPP improves bandwidth efficiency and fault tolerance and reduces latency.
Point-to-Point Protocol (PPP)	Link-layer protocol defined in RFC 1661 that provides multiprotocol encapsulation. PPP is used for link-layer and network-layer configuration.
shaping rate	In class of service (CoS) classification, a method of controlling the maximum rate of traffic transmitted on an interface.

Link Services Interfaces Overview

You configure the link services interface (**ls-0/0/0**) on a J-series Services Router to support multilink services and Compressed Real-Time Transport Protocol (CRTP).

The link services interface on a Services Router consists of services provided by the following interfaces on the Juniper M-series and T-series routing platforms: multilink services interface (**ml-fpc/pic/port**), link services interface (**ls-fpc/pic/port**), and link services intelligent queuing interface (**lsq-fpc/pic/port**). Although the multilink services, link services, and link services intelligent queuing (IQ) interfaces on M-series and T-series routing platforms are installed on Physical Interface Cards (PICs), the link services interface on a J-series Services Router is an internal interface only and is not associated with a physical medium or Physical Interface Module (PIM).

For information about interface names, see “Network Interface Naming” on page 47.

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

This section contains the following topics.

- Services Available on J-series Link Services Interface on page 263
- Link Services Exceptions on J-series Services Routers on page 264

- Multilink Bundles Overview on page 264
- Link Fragmentation and Interleaving Overview on page 265
- Compressed Real-Time Transport Protocol Overview on page 266
- Queuing with LFI on J-series Services Routers on page 267
- Load Balancing with LFI on page 268
- Configuring CoS Components with LFI on page 269

Services Available on J-series Link Services Interface

On a Services Router, the link services interface is a logical interface available by default. Table 75 on page 263 summarizes the services available on a J-series link services interface.

Table 75: Services Available on J-series Link Services Interface

Services	Purpose	More Information
Multilink bundles by means of MLPPP and MLFR encapsulation	Aggregates multiple constituent links into one larger logical bundle to provide additional bandwidth, load balancing, and redundancy.	<ul style="list-style-type: none"> ■ Configuring an MLPPP Bundle on page 274 ■ Configuring MLFR FRF.15 Bundles on page 284 ■ Configuring MLFR FRF.16 Bundles on page 287
Link fragmentation and interleaving (LFI)	Reduces delay and jitter on links by breaking up large data packets and interleaving delay-sensitive voice packets with the resulting smaller packets.	“Link Fragmentation and Interleaving Overview” on page 265
Compressed Real-Time Transport Protocol (CRTP)	Reduces the overhead caused by Real-Time Transport Protocol (RTP) on voice and video packets.	“Compressed Real-Time Transport Protocol Overview” on page 266
Class-of-service (CoS) classifiers, forwarding classes, schedulers and scheduler maps, and shaping rates	Provide a higher priority to delay-sensitive packets—by configuring class of service (CoS) components, such as the following: <ul style="list-style-type: none"> ■ Classifiers—To classify different type of traffic, such as voice, data and network control packets ■ Forwarding classes—To direct different types of traffic to different output queues ■ Schedulers and scheduler maps—To define properties for the output queues such as delay-buffer, transmission rate, and transmission priority ■ Shaping rate—To define certain bandwidth usage by an interface 	<ul style="list-style-type: none"> ■ Defining Classifiers and Forwarding Classes on page 277 ■ Defining and Applying Scheduler Maps on page 279 ■ Applying Shaping Rates to Interfaces on page 283 ■ (For more information about CoS) <i>J-series Services Router Advanced WAN Access Configuration Guide</i>

Link Services Exceptions on J-series Services Routers

The link and multilink services implementation on a J-series Services Router is similar to the implementation on the M-series and T-series routing platforms, with the following exceptions:

- A Services Router supports link and multilink services on the `ls-0/0/0` interface instead of the `ml-fpc/pic/port`, `lsq-fpc/pic/port`, and `ls-fpc/pic/port` interfaces.
- When LFI is enabled on a Services Router, Queue 2 is reserved for voice traffic, while all other queues perform fragmentation. Also, the queuing behavior on the link services interface and constituent links is different. For more information, see “Queuing with LFI on J-series Services Routers” on page 267.
- When LFI is enabled on a Services Router, fragmented packets are queued in a round-robin fashion on the constituent links to enable per-packet and per-fragment load balancing. For more information, see “Queuing with LFI on J-series Services Routers” on page 267.
- A Services Router supports per-unit scheduling on all types of constituent links (on all types of interfaces).
- A Services Router supports Compressed Real-Time Transport Protocol (CRTP) with MLPPP as well as PPP.
- A Services Router does not support multiclass MLPPP.
- A Services Router does not have the ability to apply fragmentation maps to specific queues to enable LFI on specific queues (a multiclass MLPPP feature).

Multilink Bundles Overview

The J-series Services Router supports MLPPP and MLFR multilink encapsulations. MLPPP enables you to bundle multiple PPP links into a single multilink bundle, and MLFR enables you to bundle multiple Frame Relay data-link connection identifiers (DLCIs) into a single multilink bundle. Multilink bundles provide additional bandwidth, load balancing, and redundancy by aggregating low-speed links, such as T1, E1, and serial links.

You configure multilink bundles as logical units or channels on the link services interface `ls-0/0/0`:

- With MLPPP and MLFR FRF.15, multilink bundles are configured as logical units on `ls-0/0/0`—for example, `ls-0/0/0.0` and `ls-0/0/0.1`.
- With MLFR FRF.16, multilink bundles are configured as channels on `ls-0/0/0`—for example, `ls-0/0/0:0` and `ls-0/0/0:1`.

After creating multilink bundles, you add constituent links to the bundle. The constituent links are the low-speed physical links that are to be aggregated. You can create 64 multilink bundles on a Services Router, and on each multilink bundle you can add up to 8 constituent links. The following rules apply when you add constituent links to a multilink bundle:

- On each multilink bundle, add only interfaces of the same type. For example, you can add either T1 or E1, but not both.

- Only interfaces with a PPP encapsulation can be added to an MLPPP bundle, and only interfaces with a Frame Relay encapsulation can be added to an MLFR bundle.
- If an interface is a member of an existing bundle and you add it to a new bundle, the interface is automatically deleted from the existing bundle and added to the new bundle.

For information about configuring MLPPP bundles, see “Configuring an MLPPP Bundle” on page 274. For information about configuring MLFR bundles, see “Configuring MLFR FRF.15 Bundles” on page 284 and “Configuring MLFR FRF.16 Bundles” on page 287.

Link Fragmentation and Interleaving Overview


As it does on any other interface, priority scheduling on a multilink bundle determines the order in which an output interface transmits traffic from an output queue. The queues are serviced in a weighted round-robin fashion. But when a queue containing large packets starts using the multilink bundle, small and delay-sensitive packets must wait their turn for transmission. Because of this delay, some slow links, such as T1 and E1, can become useless for delay-sensitive traffic.

On a Services Router, link fragmentation and interleaving (LFI) solves this problem. It reduces delay and jitter on links by fragmenting large packets and interleaving delay-sensitive packets with the resulting smaller packets for simultaneous transmission across multiple links of a multilink bundle.

Figure 38 on page 266 illustrates how LFI works on a Services Router. In this figure, Router R0 and Router R1 have LFI enabled. When Router R0 receives large and small packets, such as data and voice packets, it divides them into two categories. All voice packets and any other packets configured to be treated as voice packets, such as CRTP packets, are categorized as LFI packets and transmitted without fragmentation or an MLPPP header. The remaining non-LFI (data) packets can be fragmented or unfragmented based on the configured fragmentation threshold. The packets larger than the fragmentation threshold are fragmented. An MLPPP header (containing a multilink sequence number) is added to all non-LFI packets, fragmented and unfragmented.

The fragmentation is performed according to the fragmentation threshold that you configure. For example, if you configure a fragmentation threshold of 128 bytes, all packets larger than 128 bytes are fragmented. When Router R1 receives the packets, it sends the unfragmented voice packets immediately but buffers the packet fragments until it receives the last fragment for a packet. In this example, when Router R1 receives fragment 5, it reassembles the fragments and transmits the whole packet.

The unfragmented data packets are treated as a single fragment. Thus Router R1 does not buffer the unfragmented data packets and transmits them as it receives them.

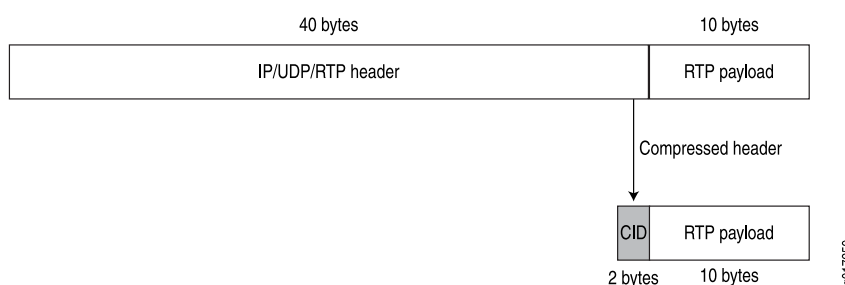
Figure 38: LFI on a Services Router

For information about configuring LFI, see “Enabling Link Fragmentation and Interleaving” on page 276.

Compressed Real-Time Transport Protocol Overview

Real-Time Transport Protocol (RTP) can help achieve interoperability among different implementations of network audio and video applications. However, in some cases, the header, which includes the IP, UDP, and RTP headers, can be too large (around 40 bytes) on networks using low-speed lines such as dial-up modems. Compressed Real-Time Transport Protocol (CRTP) can be configured to reduce network overhead on low-speed links. CRTP replaces the IP, UDP, and RTP headers with a 2-byte context ID (CID), reducing the header overhead considerably.

Figure 39 on page 266 shows how CRTP compresses the RTP headers in a voice packet and reduces a 40-byte header to a 2-byte header.

Figure 39: CRTP

On Services Routers, you can configure CRTP with MLPPP or PPP logical interface encapsulation on link services interfaces. For more information about configuring MLPPP, see “Configuring an MLPPP Bundle” on page 274.

When you configure CRTP, link fragmentation and interleaving (LFI) is automatically enabled. Real-time and non-real-time data frames are carried together on lower-speed links without causing excessive delays to the real-time traffic. For more information about LFI, see “Link Fragmentation and Interleaving Overview” on page 265.

Queuing with LFI on J-series Services Routers

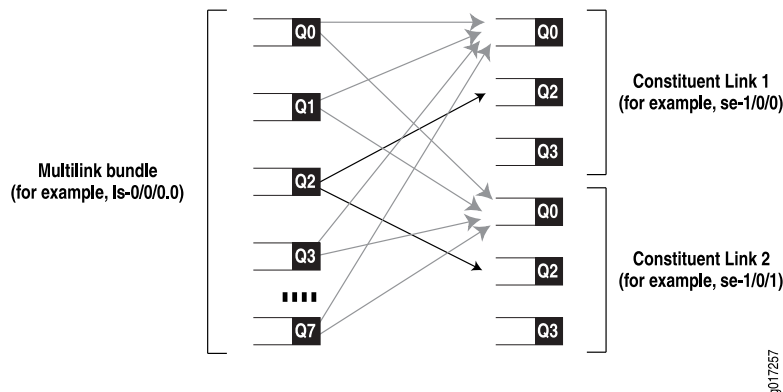
When LFI is enabled, all large packets are fragmented. These packet fragments have a multilink header that contains a multilink sequence number. The sequence numbers on the fragments must be preserved so that the remote router receiving these fragments can correctly reassemble them into a complete packet. To accommodate this requirement, the software queues all fragmented packets on constituent links of a multilink bundle to a single queue (Q0), by default.

Although they are not fragmented, data packets smaller than the fragmentation threshold are also queued to Q0.

When you configure CRTP with LFI, CRTP packets on a multilink bundle from queues other than Q2 are queued to Q2 (instead of Q0) on the constituent links. Because CRTP packets are compressed and do not require fragmentation, they are treated as LFI (voice) packets and are sent to Q2 on the constituent links.

Figure 40 on page 267 shows how traffic is queued on an MLPPP or MLFR multilink bundle and its constituent links. Irrespective of the packet queuing on the multilink bundle, the packets on the constituent links are queued according to the default setting so that traffic from all queues except Q2 is mapped to Q0.

Figure 40: Queuing on Constituent Links




- The packet fragments on Q0, Q1, Q3, Q4, Q5, Q6, and Q7 from the multilink bundle are mapped to Q0 on Constituent Links 1 and 2.
- The LFI packets (such as voice) on Q2 from the multilink bundle are mapped to Q2 on the constituent links.
- The network control packets on Q3 from the multilink bundle are mapped to Q0 on the constituent links. However, Q3 on the constituent links transmits network control packets that exchange protocol information related to constituent links—for example, packets exchanging hello messages on constituent links.

Queuing on Q0s of Constituent Links

On a multilink bundle, packet fragments from all queues except Q2 are transmitted to Q0 on constituent links. On the Q0s of constituent links, the packets are queued in a weighted round-robin fashion to enable per-fragment load balancing.

Figure 41 on page 268 shows how queuing is performed on the constituent links.

Figure 41: Queuing on Q0 of Constituent Links



Packet fragments from the multilink bundle are queued to constituent links one by one in a weighted round-robin fashion. Packet 1 from Q0 on the multilink bundle is queued to Q0 on Constituent Link 1, packet 2 from Q1 on the multilink bundle is queued to Q0 on Constituent Link 2, packet 3 from Q3 on the multilink bundle is queued to Q0 on Constituent Link 1, and so on.

Queuing on Q2s of Constituent Links

On a multilink bundle, all Q2 traffic (LFI traffic) from the multilink bundle is queued to Q2 of constituent links based on a hash computed from the source address, destination address, and IP protocol of the packet. If the IP payload is TCP or UDP traffic, the hash also includes the source port and destination port. As a result of this hash algorithm, all traffic belonging to one traffic flow is queued to Q2 of one constituent link.

Load Balancing with LFI

On link services interfaces, the traffic load is queued and balanced differently for LFI (voice and CRTP packets) and non-LFI packets (data packets) depending on the protocols configured.

Table 76 on page 269 compares queuing and load balancing for LFI and non-LFI packets when MLPPP is configured with LFI and CRTP.

Table 76: LFI Queuing and Load Balancing for Different Protocols

Packet Type	Queuing (MLPPP with LFI)	Queuing (MLPPP with CRTP)	Load Balancing
LFI (voice and CRTP) packets	All incoming packets on Q2 are treated as LFI packets	<p>The following types of incoming packets are treated as LFI packets:</p> <ul style="list-style-type: none"> ■ Packets matching Q2 (default) ■ Packets from ports configured as LFI ports ■ Packets to queues other than Q2 that are configured as LFI queues <p>NOTE: When CRTP is configured without MLPPP traffic traverses only one link thus no load balancing is performed.</p>	<p>Traffic is divided into individual traffic flows, and packets belonging to a flow traverse a single link to avoid packet-ordering issues.</p> <p>The link is selected based on a hash computed from the source address, destination address, and protocol. If the IP payload is TCP or UDP traffic, the hash also includes the source port and destination port.</p>
Non-LFI (data) packets	<p>All data packets, whether fragmented or not, are treated as non-LFI packets and queued to the Q0s of constituent links.</p> <p>(Packets smaller than the size specified in the fragmentation threshold are not fragmented but are treated as non-LFI packets.)</p>	<p>The following types of packets are treated as non-LFI packets and are queued to the Q0s of constituent links:</p> <ul style="list-style-type: none"> ■ Packets not matching Q2 ■ Packets from ports not configured as LFI ports ■ Packets queued to queues not configured for LFI ■ Packets that are not CRTP packets 	All non-LFI packets are queued to the Q0s of constituent links one by one in weighted round-robin fashion.

Configuring CoS Components with LFI

If you configure CoS components with LFI on a Services Router, we recommend that you follow certain recommendations for shaping rate, scheduling priority, and buffer size. For configuration instructions, see “Configuring MLPPP Bundles and LFI on Serial Links” on page 273. For more information about other CoS components, see the *J-series Services Router Advanced WAN Access Configuration Guide*.

Shaping Rate

When you configure LFI on a Services Router, we recommend that you configure the shaping rate on each constituent link of the multilink bundle. Shaping rate configuration on the constituent links is required to limit the jitter on the LFI queue. If you anticipate no delay-sensitive or jitter-sensitive traffic on the LFI queue, or if there is no LFI traffic at all, shaping rate configuration is optional.

For information about how to configure a shaping rate, see “Applying Shaping Rates to Interfaces” on page 283.

Scheduling Priority

Services Routers support per-unit scheduling that allows you to configure scheduler maps on each MLPPP or MLFR multilink bundle. You can also configure scheduler maps on constituent links, but you must maintain the same relative priority on the constituent links and on the multilink bundle.

Table 77 on page 270 shows an example of correct and incorrect relative priorities on a multilink bundle and its constituent link. In this example, you have assigned a high priority to LFI packets and a low priority to data packets on the multilink bundle. To maintain the relative priority on the constituent links, you can assign a high priority to the LFI packets and a medium-high priority to the data packets, but you cannot assign a medium-high priority to LFI packets and a high priority to data packets.

Table 77: Relative Priorities on Multilink Bundles and Constituent Links

Multilink Bundle	Correct Constituent Link Priorities	Incorrect Constituent Link Priorities
LFI packets—High priority	LFI packets—High priority	LFI packet—Medium-high priority
Data packets—Low priority	Data packets—Medium-high priority	Data packets—High priority

Buffer Size

All non-LFI traffic from the multilink bundle (from different queues) is transmitted to Q0 on the constituent links. On the constituent links, you must configure a large buffer size for Q0. If the Q0 buffer size on a constituent link is insufficient, the scheduler might drop overflowing packets.

Before You Begin

Before you configure a link services interface, you need to perform the following tasks:

- Install Services Router hardware. For more information, see the Getting Started Guide for your router.
- Establish basic connectivity. For more information, see the Getting Started Guide for your router.
- If you do not already have a basic understanding of physical and logical interfaces and Juniper Networks interface conventions, read “Interfaces Overview” on page 41.

Although it is not a requirement, you might also want to plan how you are going to use the link services interface on your network before you begin configuring it. Read “Link Services Interfaces Overview” on page 262 for a basic understanding of the link services interface implementation.

Configuring the Link Services Interface with Quick Configuration

You can use the services interfaces Quick Configuration pages to do the following:

- Configure the **Is-0/0/0** link services interface.
- Configure multilink logical interfaces on the **Is-0/0/0** interface. Multilink logical interfaces allow you to bundle multiple serial interfaces such as T1, T3, E1, E3, and serial interfaces into a single logical link as follows:
 - Bundle multiple Point-to-Point Protocol (PPP) links into a single Multilink Point-to-Point Protocol (MLPPP) logical link.
 - Bundle multiple Frame Relay data-link connection identifiers (DLCIs) into a single Multilink Frame Relay (MLFR) logical link.

To configure the link services interface:

1. From the Quick Configuration page, as shown in Figure 22 on page 104, select the link services interface—for example, **Is-0/0/0**—you want to configure.
2. Enter information into the Quick Configuration page, as described in Table 78 on page 271.
3. Click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
 - To apply the configuration and return to the main configuration page, click **OK**.
 - To cancel your entries and return to the main page, click **Cancel**.

Table 78: Link Services Interface Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
Add logical interfaces	Defines one or more logical units that you connect to this link services interface. You must define at least one logical unit for the link services interface.	Click Add .
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	<ol style="list-style-type: none"> 1. Type one or more IPv4 addresses and prefixes. For example: 10.10.10.10/24 2. Click Add. 3. Click OK.

Table 78: Link Services Interface Quick Configuration Summary (continued)

Field	Function	Your Action
Physical Interface Description	(Optional) Adds supplementary information about the physical link services interface.	Type a text description of the link services interface to more clearly identify it in monitoring displays.
Enable subunit queuing	Enables or disables subunit queuing on Frame Relay or VLAN IQ interfaces.	<ul style="list-style-type: none"> ■ To enable subunit queuing, select the check box. ■ To disable subunit queuing, clear the check box.
Multilink Bundle Options		
Bandwidth	Specifies the informational-only bandwidth value for the logical interface.	Type the value.
Drop Timer Period	<p>Specifies a drop timeout value (in milliseconds) to provide a recovery mechanism if individual links in the multilink bundle drop one or more packets.</p> <p>NOTE: Ensure that the value you specify is larger than the expected differential delay across the links, so that the timeout period does not elapse under normal jitter conditions, but only when there is actual packet loss.</p>	Type a value between 0 and 2000.
Encapsulation	Specifies the encapsulation type for which you want to create a multilink bundle.	<p>From the list, select one of the following:</p> <ul style="list-style-type: none"> ■ multilink-ppp—Creates a Multilink Point-to-Point Protocol (MLPPP) bundle. ■ multilink-frame-relay-end-to-end—Creates a Multilink Frame Relay (MLFR) bundle.
Fragmentation Threshold	Specifies the maximum size, in bytes, for multilink packet fragments.	Type a value that is a multiple of 64 bytes between 64 and 16320—for example, 1024.
Links needed to sustain bundle	Specifies the minimum number of links required to sustain the multilink bundle.	Type a value between 1 and 8.
MRRU	Specifies the maximum packet size, in bytes, that the multilink interface can process.	Type a value between 1500 and 4500.
Short Sequence	Sets the length of the packet sequence identification number to 12 bits.	Select this check box.

Table 78: Link Services Interface Quick Configuration Summary *(continued)*

Field	Function	Your Action
Member Interfaces	<p>Specifies the interfaces that are members of the multilink bundle.</p> <p>The Logical Interfaces list displays all the serial interfaces on the router. The Member Interfaces list displays the interfaces that are members of the multilink bundle.</p> <p>The following rules apply when you add interfaces to a multilink bundle:</p> <ul style="list-style-type: none"> ■ Only interfaces of the same type can be added to a multilink bundle. For example, a T1 and an E1 interface cannot be added to the same bundle. ■ Only interfaces with the PPP encapsulation can be added to an MLPPP bundle and interfaces with the Frame Relay encapsulation can be added to an MLFR bundle. ■ If you add an interface that is a member of an existing bundle, the interface is deleted from the existing bundle and added to the new bundle. 	<ul style="list-style-type: none"> ■ To add an interface in the multilink bundle, select the interface in the Logical Interfaces list and click the left arrow button to add it in the Member Interfaces list. ■ To remove an interface from the multilink bundle, select the interface in the Member Interfaces list and click the right arrow button to remove it from the Member Interfaces list.


Configuring the Link Services Interface with a Configuration Editor

This section contains the following topics:

- Configuring MLPPP Bundles and LFI on Serial Links on page 273
- Configuring MLFR FRF.15 Bundles on page 284
- Configuring MLFR FRF.16 Bundles on page 287
- Configuring CRTP on page 289

Configuring MLPPP Bundles and LFI on Serial Links

Figure 42 on page 274 shows a network topology that is used as an example in this section. In this example, your company's branch office is connected to its main branch using J-series Services Routers R0 and R1. You transmit data and voice traffic on two low-speed 1-Mbps serial links. To increase bandwidth, you configure MLPPP and join the two serial links **se-1/0/0** and **se-1/0/1** into a multilink bundle **ls-0/0/0.0**. Then you configure LFI and CoS on R0 and R1 to enable them to transmit voice packets ahead of data packets.

Figure 42: Configuring MLPPP and LFI on Serial Links

Configuring a multilink bundle on the two serial links increases the bandwidth by 70 percent from approximately 1 Mbps to 1.7 Mbps and prepends each packet with a multilink header as specified in the FRF.12 standard. To increase the bandwidth further, you can add up to 8 serial links to the bundle. In addition to a higher bandwidth, configuring the multilink bundle provides load balancing and redundancy. If one of the serial links fails, traffic continues to be transmitted on the other links without any interruption. In contrast, independent links require routing policies for load balancing and redundancy. Independent links also require IP addresses for each link as opposed to one IP address for the bundle. In the routing table, the multilink bundle is represented as a single interface.

This example uses MLPPP for providing multilink services. For information about configuring MLFR, see “Configuring MLFR FRF.15 Bundles” on page 284 and “Configuring MLFR FRF.16 Bundles” on page 287.

You can use the LFI and CoS configurations provided in this example with MLFR FRF.15 and MLFR FRF.16 bundles, too. You can also use the same LFI and CoS configurations for other interfaces, such as on T1 or E1.

To configure MLPPP bundles and LFI on a Services Router, perform the following tasks:

- Configuring an MLPPP Bundle on page 274
- Enabling Link Fragmentation and Interleaving on page 276
- Defining Classifiers and Forwarding Classes on page 277
- Defining and Applying Scheduler Maps on page 279
- Applying Shaping Rates to Interfaces on page 283

Configuring an MLPPP Bundle

In this example, you create an MLPPP bundle (ls-0/0/0.0) at the logical unit level of the link services interface (ls-0/0/0) on routers R0 and R1. Then you add the two serial interfaces **se-1/0/0** and **se-1/0/1** as constituent links to the multilink bundle. Adding multiple links does not require you to configure and manage more addresses.

To configure an MLPPP bundle on a Services Router:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 79 on page 275 on Router R0 and Router R1.

3. Go on to “Enabling Link Fragmentation and Interleaving” on page 276.

Table 79: Configuring an MLPPP Bundle

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy. Specify the link services interface to be configured.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Configure or Edit. 3. Next to Interface, click Add new entry. 4. In the Interface name box, type ls-0/0/0. 5. Click OK. 	From the [edit] hierarchy level, enter edit interfaces ls-0/0/0
Configure a logical unit on the ls-0/0/0 interface and define the family type—for example, Inet . Configure an IP address for the multilink bundle at the unit level of the link services interface.	<ol style="list-style-type: none"> 1. Next to ls-0/0/0, click Edit. 2. Next to Unit, click Add new entry. 3. In the Interface unit number box, type 0. 4. Under Family, select Inet and click Configure. 5. Next to Address, click Add new entry. 6. In the Source box, type the appropriate source address: <ul style="list-style-type: none"> ■ On R0—10.0.0.10/24 ■ On R1—10.0.0.9/24 7. Click OK until you return to the Interfaces page. 	Set the appropriate source address for the interface: <ul style="list-style-type: none"> ■ On R0, enter set unit 0 family inet address 10.0.0.10/24 ■ On R1, enter set unit 0 family inet address 10.0.0.9/24
From the Interfaces level in the configuration hierarchy, specify the names of the constituent links to be added to the multilink bundle—for example, se-1/0/0 and se-1/0/1 .	<ol style="list-style-type: none"> 1. On the Interfaces page, Next to Interface, click Add new entry. 2. In the Interface name box, type the name of the interface to be added to the multilink bundle—for example se-1/0/0 or se-1/0/1. 3. Click OK. 4. Click Edit next to the appropriate interface name—for example, se-1/0/0 or se-1/0/1. 	From the [edit] hierarchy level, add the constituent links to the multilink bundle. <ul style="list-style-type: none"> ■ To add se-1/0/0 to the multilink bundle, enter edit interfaces se-1/0/0 ■ To add se-1/0/1 to the multilink bundle, enter edit interfaces se-1/0/1

Table 79: Configuring an MLPPP Bundle *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the multilink bundle by specifying a logical unit on each constituent link and defining it as an MLPPP bundle—for example, ls-0/0/0.0.	<ol style="list-style-type: none"> Next to Unit, click Add new entry. In the Interface unit number box, type 0. Under Family, select Mlppp and click Configure. In the Bundle box, type ls-0/0/0.0. Click OK until you return to the Interfaces page. 	<pre>Enter set unit 0 family mlppp bundle ls-0/0/0.0</pre>
<p>Set the serial options to the same values for both interfaces on R0—se-1/0/0 and se-1/0/1.</p> <p>For more information about serial options, see “Configuring Serial Interfaces with Quick Configuration” on page 127.</p> <p>NOTE: In this example, R0 is set as a data circuit-terminating equipment (DCE) device. The serial options are not set for interfaces on R1. You can set the serial options according to your network setup.</p>	<ol style="list-style-type: none"> On the Interfaces page, click Edit. Next to the interface that you want to configure (se-1/0/0 or se-1/0/1), click Edit. Next to Serial options, click Configure. From the Clocking mode list, select dce. From the Clock rate list, select 2.0mhz. Click OK twice. 	<ol style="list-style-type: none"> On R0, from the [edit] hierarchy level, set serial options for the interface. <ul style="list-style-type: none"> To set options on se-1/0/0, enter <code>edit interfaces se-1/0/0</code> To set options on se-1/0/1, enter <code>edit interfaces se-1/0/1</code> Enter <code>set serial-options clocking-mode dce clock-rate 2.0mhz</code>

Enabling Link Fragmentation and Interleaving

To configure link fragmentation and interleaving (LFI), you define the MLPPP encapsulation type and enable fragmentation and interleaving of packets by specifying the following properties—the fragmentation threshold and fragment interleaving. In this example, a fragmentation threshold of 128 bytes is set on the MLPPP bundle that applies to all traffic on both constituent links, so that any packet larger than 128 bytes transmitted on these links is fragmented.

For more information about LFI, see “Link Fragmentation and Interleaving Overview” on page 265.

To enable LFI:

- Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
- Perform the configuration tasks described in Table 80 on page 277 on Router R0 and Router R1.
- Go on to “Defining Classifiers and Forwarding Classes” on page 277.

Table 80: Enabling LFI

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration .	From the [edit] hierarchy level, enter
Specify the link services interface for fragmentation.	2. Next to Interfaces, click Edit . 3. Under Interface, next to ls-0/0/0, click Edit .	edit interfaces ls-0/0/0
Specify the multilink encapsulation type, enable LFI, and set the fragmentation threshold for the multilink interface.	1. Under Unit, next to 0, click Edit . 2. From the Encapsulation list, select multilink-ppp as the encapsulation type.	Enter set unit 0 encapsulation multilink-ppp fragment-threshold 128 interleave-fragments
Fragment Threshold—Set the maximum size, in bytes, for multilink packet fragments—for example, 128 . 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).	3. In the Fragment threshold box, type 128 . 4. Select Interleave fragments . 5. Click OK .	
Interleave Fragments—Specify interleaving packet fragments with delay-sensitive (LFI) packets.		

Defining Classifiers and Forwarding Classes

By defining classifiers you associate incoming packets with a forwarding class and loss priority. Based on the associated forwarding class, you assign packets to output queues. To configure classifiers, you specify the bit pattern for the different types of traffic. The classifier takes this bit pattern and attempts to match it to the type of packet arriving on the interface. If the information in the packet's header matches the specified pattern, the packet is sent to the appropriate queue, defined by the forwarding class associated with the classifier.

In this example, an IP precedence classifier, **classify_input**, is assigned to all incoming traffic. The precedence bit value in the type of service (ToS) field is assumed to be **000** for all incoming data traffic and **010** for all incoming voice traffic. This classifier assigns all data traffic to Q0 and all voice traffic to Q2. On a Services Router, when LFI is enabled, all traffic assigned to Q2 is treated as LFI (voice) traffic. You do not need to assign network control traffic to a queue explicitly, because it is assigned to Q3 by default.

For more information about configuring CoS components, see the *J-series Services Router Advanced WAN Access Configuration Guide*.

To define classifiers and forwarding classes:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.

2. Perform the configuration tasks described in Table 81 on page 278 on Router R0 and Router R1.
3. Go on to “Defining and Applying Scheduler Maps” on page 279.

Table 81: Defining Classifiers and Forwarding Classes

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Class of service level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Class of service, click Configure or Edit. 	From the [edit] hierarchy level, enter edit class-of-service
Configure a behavior aggregate (BA) classifier for classifying packets. In this example, you specify the default IP precedence classifier, which maps IP precedence bits to forwarding classes and loss priorities.	<ol style="list-style-type: none"> 1. Next to Classifiers, click Configure. 2. Next to Inet precedence, click Add new entry. 3. In the Name box, type classify_input. 	Enter edit classifiers inet-precedence classify_input
For the classifier to assign an output queue to each packet, it must associate the packet with a forwarding class. Assign packets with IP precedence bits 000 to the DATA forwarding class, and specify a low loss priority.	<ol style="list-style-type: none"> 1. On the Inet precedence page, next to Forwarding class, click Add new entry. 2. In the Class name box, type DATA. 3. Next to Loss priority, click Add new entry. 4. From the Loss val list, select low. 5. Next to Code points, click Add new entry. 6. In the Value box, type 000. 7. Click OK until you return to the Class of service page. 	Enter set forwarding-class DATA loss-priority low code-points 000
Assign packets with IP precedence bits 010 to the VOICE forwarding class, and specify a low loss priority.	<ol style="list-style-type: none"> 1. Next to Forwarding class, click Add new entry. 2. In the Class name box, type VOICE. 3. Next to Loss priority, click Add new entry. 4. From the Loss val list, select low. 5. Next to Code points, click Add new entry. 6. In the Value box, type 010. 7. Click OK until you return to the Inet precedence page. 	Enter set forwarding-class VOICE loss-priority low code-points 010

Table 81: Defining Classifiers and Forwarding Classes (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
Assign each forwarding class one-to-one with the output queues. <ul style="list-style-type: none"> ■ DATA—Assign to Queue 0. ■ VOICE—Assign to Queue 2. ■ NC (Network Control)—Assign to Queue 3. NC is assigned to Queue 3 by default. 	<ol style="list-style-type: none"> 1. On the Class of service page, next to Forwarding classes, click Configure. 2. Next to Queue, click Add new entry. 3. In the Queue num box, type 0. 4. In the Class name box, type DATA. 5. Click OK. 6. Next to Queue, click Add new entry. 7. In the Queue num box, type 2. 8. In the Class name box, type VOICE. 9. Click OK. 10. Next to Queue, click Add new entry. 11. In the Queue num box, type 3. 12. In the Class name box, type NC. 13. Click OK until you return to the Class of service page. 	<p>From the [edit class-of-service] hierarchy level, enter</p> <p>set forwarding-classes queue 0 DATA</p> <p>set forwarding-classes queue 2 VOICE</p> <p>set forwarding-classes queue 3 NC</p>
Apply the behavior aggregate classifier to the incoming interface.	<ol style="list-style-type: none"> 1. On the Class of service page, next to Interfaces, click Add new entry. 2. In the Interface name box, type <code>ge-0/0/1</code>. 3. Next to Unit, click Add new entry. 4. In the Unit number box, type 0. 5. Next to Classifiers, click Configure. 6. Under Inet precedence, in the Classifier name box, type <code>classify_input</code>. 7. Click OK. 	<ol style="list-style-type: none"> 1. From the [edit class-of-service] hierarchy level, enter <code>edit interfaces ge-0/0/1</code> 2. Enter <code>set unit 0 classifiers inet-precedence classify_input</code>

Defining and Applying Scheduler Maps

By defining schedulers you configure the properties of output queues that determine the transmission service level for each queue. These properties include the amount of interface bandwidth assigned to the queue, the size of the memory buffer allocated for storing packets, and the priority of the queue. After defining schedulers you associate them with forwarding classes by means of scheduler maps. You then associate each scheduler map with an interface, thereby configuring the hardware queues and packet schedulers that operate according to this mapping.

In this example, you define and apply scheduler maps as follows:

- Enable per-unit scheduling that allows configuration of scheduler maps on the bundle.
- Create three schedulers—**DATA**, **VOICE**, and **NC**. Define the **VOICE** and **NC** schedulers to have a high priority and the **DATA** scheduler to have the default priority (low). These priority assignments allow all voice and network control traffic to be transmitted ahead of data packets. For more information about scheduling priorities, see “Queuing with LFI on J-series Services Routers” on page 267.
- Create a scheduler map **s_map** that associates these schedulers with corresponding forwarding classes.
- Apply the scheduler map to the multilink bundle and the serial interfaces.

To define and apply scheduler maps:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 82 on page 280 on Router R0 and Router R1.
3. Go on to “Applying Shaping Rates to Interfaces” on page 283.

Table 82: Defining and Applying Scheduler Maps

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interface level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Configure or Edit. 	From the [edit] hierarchy level, enter edit interfaces
To configure CoS components for each multilink bundle, enable per-unit scheduling on the interface.	<ol style="list-style-type: none"> 1. Under Interfaces, select ls-0/0/0. 2. From the Scheduler type list, select Per unit scheduler. 3. Click OK. 4. Under Interfaces, select se-1/0/0. 5. From the Scheduler type list, select Per unit scheduler. 6. Click OK. 7. Under Interfaces, select se-1/0/1. 8. From the Scheduler type list, select Per unit scheduler. 9. Click OK. 	<p>Enter</p> <p>set ls-0/0/0 per-unit-scheduler</p> <p>set se-1/0/0 per-unit-scheduler</p> <p>set se-1/0/1 per-unit-scheduler</p>

Table 82: Defining and Applying Scheduler Maps *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the Class of Service configuration hierarchy and specify the link services interface to be configured.	<ol style="list-style-type: none"> On the Class of service page, next to Interfaces, click Add new entry. In the Interface name box, type <code>ls-0/0/0</code>. 	<p>From the [edit class-of-service] hierarchy level, enter</p> <pre>edit interfaces ls-0/0/0</pre>
Define a scheduler map—for example, <code>s_map</code> .	<ol style="list-style-type: none"> Next to Unit, type Add new entry. In the Unit number box, type 0. In the Scheduler map box, type <code>s_map</code>. Click OK twice. 	<p>Enter</p> <pre>set unit 0 scheduler-map s_map</pre>
Apply the scheduler map to the constituent links of the multilink bundle—for example, <code>se-1/0/0</code> and <code>se-1/0/1</code> .	<ol style="list-style-type: none"> On the Class of service page, next to Interfaces, click Add new entry. In the Interface name box, type the name of the interface on which scheduler map <code>s_map</code> is to be applied—for example, <code>se-1/0/0</code> or <code>se-1/0/1</code>. Next to Unit, type Add new entry. In the Unit number box, type 0. In the Scheduler map box, type <code>s_map</code>. Click OK twice. 	<ol style="list-style-type: none"> From the [edit] hierarchy level, specify the interface to be configured. <ul style="list-style-type: none"> To apply the scheduler map to <code>se-1/0/0</code>, enter <pre>edit interfaces se-1/0/0</pre> To apply the scheduler map to <code>se-1/0/1</code>, enter <pre>edit interfaces se-1/0/1</pre> Apply the scheduler map to the logical interface. <pre>set unit 0 scheduler-map s_map</pre>

Table 82: Defining and Applying Scheduler Maps (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Associate a scheduler with each forwarding class.</p> <ul style="list-style-type: none"> ■ DATA—A scheduler associated with the DATA forwarding class. ■ VOICE—A scheduler associated with the VOICE forwarding class. ■ NC—A scheduler associated with the NC forwarding class. <p>A scheduler receives the forwarding class and loss priority settings, and queues the outgoing packet based on those settings.</p>	<ol style="list-style-type: none"> 1. On the Class of service page, next to Scheduler maps, click Add new entry. 2. In the Map name box, type s_map. 3. Next to Forwarding class, click Add new entry. 4. In the Class name box, type DATA. 5. In the Scheduler box, type DATA. 6. Click OK. 7. Next to Forwarding class, click Add new entry. 8. In the Class name box, type VOICE. 9. In the Scheduler box, type VOICE. 10. Click OK. 11. Next to Forwarding class, click Add new entry. 12. In the Class name box, type NC. 13. In the Scheduler box, type NC. 14. Click OK until you return to the Class of service page. 	<p>From the [edit class-of-service] hierarchy level, enter</p> <p>set scheduler-maps s_map forwarding-class DATA scheduler DATA</p> <p>set scheduler-maps s_map forwarding-class VOICE scheduler VOICE</p> <p>set scheduler-maps s_map forwarding-class NC scheduler NC</p>
<p>Define the properties of output queues for the DATA scheduler:</p> <ul style="list-style-type: none"> ■ Transmit rate—Specify a percentage of transmission capacity—49. ■ Buffer size—Specify a percentage of total buffer—49. ■ Priority—Do not specify the transmission priority for the DATA scheduler to apply the default setting—low. <p>For more information about transmit rate and buffer size, see the <i>J-series Services Router Advanced WAN Access Configuration Guide</i>.</p>	<ol style="list-style-type: none"> 1. On the Class of service page, next to Schedulers, click Add new entry. 2. In the Scheduler name box, type DATA. 3. Next to Transmit rate, click Configure. 4. From the Transmit rate choice list, select Percent. 5. In the Percent box, type 49. 6. Click OK. 7. Next to Buffer size, click Configure. 8. From the Buffer size choice list, select Percent. 9. In the Percent box, type 49. 10. Click OK twice. 	<p>Enter</p> <p>set schedulers DATA transmit-rate percent 49</p> <p>set schedulers DATA buffer-size percent 49</p>

Table 82: Defining and Applying Scheduler Maps *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
Define the properties of output queues for the VOICE scheduler:	1. On the Class of service page, next to Schedulers, click Add new entry .	Enter
■ Transmit rate—Specify a percentage of transmission capacity—50.	2. In the Scheduler name box, type VOICE .	set schedulers VOICE transmit-rate percent 50
■ Buffer size—Specify a percentage of total buffer—5.	3. Next to Transmit rate, click Configure .	set schedulers VOICE buffer-size percent 5
■ Priority—Specify a transmission priority—high.	4. From the Transmit rate choice list, select Percent .	set schedulers VOICE priority high
	5. In the Percent box, type 50.	
	6. Click OK .	
	7. Next to Buffer size, click Configure .	
	8. From the Buffer size choice list, select Percent .	
	9. In the Percent box, type 5.	
	10. Click OK .	
	11. In the Priority box, type high .	
	12. Click OK .	
Define the properties of output queues for the NC scheduler:	1. On the Class of service page, next to Schedulers, click Add new entry .	Enter
■ Transmit rate—Specify a percentage of transmission capacity—1.	2. In the Scheduler name box, type NC .	set schedulers NC transmit-rate percent 1
■ Buffer size—Specify a percentage of total buffer—1.	3. Next to Transmit rate, click Configure .	set schedulers NC buffer-size percent 1
■ Priority—Specify a transmission priority—high.	4. From the Transmit rate choice list, select Percent .	set schedulers NC priority high
	5. In the Percent box, type 1.	
	6. Click OK .	
	7. Next to Buffer size, click Configure .	
	8. From the Buffer size choice list, select Percent .	
	9. In the Percent box, type 1.	
	10. Click OK .	
	11. In the Priority box, type high .	
	12. Click OK .	

Applying Shaping Rates to Interfaces

To control the voice traffic latency within acceptable limits, you configure the shaping rate on constituent links of the MLPPP bundle. Shaping rate at the interface level is

required only when you enable LFI. To apply shaping rates to interfaces, you have to first enable per-unit scheduling. For information about shaping rates and LFI, see “Configuring CoS Components with LFI” on page 269.

You must configure the shaping rate to be equal to the combined physical interface bandwidth for the constituent links. In this example, the combined bandwidth capacity of the two constituent links—**se-1/0/0** and **se-1/0/1**—is 2 Mbps. Hence, configure a shaping rate of 2 Mbps on each constituent link.

To apply a shaping rate to the constituent links of the multilink bundle:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 83 on page 284 on Router R0 and Router R1.
3. Go on to “Verifying the Link Services Interface Configuration” on page 291, to verify your configuration.

Table 83: Applying Shaping Rate to Interfaces

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Class of service level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Class of service, click Edit. 	<p>From the [edit] hierarchy level, enter</p> <p>edit class-of-service</p>
Apply the shaping rate to the constituent links of the multilink bundle—for example, se-1/0/0 and se-1/0/1 . The shaping rate specifies the amount of bandwidth to be allocated for this multilink bundle.	<ol style="list-style-type: none"> 1. Under Interfaces, select the name of the interface on which you want to apply the shaping rate—se-1/0/0 or se-1/0/1. 2. Next to Unit 0, click Edit. 3. Select Shaping rate, and click Configure. 4. From the Shaping rate choice list, select Rate. 5. In the Rate box, type 2000000. 6. Click OK. 	<ol style="list-style-type: none"> 1. Set the shaping rate on both the constituent links: <ul style="list-style-type: none"> ■ To set the shaping rate for se-1/0/0, enter edit interfaces se-1/0/0 ■ To set the shaping rate for se-1/0/1, enter edit interfaces se-1/0/1 2. Set the shaping rate: <p>set unit 0 shaping-rate 2000000</p>

Configuring MLFR FRF.15 Bundles

J-series Services Routers support Multilink Frame Relay end-to-end (MLFR FRF.15) on the link services interface **ls-0/0/0**.

With MLFR FRF.15, multilink bundles are configured as logical units on the link services interface, such as **ls-0/0/0.0**. MLFR FRF.15 bundles combine multiple permanent virtual circuits (PVCs) into one aggregated virtual circuit (AVC). This process provides fragmentation over multiple PVCs on one end and reassembly of

the AVC on the other end. For more information about multilink bundles, see “Multilink Bundles Overview” on page 264.

You can configure LFI and CoS with MLFR in the same way that you configure them with MLPPP. For information about configuring LFI and CoS, see “Configuring MLPPP Bundles and LFI on Serial Links” on page 273.

In this example, you aggregate two T1 links to create an MLFR FRF.15 bundle on two J-series Services Routers—Router R0 and Router R1.

To configure an MLFR FRF.15 bundle:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor on Router R0 and Router R1.
- 2. Perform the configuration tasks described in Table 84 on page 285.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to “Verifying the Link Services Interface Configuration” on page 291, to verify your configuration.

Table 84: Configuring MLFR FRF.15 Bundles

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy. Specify the link services interface as an interface to be configured.	<div>1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration.</div> <div>2. Next to Interfaces, click Configure or Edit.</div> <div>3. Next to Interface, click Add new entry.</div> <div>4. In the Interface name box, type ls-0/0/0.</div> <div>5. Click OK.</div>	From the [edit] hierarchy level, enter edit interfaces ls-0/0/0

Table 84: Configuring MLFR FRF.15 Bundles *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a logical unit on the ls-0/0/0 interface, and define the family type—for example, Inet .	1. On the Interfaces page, next to ls-0/0/0 , click Edit .	Set the appropriate source address for the interface:
Configure an IP address for the multilink bundle on the unit level of the link services interface.	2. Next to Unit, click Add new entry .	■ On R0, enter set unit 0 family inet address 10.0.0.4/24
	3. In the Interface unit number box, type 0.	■ On R1, enter set unit 0 family inet address 10.0.0.5/24
	4. Under Family, select Inet and click Configure .	
	5. Next to Address, click Add new entry .	
	6. In the Source box, type the appropriate source address:	
	■ On R0—10.0.0.4/24	
	■ On R1—10.0.0.5/24	
	7. Click OK until you return to the Interfaces page.	
Define the multilink bundle as an MLFR FRF.15 bundle by specifying the Multilink Frame Relay end-to-end encapsulation type.	1. On the Interfaces page, next to ls-0/0/0 , click Edit .	From the [edit interfaces ls-0/0/0] hierarchy level, enter
	2. Under Unit, next to 0, click Edit .	set unit 0 encapsulation
	3. From the Encapsulation list, select multilink-frame-relay-end-to-end .	multilink-frame-relay-end-to-end
	4. Click OK until you return to the Interfaces page.	
Specify the names of the constituent links to be added to the multilink bundle—for example, t1-2/0/0 and t1-2/0/1 .	1. On the Interfaces page, next to Interface, click Add new entry .	1. From the [edit] hierarchy level, enter
Define the Frame Relay encapsulation type.	2. In the Interface name box, type the name of the interface:	■ For configuring t1-2/0/0 edit interfaces t1-2/0/0
	■ To configure t1-2/0/0 , type t1-2/0/0 .	■ For configuring t1-2/0/1 edit interfaces t1-2/0/1
	■ To configure t1-2/0/1 , type t1-2/0/1 .	2. Enter
	3. Click OK .	set encapsulation frame-relay
	4. Next to the interface you want to configure, click Edit .	
	5. From the Encapsulation list, select frame-relay .	

Table 84: Configuring MLFR FRF.15 Bundles (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define R0 to be a data circuit-terminating equipment (DCE) device. R1 performs as a data terminal equipment (DTE) device, which is the default with Frame Relay encapsulation.	On R0 only, select Dce .	On R0 only, enter set dce
For more information about DCE and DTE, see “Serial Interface Overview” on page 66.		
On the logical unit level of the interface, specify the data-link connection identifier (DLCI). The DLCI field identifies which logical circuit the data travels over. DLCI is a value from 16 through 1022—for example, 100. (Numbers 1 through 15 are reserved for future use.)	1. Next to Unit, click Add new entry . 2. In the Interface unit number box, type 0. 3. In the DlcI box, type 100.	Enter set unit 0 dlci 100 family mlfr-end-to-end bundle ls-0/0/0.0
Specify the multilink bundle to which the interface is to be added as a constituent link—ls-0/0/0.0.	4. Under Family, select mlfr-end-to-end and click Configure . 5. In the Bundle box, type ls-0/0/0.0. 6. Click OK .	

Configuring MLFR FRF.16 Bundles

J-series Services Routers support Multilink Frame Relay (MLFR) user-to-network interface (UNI) network-to-network interface (NNI) (MLFR FRF.16) on the link services interface ls-0/0/0.

MLFR FRF.16 configures multilink bundles as channels on the link services interface, such as ls-0/0/0:0. A multilink bundle carries Frame Relay permanent virtual circuits (PVCs), identified by their data-link connection identifiers (DLCIs). Each DLCI is configured at the logical unit level of the link services interface and is also referred as a logical interface. Packet fragmentation and reassembly occur on each virtual circuit. For more information about multilink bundles, see “Multilink Bundles Overview” on page 264.

You can configure LFI and CoS with MLFR in the same way that you configure them with MLPPP. For information about configuring LFI and CoS, see “Configuring MLPPP Bundles and LFI on Serial Links” on page 273.

In this example, you aggregate two T1 interfaces to create an MLFR FRF.16 bundle on two J-series Services Routers—Router R0 and Router R1.

To configure an MLFR FRF.16 bundle:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor on Router R0 and Router R1.
2. Perform the configuration tasks described in Table 85 on page 288.

3. If you are finished configuring the router, commit the configuration.
4. Go on to “Verifying the Link Services Interface Configuration” on page 291, to verify your configuration.

Table 85: Configuring MLFR FRF.16 Bundles

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Chassis level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Chassis, click Configure or Edit. 	From the [edit] hierarchy level, enter edit chassis
Specify the number of multilink frame relay UNI NNI (FRF.16) bundles to be created on the interface. You can specify a number from 1 through 255.	<ol style="list-style-type: none"> 1. Next to Fpc, click Add new entry. 2. In the Slot box, type 0. 3. Next to Pic, click Add new entry. 4. In the Slot box, type 0. 5. In the Mlfr uni nni bundles box, type 1. 6. Click OK. 	Enter set fpc 0 pic 0 mlfr-uni-nni-bundles 1
Specify the channel to be configured as a multilink bundle.	<ol style="list-style-type: none"> 1. On the main Configuration page, next to Interfaces, click Configure or Edit. 2. Next to Interface, click Add new entry. 3. In the Interface name box, type ls-0/0/0:0. 4. Click OK. 	From the [edit] hierarchy level, enter edit interfaces ls-0/0/0:0
Define the multilink bundle as an MLFR FRF.16 bundle by specifying the Multilink Frame Relay UNI NNI encapsulation type.	<ol style="list-style-type: none"> 1. Next to ls-0/0/0:0, click Edit. 2. From the Encapsulation list, select multilink-frame-relay-uni-nni. 	Enter set encapsulation multilink-frame-relay-uni-nni
Define R0 to be a data circuit-terminating equipment (DCE) device. R1 performs as a data terminal equipment (DTE) device, which is the default with Frame Relay encapsulation.	On R0 only, select Dce .	On R0 only, enter set dce
For more information about DCE and DTE, see “Serial Interface Overview” on page 66		

Table 85: Configuring MLFR FRF.16 Bundles (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a logical unit on the multilink bundle <code>ls-0/0/0:0</code> , and define the family type—for example, <code>Inet</code> .	1. Next to Unit, click Add new entry .	Set the appropriate address for the interface:
Assign a data link connection identifier (DLCI) to the multilink bundle. The DLCI field identifies which logical circuit the data travels over. DLCI is a value from 16 through 1022—for example, 400. (Numbers 1 through 15 are reserved for future use.)	2. In the Interface unit number box, type 0.	■ On R0, enter set unit 0 dcli 400 family inet address 10.0.0.10/24
Assign an IP address to the multilink bundle.	3. In the Dcli box, type 400.	■ On R1, enter set unit 0 dcli 400 family inet address 10.0.0.9/24
	4. Under Family, select Inet and click Configure .	
	5. Next to Address, click Add new entry .	
	6. In the Source box, type the appropriate source address:	
	■ On R0—10.0.0.10/24	
	■ On R1—10.0.0.9/24	
	7. Click OK until you return to the Interfaces page.	
Create the T1 interfaces that are to be added as constituent links to the multilink bundle— <code>t1-2/0/0</code> and <code>t1-2/0/1</code> .	1. On the Interfaces page, next to Interface, click Add new entry .	1. From the [edit] hierarchy level, enter
Define the Frame Relay encapsulation type.	2. In the Interface name box, type the name of the interface:	■ For configuring <code>t1-2/0/0</code> edit interfaces t1-2/0/0
	■ To configure <code>t1-2/0/0</code> , type <code>t1-2/0/0</code> .	■ For configuring <code>t1-2/0/1</code> edit interfaces t1-2/0/1
	■ To configure <code>t1-2/0/1</code> , type <code>t1-2/0/1</code> .	2. Enter
	3. Click OK .	set encapsulation multilink-frame-relay-uni-nni
	4. Next to the interface you want to configure, click Edit .	
	5. From the Encapsulation list, select multilink-frame-relay-uni-nni .	
Specify the multilink bundle to which the interface is to be added as a constituent link— <code>ls-0/0/0:0</code> .	1. Next to Unit, click Add new entry .	Enter
	2. In the Interface unit number box, type 0.	set unit 0 family mlfr-uni-nni bundle ls-0/0/0:0
	3. Under Family, select mlfr-uni-nni and click Configure .	
	4. In the Bundle box, type <code>ls-0/0/0:0</code> .	
	5. Click OK .	

Configuring CRTP

Compressed Real-Time Transport Protocol (CRTP) is typically used for compressing voice and video packets. You can configure CRTP with LFI on the link services interface of a Services Router.

On the Services Router, CRTP can be configured as a compression device on a T1 or E1 interface with PPP encapsulation, using the link services interface.

For more information about configuring CRTP on a single link, see the *JUNOS Network Interfaces Configuration Guide* and the *JUNOS Services Interfaces Configuration Guide*.

To configure CRTP on the Services Router:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 86 on page 290.
3. If you are finished configuring the router, commit the configuration.

Table 86: Adding CRTP to an T1 or E1 Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	<p>From the [edit] hierarchy level, enter</p> <pre>edit interfaces interface-name</pre>
Select an E1 or T1 interface—for example, t1-1/0/0 .	<ol style="list-style-type: none"> 1. Next to a T1 or E1 interface, click Edit. 2. From the Encapsulation list, select ppp as the encapsulation type. 	<ol style="list-style-type: none"> 1. Enter <pre>set encapsulation ppp</pre>
Set PPP as the type of encapsulation for the physical interface.	<ol style="list-style-type: none"> 3. Next to Unit, click Add new entry. 4. In the Interface unit number box, type 0. 	<ol style="list-style-type: none"> 2. Enter <pre>edit unit 0</pre>
Add the link services interface, ls-0/0/0.0 , to the physical interface.	<ol style="list-style-type: none"> 1. In the Compression device box, enter ls-0/0/0.0. 2. Click OK until you return to the Interfaces page. 	<p>Enter</p> <pre>set compression-device ls-0/0/0.0</pre>
Add the link services interface, ls-0/0/0 , to the Services Router.	<ol style="list-style-type: none"> 1. Next to Interface, click Add new entry. 2. In the Interface name box, type ls-0/0/0. 3. Click OK to return to the Interfaces page. 4. On the main Interface page, next to ls-0/0/0, click Edit. 5. Next to Unit, click Add new entry. 6. In the Interface unit number box, type 0. 	<p>From the [edit interfaces] hierarchy level, enter</p> <pre>edit interfaces ls-0/0/0 unit 0</pre>

Table 86: Adding CRTP to an T1 or E1 Interface (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the link services interface, ls-0/0/0, properties.	1. Next to Compression, select yes , and then click Configure .	Enter
F-max period —Maximum number of compressed packets allowed between transmission of full headers. It has a range from 1 to 65535.	2. Select RTP , and then click Configure .	set compression rtp
	3. In the F-Max period box, type 2500.	f-max-period 2500 port
	4. Select Port, then click Configure .	minimum 2000 maximum
	5. In the Minimum value box, type 2000.	64009
Maximum and Minimum —UDP port values from 1 to 65536 reserve these ports for RTP compression. CRTP is applied to network traffic on ports within this range. This feature is applicable only to voice services interfaces.	6. In the Maximum value box, type 64009.	
	7. Click OK .	

Verifying the Link Services Interface Configuration

To verify a link services configuration, perform the following tasks:

- Displaying Multilink Bundle Configurations on page 292
- Displaying Link Services CoS Configurations on page 294
- Verifying Link Services Interface Statistics on page 296
- Verifying Link Services CoS on page 300

Displaying Multilink Bundle Configurations

Purpose Verify the multilink bundle configuration.

Action From the J-Web interface, select **Configuration > View and Edit > View Configuration Text**. Alternatively, from configuration mode in the CLI, enter the `show interfaces` command.

The sample output in this section displays the multilink bundle configurations provided in “Configuring MLPPP Bundles and LFI on Serial Links” on page 273.



NOTE: The MLFR FRF.15 and MLFR FRF.16 configurations are not displayed in this section, but you can display MLFR configurations in the same manner.

```
[edit]
user@R0# show interfaces
interfaces {
  ls-0/0/0 {
    per-unit-scheduler;
    unit 0 {
      encapsulation multilink-ppp;
      fragment-threshold 128;
      interleave-fragments;
      family inet {
        address 10.0.0.10/24;
      }
    }
  }
  ge-0/0/1 {
    unit 0 {
      family inet {
        address 192.1.1.1/24;
      }
    }
  }
  se-1/0/0 {
    per-unit-scheduler;
    dce-options {
      clocking-mode dce;
      clocking-rate 2.0mhz;
    }
    unit 0 {
      family mlppp {
        bundle ls-0/0/0.0;
      }
    }
  }
  se-1/0/1 {
    per-unit-scheduler;
    dce-options {
      clocking-mode dce;
      clocking-rate 2.0mhz;
    }
  }
}
```

```
    unit 0 {  
      family mlppp {  
        bundle ls-0/0/0.0;  
      }  
    }  
  }  
}
```

What It Means Verify that the output shows the intended multilink bundle configurations.

See Also For more information about the format of a configuration file, see “Viewing the Configuration Text” on page 9.

Displaying Link Services CoS Configurations

Purpose Displaying the CoS configurations on the link services interface.

Action From the J-Web interface, select **Configuration > View and Edit > View Configuration Text**. Alternatively, from the configuration mode in the CLI, enter the `show class-of-service` command.

The sample output in this section displays the CoS configurations provided in “Configuring MLPPP Bundles and LFI on Serial Links” on page 273.

```
[edit]
user@R0# show class-of-service
classifiers {
  inet-precedence classify_input {
    forwarding-class DATA {
      loss-priority low code-points 000;
    }
    forwarding-class VOICE {
      loss-priority low code-points 010;
    }
  }
}
forwarding-classes {
  queue 0 DATA;
  queue 2 VOICE;
  queue 3 NC;
}
interfaces {
  ls-0/0/0 {
    unit 0 {
      scheduler-map s_map;
    }
  }
  ge-0/0/1 {
    unit 0 {
      classifiers {
        inet-precedence classify_input
      }
    }
  }
  se-1/0/0 {
    unit 0 {
      scheduler-map s_map;
      shaping-rate 2000000;
    }
  }
  se-1/0/1 {
    unit 0 {
      scheduler-map s_map;
      shaping-rate 2000000;
    }
  }
}
scheduler-maps {
```



```

s_map {
    forwarding-class DATA scheduler DATA;
    forwarding-class VOICE scheduler VOICE;
    forwarding-class NC scheduler NC;
}
}
schedulers {
    DATA {
        transmit-rate percent 49;
        buffer-size percent 49;
    }
    VOICE {
        transmit-rate percent 50;
        buffer-size percent 5;
        priority high;
    }
    NC {
        transmit-rate percent 1;
        buffer-size percent 1;
        priority high;
    }
}
}

```

What It Means Verify that the output shows the intended CoS configurations.

See Also For more information about the format of a configuration file, see “Viewing the Configuration Text” on page 9.

Verifying Link Services Interface Statistics

Purpose Verify the link services interface statistics.

Action The sample output provided in this section is based on the configurations provided in “Configuring MLPPP Bundles and LFI on Serial Links” on page 273. To verify that the constituent links are added to the bundle correctly and the packets are fragmented and transmitted correctly, take the following actions:

1. On Router R0 and Router R1, the two J-series routers used in this example, configure MLPPP and LFI as described in “Configuring MLPPP Bundles and LFI on Serial Links” on page 273.
2. From the CLI, enter the **ping** command to verify that a connection is established between R0 and R1.
3. Transmit 10 data packets, 200 bytes each, from R0 to R1.
4. On R0, from the CLI, enter the **show interfaces *interface-name* statistics** command.

```
user@R0> show interfaces ls-0/0/0 statistics detail
Physical interface: ls-0/0/0, Enabled, Physical link is Up
Interface index: 134, SNMP ifIndex: 29, Generation: 135
Link-level type: LinkService, MTU: 1504
Device flags   : Present Running
Interface flags: Point-To-Point SNMP-Traps
Last flapped   : 2006-06-23 11:36:23 PDT (03:38:43 ago)
Statistics last cleared: 2006-06-23 15:13:12 PDT (00:01:54 ago)
Traffic statistics:
Input bytes   :                0                0 bps
Output bytes  :             1820                0 bps
Input packets:                0                0 pps
Output packets:             10                0 pps
...
Egress queues: 8 supported, 8 in use
Queue counters:      Queued packets  Transmitted packets      Dropped packets

  0 DATA                10                10                0

  1 expedited-fo          0                0                0

  2 VOICE                 0                0                0

  3 NC                   0                0                0

Logical interface ls-0/0/0.0 (Index 67) (SNMP ifIndex 41) (Generation 133)
Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
Bandwidth: 16mbps
Bundle options:
....
Drop timer period                0
Sequence number format           long (24 bits)
Fragmentation threshold          128
Links needed to sustain bundle   1
Interleave fragments             Enabled
Bundle errors:
Packet drops                     0 (0 bytes)
Fragment drops                   0 (0 bytes)
```

```
...
Statistics
Bundle:
  Fragments:
    Input :      0      0      0      0
    Output:     20      0    1920      0
  Packets:
    Input :      0      0      0      0
    Output:     10      0    1820      0
Link:
  se-1/0/0.0
    Input :      0      0      0      0
    Output:     10      0    1320      0
  se-1/0/1.0
    Input :      0      0      0      0
    Output:     10      0     600      0
...
Destination: 10.0.0.9/24, Local: 10.0.0.10, Broadcast: Unspecified,
Generation:144
```

What It Means This output shows a summary of interface information. Verify the following information:

- **Physical interface**—The physical interface is **Enabled**. If the interface is shown as **Disabled**, do either of the following:
 - In the CLI configuration editor, delete the **disable** statement at the **[edit interfaces *interface-name*]** level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the **Interfaces > *interface-name*** page.
- **Physical link**—The physical link is **Up**. A link state of **Down** indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- **Last flapped**—The **Last Flapped** time is an expected value. The **Last Flapped** time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- **Traffic statistics**—Number and rate of bytes and packets received and transmitted on the interface. Verify that the number of inbound and outbound bytes and packets match the expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics *interface-name*** command.
- **Queue counters**—Name and number of queues are as configured. This sample output shows that 10 data packets were transmitted and no packets were dropped.
- **Logical interface**—Name of the multilink bundle you configured—**ls-0/0/0.0**.
- **Bundle options**—Fragmentation threshold is correctly configured, and fragment interleaving is enabled.
- **Bundle errors**—Any packets and fragments dropped by the bundle.
- **Statistics**—The fragments and packets are received and transmitted correctly by the router. All references to traffic direction (input or output) are defined with respect to the router. Input fragments received by the router are assembled into input packets. Output packets are segmented into output fragments for transmission out of the router.

In this example, 10 data packets of 200 bytes were transmitted. Because the fragmentation threshold is set to 128 bytes, all data packets were fragmented into two fragments. The sample output shows that 10 packets and 20 fragments were transmitted correctly.

- **Link**—The constituent links are added to this bundle and are receiving and transmitting fragments and packets correctly. The combined number of fragments transmitted on the constituent links must be equal to the number of fragments transmitted from the bundle. This sample output shows that the bundle transmitted 20 fragments and the two constituent links **se-1/0/0.0** and **se-1/0/1.0.0** correctly transmitted **10+10=20** fragments.
- **Destination and Local**—IP address of the remote side of the multilink bundle and the local side of the multilink bundle. This sample output shows that the destination address is the address on R1 and the local address is the address on R0.

See Also For a complete description of `show interfaces` output, see the *JUNOS Interfaces Command Reference*.

Verifying Link Services CoS

Purpose Verify CoS configurations on the link services interface.

Action From the CLI, enter the following commands:

- `show class-of-service interface interface-name`
- `show class-of-service classifier name classifier-name`
- `show class-of-service scheduler-map scheduler-map-name`

The sample output provided in this section is based on the configurations provided in “Configuring MLPPP Bundles and LFI on Serial Links” on page 273.

```
user@R0> show class-of-service interface ls-0/0/0
```

```
Physical interface: ls-0/0/0, Index: 136
```

```
Queues supported: 8, Queues in use: 4
```

```
Scheduler map: [default], Index: 2
```

```
Input scheduler map: [default], Index: 3
```

```
Chassis scheduler map: [default-chassis], Index: 4
```

```
Logical interface: ls-0/0/0.0, Index: 69
```

Object	Name	Type	Index
Scheduler-map	s_map	Output	16206
Classifier	ipprec-compatibility	ip	12

```
user@R0> show class-of-service interface ge-0/0/1
```

```
Physical interface: ge-0/0/1, Index: 140
```

```
Queues supported: 8, Queues in use: 4
```

```
Scheduler map: [default], Index: 2
```

```
Input scheduler map: [default], Index: 3
```

```
Logical interface: ge-0/0/1.0, Index: 68
```

Object	Name	Type	Index
Classifier	classify_input	ip	4330

```
user@R0> show class-of-service classifier name classify_input
```

```
Classifier: classify_input, Code point type: inet-precedence, Index: 4330
```

Code point	Forwarding class	Loss priority
000	DATA	low
010	VOICE	low

```
user@R0> show class-of-service scheduler-map s_map
```

```
Scheduler map: s_map, Index: 16206
```

```
Scheduler: DATA, Forwarding class: DATA, Index: 3810
```

```
Transmit rate: 49 percent, Rate Limit: none, Buffer size: 49 percent,
```

```
Priority:low
```

```
Drop profiles:
```

Loss priority	Protocol	Index	Name
Low	any	1	[default-drop-profile]
Medium low	any	1	[default-drop-profile]
Medium high	any	1	[default-drop-profile]
High	any	1	[default-drop-profile]

Scheduler: VOICE, Forwarding class: VOICE, Index: 43363
 Transmit rate: 50 percent, Rate Limit: none, Buffer size: 5 percent,
 Priority:high

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	1	[default-drop-profile]
Medium low	any	1	[default-drop-profile]
Medium high	any	1	[default-drop-profile]
High	any	1	[default-drop-profile]

Scheduler: NC, Forwarding class: NC, Index: 2435

Transmit rate: 1 percent, Rate Limit: none, Buffer size: 1 percent, Priority:high

Drop profiles:

Loss priority	Protocol	Index	Name
Low	any	1	[default-drop-profile]
Medium low	any	1	[default-drop-profile]
Medium high	any	1	[default-drop-profile]
High	any	1	[default-drop-profile]

What It Means These output examples show a summary of configured CoS components. Verify the following information:

- **Logical Interface**—Name of the multilink bundle and the CoS components applied to the bundle. The sample output shows that the multilink bundle is `ls-0/0/0.0`, and the CoS scheduler-map `s_map` is applied to it.
- **Classifier**—Code points, forwarding classes, and loss priorities assigned to the classifier. The sample output shows that a default classifier, `ipprec-compatibility`, was applied to the `ls-0/0/0` interface and the classifier `classify_input` was applied to the `ge-0/0/1` interface.
- **Scheduler**—Transmit rate, buffer size, priority, and loss priority assigned to each scheduler. The sample output displays the data, voice, and network control schedulers with all the configured values.

See Also For complete descriptions of `show class-of-service` commands and output, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Frequently Asked Questions About the Link Services Interface

Use answers to the following questions to solve configuration problems on a link services interface:

- Which CoS Components Are Applied to the Constituent Links? on page 302
- What Causes Jitter and Latency on the Multilink Bundle? on page 303

- Are LFI and Load Balancing Working Correctly? on page 304
- Why Are Packets Dropped on a PVC Between a J-series Router and Another Vendor? on page 310

Which CoS Components Are Applied to the Constituent Links?

Problem—I have configured a multilink bundle, but I also have traffic without MLPPP encapsulation passing through constituent links of the multilink bundle. Do I apply all CoS components to the constituent links, or is applying them to the multilink bundle enough?

Solution—On a J-series Services Router you can apply a scheduler map to the multilink bundle and its constituent links. Although you can apply several CoS components with the scheduler map, configure only the ones that are required. We recommend that you keep the configuration on the constituent links simple to avoid unnecessary delay in transmission.

Table 87 on page 302 shows the CoS components to be applied on a multilink bundle and its constituent links. For more information, see the *JUNOS Class of Service Configuration Guide*.

Table 87: CoS Components Applied on Multilink Bundles and Constituent Links

Cos Component	Multilink Bundle	Constituent Links	Explanation
Classifier	Yes	No	CoS classification takes place on the incoming side of the interface, not on the transmitting side, so no classifiers are needed on constituent links.
Forwarding class	Yes	No	Forwarding class is associated with a queue, and the queue is applied to the interface by a scheduler map. The queue assignment is predetermined on the constituent links. All packets from Q2 of the multilink bundle are assigned to Q2 of the constituent link, and packets from all the other queues are queued to Q0 of the constituent link.

Table 87: CoS Components Applied on Multilink Bundles and Constituent Links (*continued*)

Cos Component	Multilink Bundle	Constituent Links	Explanation
Scheduler map	Yes	Yes	<p>Apply scheduler maps on the multilink bundle and the constituent links, as follows:</p> <ul style="list-style-type: none"> ■ Transmit rate—Make sure that the relative order of the transmit rate configured on Q0 and Q2 is the same on the constituent links as on the multilink bundle. ■ Scheduler priority—Make sure that the relative order of the scheduler priority configured on Q0 and Q2 is the same on the constituent links as on the multilink bundle. ■ Buffer size—Because all non-LFI packets from the multilink bundle transit Q0 of constituent links, make sure that the buffer size on Q0 of the constituent links is large enough. ■ RED drop profile—Configure a RED drop profile on the multilink bundle only. Configuring the RED drop profile on the constituent links applies a back pressure mechanism that changes the buffer size and introduces variation. Because this behavior might cause fragment drops on the constituent links, make sure to leave the RED drop profile at the default settings on the constituent links.
Shaping rate for a per-unit scheduler or an interface-level scheduler	No	Yes	Because per-unit scheduling is applied only at the end point, apply this shaping rate to the constituent links only. Any configuration applied earlier is overwritten by the constituent link configuration.
Transmit-rate exact or queue-level shaping	Yes	No	The interface-level shaping applied on the constituent links overrides any shaping on the queue. Thus apply transmit-rate exact shaping on the multilink bundle only.
Rewrite rules	Yes	No	Rewrite bits are copied from the packet into the fragments automatically during fragmentation. Thus what you configure on the multilink bundle is carried on the fragments to the constituent links.
Virtual channel group	Yes	No	Virtual channel groups are identified through firewall filter rules that are applied on packets only before the multilink bundle. Thus you do not need to apply the virtual channel group configuration to the constituent links.

What Causes Jitter and Latency on the Multilink Bundle?

Problem—To test jitter and latency on a J-series Services Router, I sent three streams of IP packets. All packets have the same IP precedence settings. After I configured LFI and CRTP, the latency increased even over a non-congested link. How can I reduce jitter and latency?

Solution—To reduce jitter and latency do the following:

1. Make sure that you have configured a shaping rate on each constituent link.
2. Make sure that you have not configured a shaping rate on the link services interface.
3. Make sure that the configured shaping rate value is equal to the physical interface bandwidth. For more information, see “Applying Shaping Rates to Interfaces” on page 283.
4. If shaping rates are configured correctly, and jitter still persists, contact the Juniper Networks Technical Assistance Center (JTAC). (See “Requesting Support” on page xxv.)

Are LFI and Load Balancing Working Correctly?

Problem—I have a single network that supports multiple services. My network transmits data and delay-sensitive voice traffic. I configured MLPPP and LFI to make sure that voice packets are transmitted across the network with very little delay and jitter. How can I find out if voice packets are being treated as LFI packets and load balancing is performed correctly?

Solution—When LFI is enabled, data (non-LFI) packets are encapsulated with an MLPPP header and fragmented to packets of a specified size. The delay-sensitive, voice (LFI) packets are PPP-encapsulated and interleaved between data packet fragments. Queuing and load balancing are performed differently for LFI and non-LFI packets. For more information, see “Load Balancing with LFI” on page 268.

To verify that LFI is performed correctly, determine that packets are fragmented and encapsulated as configured. After you know whether a packet is treated as an LFI packet or a non-LFI packet, you can confirm whether the load balancing is performed correctly.

Solution Scenario—Suppose two Services Routers R0 and R1 are connected by a multilink bundle `ls-0/0/0.0` that aggregates two serial links, `se-1/0/0` and `se-1/0/1`. On R0 and R1, MLPPP and LFI are enabled on the link services interface and the fragmentation threshold is set to 128 bytes.

In this example, we used a packet generator to generate voice and data streams. You can use the packet capture feature to capture and analyze the packets on the incoming interface. For more information, see the *J-series Services Router Administration Guide*.

The following two data streams were sent on the multilink bundle:

- 100 data packets of 200 bytes (larger than the fragmentation threshold)
- 500 data packets of 60 bytes (smaller than the fragmentation threshold)

The following two voice streams were sent on the multilink bundle:

- 100 voice packets of 200 bytes from source port 100
- 300 voice packets of 200 bytes from source port 200

To confirm that LFI and load balancing are performed correctly, first verify that the link services interface is performing packet fragmentation as configured. Second, verify that the interface is encapsulating packets as configured. Finally, use the results to verify load balancing.



NOTE: Only the significant portions of command output are displayed and described in this example. For more information, see “Verifying the Link Services Interface Configuration” on page 291.

Step 1: Verifying Packet Fragmentation

From the CLI, enter the `show interfaces ls-0/0/0` command, to check that large packets are fragmented correctly.

```
user@R0#> show interfaces ls-0/0/0
Physical interface: ls-0/0/0, Enabled, Physical link is Up
  Interface index: 136, SNMP ifIndex: 29
  Link-level type: LinkService, MTU: 1504
  Device flags   : Present Running
  Interface flags: Point-To-Point SNMP-Traps
  Last flapped   : 2006-08-01 10:45:13 PDT (2w0d 06:06 ago)
  Input rate     : 0 bps (0 pps)
  Output rate    : 0 bps (0 pps)

Logical interface ls-0/0/0.0 (Index 69) (SNMP ifIndex 42)
  Flags: Point-To-Point SNMP-Traps 0x4000 Encapsulation: Multilink-PPP
  Bandwidth: 16mbps
  Statistics          Frames      fps      Bytes      bps
  Bundle:
    Fragments:
      Input  :           0         0          0         0
      Output:        1100         0     118800         0
    Packets:
      Input  :           0         0          0         0
      Output:        1000         0     112000         0
  ...
  Protocol inet, MTU: 1500
  Flags: None
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 9.9.9/24, Local: 9.9.9.10
```

What It Means—The output shows a summary of packets transiting the router on the multilink bundle. Verify the following information on the multilink bundle:

- The total number of transiting packets = 1000
- The total number of transiting fragments = 1100
- The number of data packets that were fragmented = 100

The total number of packets sent (600 + 400) on the multilink bundle match the number of transiting packets (1000), indicating that no packets were dropped.

The number of transiting fragments exceeds the number of transiting packets by 100, indicating that 100 large data packets were correctly fragmented.

Corrective Action—If the packets are not fragmented correctly, check your fragmentation threshold configuration. Packets smaller than the specified fragmentation threshold are not fragmented. For information about configuring the fragmentation threshold, see “Configuring the Link Services Interface with a Configuration Editor” on page 273.

Step 2: Verifying Packet Encapsulation

To find out whether a packet is treated as an LFI or non-LFI packet, determine its encapsulation type. LFI packets are PPP encapsulated, and non-LFI packets are encapsulated with both PPP and MLPPP. PPP and MLPPP encapsulations have different overheads resulting in different-sized packets. You can compare packet sizes to determine the encapsulation type.

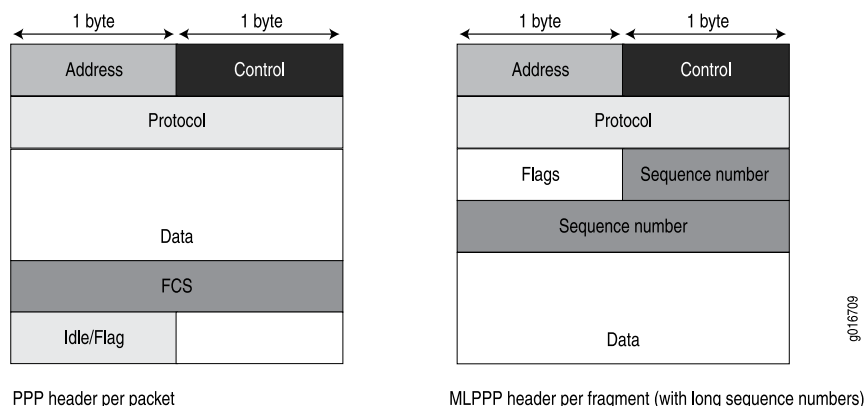
A small unfragmented data packet contains a PPP header and a single MLPPP header. In a large fragmented data packet, the first fragment contains a PPP header and an MLPPP header, but the consecutive fragments contain only an MLPPP header.

PPP and MLPPP encapsulations add the following number of bytes to a packet:

- PPP encapsulation adds 7 bytes:
 - 4 bytes of header + 2 bytes of frame check sequence (FCS) + 1 byte that is idle or contains a flag
- MLPPP encapsulation adds between 6 and 8 bytes:
 - 4 bytes of PPP header + 2 to 4 bytes of multilink header

Figure 43 on page 306 shows the overhead added to PPP and MLPPP headers.

Figure 43: PPP and MLPPP Headers



For CRTP packets, the encapsulation overhead and packet size are even smaller than for an LFI packet. For more information, see “Configuring CRTP” on page 289.

Table 88 on page 307 shows the encapsulation overhead for a data packet and a voice packet of 70 bytes each. After encapsulation, the size of the data packet is larger than the size of the voice packet.

Table 88: PPP and MLPPP Encapsulation Overhead

Packet Type	Encapsulation	Initial Packet Size	Encapsulation Overhead	Packet Size after Encapsulation
Voice packet (LFI)	PPP	70 bytes	$4 + 2 + 1 = 7$ bytes	77 bytes
Data fragment (non-LFI) with short sequence	MLPPP	70 bytes	$4 + 2 + 1 + 4 + 2 = 13$ bytes	83 bytes
Data fragment (non-LFI) with long sequence	MLPPP	70 bytes	$4 + 2 + 1 + 4 + 4 = 15$ bytes	85 bytes

From the CLI, enter the `show interfaces queue` command to display the size of transmitted packet on each queue. Divide the number of bytes transmitted by the number of packets to obtain the size of the packets and determine the encapsulation type.

Step 3: Verifying Load Balancing

From the CLI, enter the `show interfaces queue` command on the multilink bundle and its constituent links to confirm whether load balancing is performed accordingly on the packets.

```

user@R0> show interfaces queue ls-0/0/0
Physical interface: ls-0/0/0, Enabled, Physical link is Up
  Interface index: 136, SNMP ifIndex: 29
Forwarding classes: 8 supported, 8 in use
Egress queues: 8 supported, 8 in use
Queue: 0, Forwarding classes: DATA
  Queued:
    Packets      :           600          0 pps
    Bytes        :          44800          0 bps
  Transmitted:
    Packets      :           600          0 pps
    Bytes        :          44800          0 bps
    Tail-dropped packets :           0          0 pps
    RED-dropped packets  :           0          0 pps
  ...
Queue: 1, Forwarding classes: expedited-forwarding
  Queued:
    Packets      :              0          0 pps
    Bytes        :              0          0 bps
  ...
Queue: 2, Forwarding classes: VOICE
  Queued:
    Packets      :           400          0 pps
    Bytes        :          61344          0 bps
  Transmitted:
    Packets      :           400          0 pps
    Bytes        :          61344          0 bps
  ...
Queue: 3, Forwarding classes: NC
  Queued:
    Packets      :              0          0 pps

```

```

Bytes          :          0          0 bps
...

```

```

user@R0> show interfaces queue se-1/0/0

```

```

Physical interface: se-1/0/0, Enabled, Physical link is Up

```

```

Interface index: 141, SNMP ifIndex: 35

```

```

Forwarding classes: 8 supported, 8 in use

```

```

Egress queues: 8 supported, 8 in use

```

```

Queue: 0, Forwarding classes: DATA

```

```

  Queued:

```

```

    Packets      :          350          0 pps

```

```

    Bytes        :        24350          0 bps

```

```

  Transmitted:

```

```

    Packets      :          350          0 pps

```

```

    Bytes        :        24350          0 bps

```

```

  ..

```

```

Queue: 1, Forwarding classes: expedited-forwarding

```

```

  Queued:

```

```

    Packets      :              0          0 pps

```

```

    Bytes        :              0          0 bps

```

```

  ...

```

```

Queue: 2, Forwarding classes: VOICE

```

```

  Queued:

```

```

    Packets      :           100          0 pps

```

```

    Bytes        :       15272          0 bps

```

```

  Transmitted:

```

```

    Packets      :           100          0 pps

```

```

    Bytes        :       15272          0 bps

```

```

  ...

```

```

Queue: 3, Forwarding classes: NC

```

```

  Queued:

```

```

    Packets      :           19          0 pps

```

```

    Bytes        :          247          0 bps

```

```

  Transmitted:

```

```

    Packets      :           19          0 pps

```

```

    Bytes        :          247          0 bps

```

```

  ...

```

```

user@R0> show interfaces queue se-1/0/1

```

```

Physical interface: se-1/0/1, Enabled, Physical link is Up

```

```

Interface index: 142, SNMP ifIndex: 38

```

```

Forwarding classes: 8 supported, 8 in use

```

```

Egress queues: 8 supported, 8 in use

```

```

Queue: 0, Forwarding classes: DATA

```

```

  Queued:

```

```

    Packets      :          350          0 pps

```

```

    Bytes        :        24350          0 bps

```

```

  Transmitted:

```

```

    Packets      :          350          0 pps

```

```

    Bytes        :        24350          0 bps

```

```

  ...

```

```

Queue: 1, Forwarding classes: expedited-forwarding

```

```

  Queued:

```

```

    Packets      :              0          0 pps

```

```

    Bytes        :              0          0 bps

```

```

  ...

```

```

Queue: 2, Forwarding classes: VOICE

```

```

  Queued:

```

```

    Packets      :           300          0 pps

```

```

    Bytes        :       45672          0 bps

```

```

  Transmitted:

```

```

Packets          :          300          0 pps
Bytes            :          45672        0 bps
...
Queue: 3, Forwarding classes: NC
Queued:
Packets          :          18          0 pps
Bytes            :          234          0 bps
Transmitted:
Packets          :          18          0 pps
Bytes            :          234          0 bps

```

What It Means—The output from these commands shows the packets transmitted and queued on each queue of the link services interface and its constituent links. Table 89 on page 309 shows a summary of these values. (Because the number of transmitted packets equaled the number of queued packets on all the links, this table shows only the queued packets.)

Table 89: Number of Packets Transmitted on a Queue

Packets Queued	Bundle ls-0/0/0.0	Constituent Link se-1/0/0	Constituent Link se-1/0/1	Explanation
Packets on Q0	600	350	350	The total number of packets transiting the constituent links (350 + 350 = 700) exceeded the number of packets queued (600) on the multilink bundle.
Packets on Q2	400	100	300	The total number of packets transiting the constituent links equaled the number of packets on the bundle.
Packets on Q3	0	19	18	The packets transiting Q3 of the constituent links are for keepalive messages exchanged between constituent links. Thus no packets were counted on Q3 of the bundle.

On the multilink bundle, verify the following:

- The number of packets queued matches the number transmitted. If the numbers match, no packets were dropped. If more packets were queued than were transmitted, packets were dropped because the buffer was too small. The buffer size on the constituent links controls congestion at the output stage. To correct this problem, increase the buffer size on the constituent links. For more information, see “Defining and Applying Scheduler Maps” on page 279.
- The number of packets transiting Q0 (600) matches the number of large and small data packets received (100 + 500) on the multilink bundle. If the numbers match, all data packets correctly transited Q0.
- The number of packets transiting Q2 on the multilink bundle (400) matches the number of voice packets received on the multilink bundle. If the numbers match, all voice LFI packets correctly transited Q2.

On the constituent links, verify the following:

- The total number of packets transiting Q0 (350 + 350) matches the number of data packets and data fragments (500 + 200). If the numbers match, all the data packets after fragmentation correctly transited Q0 of the constituent links.

Packets transited both constituent links, indicating that load balancing was correctly performed on non-LFI packets.

- The total number of packets transiting Q2 (300 + 100) on constituent links matches the number of voice packets received (400) on the multilink bundle. If the numbers match, all voice LFI packets correctly transited Q2.

LFI packets from source port 100 transited se-1/0/0, and LFI packets from source port 200 transited se-1/0/1. Thus all LFI (Q2) packets were hashed based on the source port and correctly transited both constituent links.

Corrective Action—If the packets transited only one link, take the following steps to resolve the problem:

1. Determine whether the physical link is **up** (operational) or **down** (unavailable). An unavailable link indicates a problem with the PIM, interface port, or physical connection (link-layer errors). If the link is operational, move to the next step.
2. Verify that the classifiers are correctly defined for non-LFI packets. Make sure that non-LFI packets are not configured to be queued to Q2. All packets queued to Q2 are treated as LFI packets.
3. Verify that at least one of the following values is different in the LFI packets: source address, destination address, IP protocol, source port, or destination port. If the same values are configured for all LFI packets, the packets are all hashed to the same flow and transit the same link.

Why Are Packets Dropped on a PVC Between a J-series Router and Another Vendor?

Problem—I configured a permanent virtual circuit (PVC) between T1, E1, T3, or E3 interfaces on a J-series Services Router and another vendor's router, and packets are being dropped and ping fails.

Solution—If the other vendor's router does not have the same FRF.12 support as the J-series router or supports FRF.12 in a different way, the J-series interface on the PVC might discard a fragmented packet containing FRF.12 headers and count it as a "Policed Discard." As a workaround for this problem, configure multilink bundles on both peers, and configure fragmentation thresholds on the multilink bundles.

Chapter 9

Configuring VoIP

J4350 and J6350 Services Routers support voice over IP (VoIP) connectivity for branch offices with the Avaya IG550 Integrated Gateway. The Avaya IG550 Integrated Gateway consists of four VoIP modules—a TGM550 Telephony Gateway Module and three types of Telephony Interface Modules (TIMs).

The VoIP modules installed in a Services Router at a branch office connect the IP and analog telephones and trunk lines at the branch to headquarters and to the public-switched telephone network.

You can use either J-Web Quick Configuration or a configuration editor to configure VoIP on the Services Router. Alternatively, you can download a complete router configuration that includes VoIP from an Electronic Preinstallation Worksheet (EPW) and a Disk-on-Key USB memory stick.

This chapter contains the following topics:

- VoIP Terms on page 311
- VoIP Overview on page 313
- VoIP Configuration Overview on page 321
- Before You Begin on page 322
- Configuring VoIP Interfaces with EPW and Disk-on-Key on page 323
- Configuring VoIP Interfaces with Quick Configuration on page 325
- Configuring VoIP with a Configuration Editor on page 327
- Accessing and Administering the TGM550 CLI on page 333
- Verifying the VoIP Configuration on page 338
- Frequently Asked Questions About the VoIP Interface on page 342

VoIP Terms

Before configuring VoIP, become familiar with the terms defined in Table 90 on page 312.

Table 90: VoIP Terminology

Term	Definition
bearer bandwidth limit (BBL)	Maximum bandwidth available for voice traffic on an interface when dynamic call admission control is configured on the interface. See also <i>dynamic CAC</i> .
call admission control (CAC)	Method of limiting voice traffic over a particular link in a network. See also <i>dynamic CAC</i> .
centralized automatic message accounting (CAMA)	Recording of toll calls at a central point.
direct inward dialing (DID)	Feature of a trunk line that allows incoming calls to be routed directly to selected stations without help from an attendant.
direct outward dialing (DOD)	Feature of a trunk line that allows outgoing calls to be routed directly without help from an attendant.
direct inward and outward dialing (DIOD)	Feature of a trunk line that allows both incoming and outgoing calls to be routed directly without help from an attendant. See also <i>direct inward dialing (DID)</i> and <i>direct outward dialing (DOD)</i> .
Disk-on-Key	Memory device (stick) that plugs into a USB port to load a complete JUNOS configuration with VoIP onto a Services Router. You must first use an Electronic Preinstallation Worksheet (EPW) to download the configuration to the Disk-on-Key device. The EPW and Disk-on-Key device provide an alternative method to configure the router for VoIP.
dynamic CAC	Application that blocks calls on a WAN interface when the bandwidth is exhausted. See also <i>call admission control (CAC)</i> .
Electronic Preinstallation Worksheet (EPW)	Customized Microsoft Excel spreadsheet used with a Disk-on-Key USB memory stick to configure VoIP on a Services Router. You download the EPW from an Avaya Web site.
emergency transfer relay (ETR)	Feature that provides an emergency link between the telephone connected to the first LINE port on the TGM550 and the trunk connected to the TRUNK port on the TGM550 if power is disconnected from the Services Router or if the TGM550 becomes unregistered from its Media Gateway Controller (MGC).
IEEE 802.1p standard	IEEE standard for a Layer 2 frame structure that supports virtual LAN (VLAN) identification and class-of-service (CoS) traffic classification.
IEEE 802.3af standard	IEEE standard that defines a method for powering network devices via Ethernet cable. Also known as Power over Ethernet (PoE), this standard enables remote devices (such as VoIP telephones) to operate without a separate, external power source. See also <i>Power over Ethernet (PoE)</i> .
ITU H.248 standard	International Telecommunications Union (ITU) standard for communication between a gateway controller and a media gateway.
ITU H.323 standard	International Telecommunications Union (ITU) standard for packet-based multimedia communications over networks that do not guarantee class of service (CoS), such as IP networks. H323, modeled after ISDN PRI, is the standard for voice over IP (VoIP) and conferencing.

Table 90: VoIP Terminology (continued)

Term	Definition
Media Gateway Controller (MGC)	Avaya media server that controls the parts of the call state that pertain to connection control for media channels in a media gateway. The MGC is the controlling entity in an H.248 relationship.
Power over Ethernet (PoE)	Electrical current run to networking devices over Ethernet Category 5 or higher data cables. No extra AC power cord or outlets are needed at the product location.
public switched telephone network (PSTN)	The public worldwide voice telephone network.
standard local survivability (SLS)	Configurable software feature that enables a TGM550 to provide limited Media Gateway Controller (MGC) functionality when no link is available to a registered MGC.
time-division multiplexing (TDM)	A form of multiplexing that divides a transmission channel into successive time slots.
TGM550	Avaya Telephony Gateway Module. Avaya VoIP H.248 media gateway module installed in a Services Router along with one or more Telephony Interface Modules (TIMs) to connect VoIP and legacy analog telephones and trunks over IP networks. Only the TGM550 has an interface configurable through the J-Web interface or JUNOS CLI. The TIMs are configured and administered from the TGM550 CLI.
TIM510	Avaya E1/T1 Telephony Interface Module. Avaya VoIP module installed in a Services Router to provide an E1 or T1 trunk connection over the Internet to a telephone central office (CO). A TIM510 is configured and administered from a TGM550 installed in the same router.
TIM514	Avaya Analog Telephony Interface Module. Avaya VoIP module installed in a Services Router to connect individual telephones or trunk lines to the Internet. A TIM514 is configured and administered from a TGM550 installed in the same router.
TIM521	Avaya BRI Telephony Interface Module. Avaya VoIP module installed in a Services Router to connect ISDN Basic Rate Interface (BRI) trunk lines to a telephone central office (CO) over the Internet for data or voice transmission. A TIM521 is configured and administered from a TGM550 installed in the same router.

VoIP Overview

This section contains the following topics.

- About the Avaya IG550 Integrated Gateway on page 314
- VoIP Interfaces on page 315
- Avaya VoIP Modules Overview on page 316
- Media Gateway Controller on page 317
- Avaya Communication Manager on page 318


- Dynamic Call Admission Control Overview on page 318
- TGM550 Firmware Compatibility with JUNOS Internet Software on page 320
- TGM550 IP Addressing Guidelines on page 320

About the Avaya IG550 Integrated Gateway

The Avaya IG550 Integrated Gateway consists of the TGM550 Telephony Gateway Module and one or more Telephony Interface Modules (TIMs) that are installed in the slots on the J4350 and J6350 Services Routers to provide VoIP connectivity. The TGM550 is an H.248 media gateway that works with the TIMs to connect IP and legacy analog telephones and trunks over IP networks and enable IP telephones to communicate through analog telephone lines and trunks.

The TGM550 is also connected over a LAN or WAN to a Media Gateway Controller (MGC)—an Avaya media server running Avaya Communication Manager (CM) call processing software. The telephony services on the TGM550 are managed by an MGC located at headquarters or in a branch office. When the primary MGC is located at a remote location, the TGM550 uses standard local survivability (SLS) for partial MGC backup in the event that the connection to the primary MGC is lost. J-series Services Routers can thereby provide reliable telephony services to branch offices.

Figure 44 on page 315 shows a typical VoIP topology. The small branch office shown in the expanded illustration on the right is connected over the corporate WAN to the head office through a J6300 Services Router with VoIP modules installed. The Avaya Media Gateway Controller, S8700 Media Server, and integrated Management tools at the head office manage telephony services for headquarters and the branch offices on the WAN, connecting the corporation's legacy analog telephones, VoIP telephones, PCs, and fax machines to the PSTN.

Figure 44: Typical VoIP Topology

VoIP Interfaces

Four types of interfaces on Avaya VoIP modules provide VoIP connectivity on J4350 and J6350 Services Routers:

- Analog telephone or trunk port
- T1 port
- E1 port
- ISDN BRI telephone or trunk port

These interfaces are available on the field-replaceable Avaya VoIP modules listed in Table 91 on page 316. For more information about interface names, see “Network Interface Naming” on page 47. For more information about the modules, see “Avaya VoIP Modules Overview” on page 316.

Table 91: Interfaces on Avaya VoIP Modules

Module Name	Description	VoIP Interfaces	JUNOS Interface (type-pim/0/port)
TGM550	Avaya Telephony Gateway Module (TGM)	<ul style="list-style-type: none"> ■ Two analog telephone ports ■ Two analog trunk ports ■ One serial port for console access 	<p><code>vp-pim/0/0</code></p> <p>On a VoIP interface, the port is always 0.</p>
TIM510	Avaya E1/T1 Telephony Interface Module (TIM)	One E1/T1 trunk port providing up to 30 E1 or 24 T1 channels	–
TIM514	Avaya Analog TIM	<ul style="list-style-type: none"> ■ Four analog telephone ports ■ Four analog trunk ports 	–
TIM521	Avaya BRI TIM	Four ISDN BRI trunk ports providing up to eight channels	–

Only the TGM550 has a JUNOS interface. Because the TIMs do not have corresponding physical interfaces, you cannot configure or administer them with the J-Web interface or the JUNOS CLI. However, you can display TGM550 and TIM status from J-Web Monitor > Chassis pages and with the CLI **show chassis** commands.



NOTE: TIMs are configured and administered from the TGM550 CLI. For more information, see the *Administration Guide and CLI Reference for the Avaya IG550 Integrated Gateway*.



CAUTION: The TGM550 and TIMs are not hot-swappable. You must power off the router before installing or removing the Avaya VoIP modules. Ensure that the Avaya VoIP modules are installed in the router chassis before booting up the system.

Avaya VoIP Modules Overview

A TGM550 and one or more TIMs installed in a Services Router provide telephony exchange services to a branch office over IP networks. Different TIMs have access ports for different types of VoIP and analog telephones and telephone lines. You connect the telephones and lines to the ports on the TGM550 and the TIMs. VoIP telephones require connection to a Power over Ethernet (PoE) adapter or switch that is plugged into an Ethernet port on the Services Router.

VoIP capabilities on the TGM550 enable the Services Router to provide VoIP services to telephones and trunks that do not directly support VoIP. The TGM550 translates voice and signaling data between VoIP and the system used by the telephones and

trunks. TIMs convert the voice path of traditional circuits such as analog trunk and T1 or E1 to a TDM bus inside the router. The TGM550 then converts the voice path from the TDM bus to compressed or uncompressed and packetized VoIP on an Ethernet connection.

Media Gateway Controller

A Media Gateway Controller (MGC) is a media server (call controller) that controls telephone services on the TGM550. An Avaya media server running Avaya Communication Manager (CM) software acts as an MGC for the TGM550.

The following media servers running Avaya Communication Manager can be used as an MGC with the TGM550:

- Avaya S8300 Media Server—Controls up to 49 TGM550s.
- Avaya S8400 Media Server—Controls up to 5 TGM550s.
- Avaya S8500 Media Server—Controls up to 250 TGM550s.
- Avaya S8700 Media Server—Controls up to 250 TGM550s.
- Avaya S8710 Media Server—Controls up to 250 TGM550s.
- Avaya S8720 Media Server—Controls up to 250 TGM550s.

To provide telephony services, the TGM550 must be registered with at least one Media Gateway Controller (MGC). You can configure the IP addresses of up to four MGCs that the TGM550 can connect to in the event of a link failure. The MGC list consists of the IP addresses of the MGCs to connect to and the order in which to reestablish the H.248 link. The first MGC on the list is the primary MGC. The TGM550 searches for the primary MGC first. If it cannot connect to the primary MGC or loses its connection to the primary MGC, it attempts to connect to the next MGC in the list, and so on.



NOTE: The MGC list is stored in the TGM550. It is not written to the JUNOS configuration file.

You must also administer Avaya Communication Manager on the configured Media Gateway Controllers to support the TGM550. For more information, see the following Avaya IG550 Integrated Gateway manuals at <http://support.avaya.com>:

- *Installing and Configuring the Avaya IG550 Integrated Gateway*
- *Administration Guide and CLI Reference for the Avaya IG550 Integrated Gateway*
- *Administrator Guide for Avaya Communication Manager*
- *Avaya Maintenance Procedures for Communication Manager, Media Servers, and Media Gateways*
- *Avaya Maintenance Commands for Communication Manager, Media Servers, and Media Gateways*
- *Avaya Maintenance Alarms for Communication Manager, Media Servers, and Media Gateways*

Avaya Communication Manager

Avaya Communication Manager (CM) software manages the Media Gateway Controller (MGC). Avaya CM allows you to do the following:

- Assign numbers to local telephones.
- Determine where to connect your telephone call based on the number you dial.
- Play dial tones, busy signals, and prerecorded voice announcements.
- Allow or prohibit access to outside lines for specific telephones.
- Assign telephone numbers and buttons to special features.
- Exchange call switching information with older telephone switches that do not support VoIP.



NOTE: The TGM550 supports Avaya Communication Manager (CM) release 4.0 and later releases. The TGM550 does not support Avaya Communication Manager (CM) releases earlier than release 4.0.

For more information about Avaya CM, see the *Administrator Guide for Avaya Communication Manager*.

Dynamic Call Admission Control Overview

Dynamic call admission control (CAC) enables the Media Gateway Controller (MGC) to automatically assign the bandwidth available for voice traffic on WAN interfaces and block new calls when the existing call bandwidth is completely engaged. You configure dynamic CAC on a high-bandwidth primary interface and on one or more backup interfaces with less bandwidth.

Without dynamic CAC, the MGC cannot detect the switchover to the backup link or the resulting changes in network topology and available bandwidth. As a result, the MGC continues to admit calls at the bandwidth of the primary link, causing network congestion and possible jitter, delay, and loss of calls.

Supported Interfaces

Dynamic CAC must be configured on each Services Router interface responsible for providing call bandwidth. You can configure dynamic CAC on the following types of interfaces on Services Routers:

- ADSL
- E1
- E3
- Fast Ethernet
- Gigabit Ethernet
- GRE

- G.SHDSL
- ISDN BRI
- Serial interfaces
- T1
- T3

Bearer Bandwidth Limit and Activation Priority

The dynamic CAC bearer bandwidth limit (BBL) configured on an interface specifies 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.

You configure the dynamic CAC activation priority value on interfaces to specify the order in which the interfaces are used for providing call bandwidth.

Rules for Determining Reported BBL

To assess the WAN interfaces that have an activation priority value and determine a single BBL to report to the MGC, the TGM550 uses the following rules. The reported BBL (RBBL) allows the MGC to automatically control the call bandwidth when interfaces responsible for providing call bandwidth become available or unavailable.

- Report the BBL of the active interface with the highest activation priority. For example, if one interface has the activation priority of 200 and a BBL of 1500 Kbps and another interface has the activation priority of 100 and a BBL of 1000 Kbps, the RBBL is 1500 Kbps.
- If more than one active interface has the same activation priority, report a BBL that is the number of interfaces times their lowest BBL. For example, if two interfaces with the same activation priority have BBLs of 2000 Kbps and 1500 Kbps, the RBBL is 3000 Kbps (2 x 1500 Kbps).
- If the interface with the highest activation priority is unavailable, report the BBL of the active interface with the next highest activation priority.
- If all the interfaces on which dynamic CAC is configured are inactive, report a BBL of 0. The MGC does not allow calls to go through when the RBBL is 0.



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*.

TGM550 Firmware Compatibility with JUNOS Internet Software

The TGM550 firmware version must be compatible with the JUNOS software version installed on the Services Router. For compatibility information, see the *Communication Manager Software & Firmware Compatibility Matrix* at <http://support.avaya.com>.



CAUTION: If the TGM550 firmware version is not compatible with the JUNOS Internet software version on the router, the router does not detect the VoIP interface (`vp-pim/0/0`) and the interface is unavailable. For more information, see “TGM550 Is Installed But the VoIP Interface Is Unavailable” on page 342.

If you are upgrading both the TGM550 firmware and the JUNOS software on the router, first upgrade the TGM550 firmware, and then upgrade the JUNOS software.

For information about upgrading JUNOS software, see the *J-series Services Router Administration Guide*. For information about upgrading the TGM550 firmware, see the *Administration Guide and CLI Reference for the Avaya IG550 Integrated Gateway*.

TGM550 IP Addressing Guidelines

For operational purposes, the TGM550 is identified as a host on the Services Router. Hence, the TGM550 needs to be assigned an IP address that is reachable both externally and internally from the Services Router. The TGM550 uses this IP address to identify itself when communicating with other devices, particularly the Media Gateway Controller (MGC).

To assign the IP address for the TGM550, you configure the destination address on the `vp-pim/0/0` interface. For information about configuring the `vp-pim/0/0` interface, see “Configuring VoIP Interfaces with Quick Configuration” on page 325 or “Configuring the VoIP Interface (Required)” on page 327.



CAUTION: Applying a new or modified IP address resets the TGM550. Before modifying the IP address, take the following precautions:

- Log into the TGM550 and enter `copy running-config startup-config` to save the TGM550 configuration. (For login instructions, see “Accessing and Administering the TGM550 CLI” on page 333.)
 - Ensure that the TGM550 is not currently handling voice traffic.
-

To enable easier administration of the TGM550, we recommend the following guidelines for assigning the IP address of the TGM550:

- Assign an address from one of the subnets that is already configured in the branch office where the Services Router is installed.
- Decide on a block of IP addresses for VoIP services, and assign an IP address from that block to the TGM550.
- Do not assign the following IP addresses to the TGM550:

- A broadcast address (255.255.255.255)
- A class E address (240.0.0.0 to 255.255.255.254)
- A loopback address (127.0.0.0 to 127.255.255.255)
- A multicast address (224.0.0.0 to 239.255.255.255)
- An address with 0 as the first byte or an address with 0 or 255 as the last byte

VoIP Configuration Overview

To configure VoIP, you perform the tasks listed in Table 92 on page 321. For instructions, see the cross-references in the table.

Table 92: VoIP Configuration Overview

Task	Instructions
Perform prerequisite tasks.	“Before You Begin” on page 322
On the Services Router	
1. Make the following assignments: <ul style="list-style-type: none"> ■ Assign an IP address to the Services Router VoIP interface (<code>vp-pim/0/0</code>). ■ Assign an IP address to the TGM550. ■ Assign a Media Gateway Controller (MGC) list for the TGM550. 	<ul style="list-style-type: none"> ■ Configuring VoIP Interfaces with EPW and Disk-on-Key on page 323 ■ Configuring VoIP Interfaces with Quick Configuration on page 325 ■ “Configuring the VoIP Interface (Required)” on page 327 and Configuring the Media Gateway Controller List (Required) on page 329
2. Optionally, configure dynamic CAC on Services Router WAN interfaces. (You must also configure CAC-BL on Avaya Communication Manager (CM)).	“Configuring Dynamic Call Admission Control on WAN Interfaces (Optional)” on page 330
3. Verify the VoIP configuration on the router.	“Verifying the VoIP Configuration” on page 338
4. Perform administrative tasks as necessary. <ul style="list-style-type: none"> ■ Modify the TGM550 IP address ■ Monitor the TGM550 Media Gateway. 	<ul style="list-style-type: none"> ■ Modifying the IP Address of the TGM550 on page 332 ■ For information about monitoring the TGM550 media gateway, see the <i>J-series Services Router Administration Guide</i>
On the TGM550	
1. If you have not already done so, establish a user account (username and password) on the TGM550 with your system administrator.	See the <i>Administration Guide and CLI Reference for the Avaya IG550 Integrated Gateway</i> .
2. If you have not already done so, connect to the TGM550.	See “Connecting Through the TGM550 Console Port” on page 334, “Connecting to the TGM550 with SSH” on page 335 or “Connecting to the TGM550 with Telnet” on page 336.

Table 92: VoIP Configuration Overview (continued)

Task	Instructions
3. Configure the TGM550 and one or more TIMs as necessary for your network.	<p>See the following Avaya manuals:</p> <ul style="list-style-type: none"> ■ <i>Installing and Configuring the Avaya IG550 Integrated Gateway</i> ■ <i>Administration Guide and CLI Reference for the Avaya IG550 Integrated Gateway</i>
4. Perform administrative tasks as necessary.	<ul style="list-style-type: none"> ■ Accessing the Services Router from the TGM550 on page 337 ■ Resetting the TGM550 on page 337 ■ Saving the TGM550 Configuration on page 338 <p>In addition, see the <i>Administration Guide and CLI Reference for the Avaya IG550 Integrated Gateway</i>.</p>
On Avaya Devices	
Configure and administer Avaya Communication Manager (CM) on the configured Media Gateway Controllers (MGCs) to support the IG550 Integrated Gateway.	<p>See the following Avaya manuals:</p> <ul style="list-style-type: none"> ■ <i>Installing and Configuring the Avaya IG550 Media Gateway</i> ■ <i>Administration Guide and CLI Reference for the Avaya IG550 Integrated Gateway</i> ■ <i>Administrator Guide for Avaya Communication Manager</i> ■ <i>Avaya Maintenance Procedures for Communication Manager, Media Servers, and Media Gateways</i> ■ <i>Avaya Maintenance Commands for Communication Manager, Media Servers, and Media Gateways</i> ■ <i>Avaya Maintenance Alarms for Communication Manager, Media Servers, and Media Gateways</i>

Before You Begin

Before you configure VoIP interfaces, you need to perform the following tasks:

- Install Services Router hardware, including the TGM550 and the TIMs. Before power is connected, ensure that the router is grounded with a 10 AWG cable.

For installation and grounding instructions, see the *J4350 and J6350 Services Router Getting Started Guide*.



CAUTION: The original grounding cable for J-series Services Routers is 14 AWG only and must be replaced with a 10 AWG cable.

- Verify that you have connectivity to at least one Avaya media server running Avaya Communication Manager (CM) release 4.0 or later. For more information about Avaya media servers, see “Media Gateway Controller” on page 317.
- Verify that the Services Router is running JUNOS Release 8.2R1 or later.

- Download and install the most recent firmware for the TGM550. Verify that the TGM550 firmware version is compatible with the JUNOS software version installed on the Services Router. For more information, see “TGM550 Firmware Compatibility with JUNOS Internet Software” on page 320.
- If you are configuring VoIP using the Avaya Electronic Preinstallation Worksheet (EPW) and a Disk-on-Key USB memory stick, order a Disk-on-Key USB memory stick. For Disk-on-Key requirements, see “Configuring VoIP Interfaces with EPW and Disk-on-Key” on page 323.
- Establish basic connectivity. For more information, see the *J4350 and J6350 Services Router Getting Started Guide*.
- If you do not already have a basic understanding of physical and logical interfaces and Juniper Networks interface conventions, read “Interfaces Overview” on page 41.
- Applying an IP address to the TGM550 resets the module. If you are updating an existing VoIP configuration by modifying the TGM550 IP address, take the following precautions:
 - Log into the TGM550 and enter `copy running-config startup-config` to save the TGM550 configuration. (For login instructions, see “Accessing and Administering the TGM550 CLI” on page 333.)
 - Ensure that the TGM550 is not currently handling voice traffic.

Configuring VoIP Interfaces with EPW and Disk-on-Key

If you have a new J4350 or J6350 Services Router with the TGM550 and TIMs installed in the router, you can use the Avaya Electronic Preinstallation Worksheet (EPW) and a Disk-on-Key USB memory stick to configure VoIP on the router.

The EPW is a customized Microsoft Excel spreadsheet that you use to collect a complete set of VoIP configuration information and create a configuration file named `juniper-config.txt`. You can copy the `juniper-config.txt` file to a Disk-on-Key device and boot the router from the device to configure VoIP on the router.

This configuration method has the following requirements:

- A management device (PC or laptop) running Microsoft Excel version 2000 or later.
- A Disk-on-Key device with one of the following 16-bit or 32-bit file allocation table (FAT) file systems:
 - DOS 3.0 + 16-bit FAT (up to 32 MB)
 - DOS 3.31 + 16-bit FAT (over 32 MB)
 - WIN95 OSR2 FAT32
 - WIN95 OSR2 FAT32, LBA-mapped
 - WIN95 DOS 16-bit FAT, LBA-mapped

- A Services Router with the factory configuration and the TGM550 and TIMs installed. If other JUNOS configuration files exist on the Services Router, the router cannot read the `juniper-config.txt` file from the Disk-on-Key device. To remove the configuration files from the router, press and hold the **RESET CONFIG** button for 15 seconds or more, until the **STATUS LED** blinks red.



CAUTION: Pressing and holding the **RESET CONFIG** button for 15 seconds or more—until the **STATUS LED** blinks red—deletes all configurations on the router, including the backup configurations and rescue configuration, and loads and commits the factory configuration.

To configure a VoIP interface using EPW and Disk-on-Key:

1. Follow these instructions to download the EPW to a PC or laptop computer.
 - a. Go to <http://support.avaya.com>.
 - b. On the Avaya support page, click **Find Documentation and Technical Information by Product Name**.
 - c. Scroll down and click **Integrated Management — Provisioning & Installation Manager**.
 - d. Select the 4.0 release from the **Select a release** drop-down box and click **View all documents**.
 - e. Scroll down and double click the **Electronic Preinstallation Worksheet for Provisioning Installation Manager** link.
 - f. Scroll down and double click the **View XLS** link.
 - g. In the File Open window, click the **Open** button.
 - h. In the Security Warning window, open the EPW by clicking **Enable Macros**. Be sure to open the EPW in Microsoft Excel version 2000 or later versions.
2. Enter information in the individual worksheets. Ensure that all mandatory fields (highlighted in blue color) are filled in.
3. Select **File > Save**.
4. Open the InitialConfig worksheet and click **Create Configuration File**.
The Select Location page is displayed.
5. Choose a location where you want to create the configuration file.
The configuration file with the name `juniper-config.txt` is created.
6. Copy the `juniper-config.txt` file to a Disk-on-Key device.
7. Press and release the power button to power off the router. Wait for the **POWER LED** to turn off.
8. Plug the Disk-on-Key device into the USB port on the Services Router.

- 9. Power on the router by pressing the **POWER** button on the front panel. Verify that the **POWER** LED on the front panel turns green.

The router reads the `juniper-config.txt` file from the Disk-on-Key device and commits the configuration.
- 10. Remove the Disk-on-Key device.

Configuring VoIP Interfaces with Quick Configuration

You can use the VoIP Interfaces Quick Configuration pages to configure the VoIP interface on a router.

To configure a VoIP interface with Quick Configuration:

- 1. From the Quick Configuration page, as shown in Figure 22 on page 104, select the VoIP interface—for example, `vp-5/0/0`—you want to configure.

For information about interface names, see “Network Interface Naming” on page 47.
- 2. Enter information into the Quick Configuration page, as described in Table 93 on page 325.
- 3. Click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
 - To apply the configuration and return to the main configuration page, click **OK**.
 - To cancel your entries and return to the main page, click **Cancel**.
- 4. To verify that the VoIP interface is configured correctly, see “Verifying the VoIP Configuration” on page 338.

Table 93: VoIP Interface Quick Configuration Page Summary

Field	Function	Your Action
VoIP Logical Interfaces		
Add logical interfaces	Defines logical unit 0 that you connect to the physical VoIP interface. You must define one logical unit for the VoIP interface. NOTE: You cannot define more than one logical unit for the VoIP interface. The logical unit number must be 0 .	Click Add .
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.

Table 93: VoIP Interface Quick Configuration Page Summary (continued)

Field	Function	Your Action
IPv4 Address and Prefix	<p>Specifies the IPv4 address for the interface.</p> <p>The following rules apply:</p> <ul style="list-style-type: none"> ■ You cannot specify more than one IPv4 address. ■ Do not assign the following IPv4 addresses: <ul style="list-style-type: none"> ■ A broadcast address (255.255.255.255) ■ A class E address (240.0.0.0 to 255.255.255.254) ■ A loopback address (127.0.0.0 to 127.255.255.255) ■ A multicast address (224.0.0.0 to 239.255.255.255) ■ An address with 0 as the first byte or an address with 0 or 255 as the last byte ■ The VoIP interface needs a point-to-point connection to the TGM550. To configure the point-to-point connection, specify /32 as the subnet mask in the IPv4 address. 	<p>Type the IPv4 address with /32 as the subnet mask. For example:</p> <p>10.10.10.1/32</p>
Destination Address	<p>Specifies the IP address of the TGM550.</p> <p>CAUTION: Applying a new or modified IP address resets the TGM550. For existing configurations, ensure that the TGM550 configuration is saved (see “Saving the TGM550 Configuration” on page 338) and that the TGM550 module is carrying no voice traffic.</p> <p>You cannot specify more than one IP address. For more information, see “TGM550 IP Addressing Guidelines” on page 320.</p>	<p>Type the IP address of the TGM550—for example, 10.10.10.2.</p>
Physical Interface Description	<p>(Optional) Adds supplemental information about the VoIP physical interface on the router.</p>	<p>Type a text description of the physical VoIP interface in the box to clearly identify it in monitoring displays.</p>
TGM Configuration		
MGC List	<p>Specifies the IP address of at least one and up to four Media Gateway Controllers (MGCs) with which the TGM550 must be registered.</p> <p>The first MGC in the list is the primary MGC. The TGM550 searches for the primary MGC first. If it cannot connect to the primary MGC, the TGM550 searches for the next MGC on the list, and so on.</p>	<ol style="list-style-type: none"> 1. Type the IP address of the MGC. 2. Click Add. <p>To delete an IP address, select it in the MGC List box, then click Delete.</p>

Configuring VoIP with a Configuration Editor

To configure VoIP on a Services Router, perform the following tasks marked *(Required)*. Perform other tasks if needed on your network.

- Configuring the VoIP Interface (Required) on page 327
- Configuring the Media Gateway Controller List (Required) on page 329
- Configuring Dynamic Call Admission Control on WAN Interfaces (Optional) on page 330
- Modifying the IP Address of the TGM550 on page 332

Configuring the VoIP Interface (Required)

You must assign a local IP address to the `vp-pim/0/0` interface on the Services Router and also a destination IP address to the TGM550, so that they can communicate with each other.

To configure the VoIP interface on the Services Router:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 94 on page 327.
3. If you are finished configuring the router, commit the configuration.
4. Go on to “Configuring the Media Gateway Controller List (Required)” on page 329.

Table 94: Configuring the VoIP Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<div>1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration.</div> <div>2. Next to Interfaces, click Edit.</div>	From the [edit] hierarchy level, enter edit interfaces vp-3/0/0
Select the VoIP interface—for example, vp-3/0/0.	In the Interface name column, click the VoIP interface name <code>vp-3/0/0</code> .	
Create the logical unit 0. NOTE: You cannot configure more than one logical unit on the VoIP interface. The logical unit number must be 0.	<div>1. Next to Unit, click Add new entry.</div> <div>2. In the Interface unit number box, type 0.</div>	Enter edit unit 0

Table 94: Configuring the VoIP Interface *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Configure the source IPv4 address—for example, 10.10.10.1/32—for the VoIP interface.</p> <p>The following rules apply:</p> <ul style="list-style-type: none"> ■ You cannot specify more than one IPv4 address. ■ Do not assign the following IPv4 addresses: <ul style="list-style-type: none"> ■ A broadcast address (255.255.255.255) ■ A class E address (240.0.0.0 to 255.255.255.254) ■ A loopback address (127.0.0.0 to 127.255.255.255) ■ A multicast address (224.0.0.0 to 239.255.255.255) ■ An address with 0 as the first byte or an address with 0 or 255 as the last byte ■ The VoIP interface needs a point-to-point connection to the TGM550. To configure the point-to-point connection, specify /32 as the subnet mask in the IPv4 address. 	<ol style="list-style-type: none"> Under Family, select the Inet check box and click Configure. Next to Address, click Add new entry. In the Source box, type 10.10.10.1/32. 	<p>Enter</p> <p>set family inet address 10.10.10.1/32 destination 10.10.10.2</p>
<p>Configure the destination IP address—for example 10.10.10.2—for the TGM550. The TGM550 uses this IP address to identify itself when communicating with other devices, particularly the Media Gateway Controller (MGC).</p> <p>CAUTION: Applying a new or modified IP address resets the TGM550. For existing configurations, ensure that the TGM550 configuration is saved (see “Saving the TGM550 Configuration” on page 338) and that the TGM550 module is carrying no voice traffic.</p> <p>You cannot specify more than one IP address. For more information, see “TGM550 IP Addressing Guidelines” on page 320.</p>	<ol style="list-style-type: none"> In the Destination box, type 10.10.10.2. Click OK until you return to the Interfaces page. 	

Configuring the Media Gateway Controller List (Required)

To provide telephony services, the TGM550 must be registered with at least one Media Gateway Controller (MGC). You can configure the IP addresses of up to four MGCs that the TGM550 can connect to in the event of a link failure. For more information, see “Media Gateway Controller” on page 317.

In addition to configuring the MGC list from a J-Web Quick Configuration page (see Table 93 on page 325) and the JUNOS CLI, you can log in to the TGM550 and configure the list. For more information, see the *Administration for the Avaya IG550 Integrated Gateway*.

This section contains the following topics:

- Configuring an MGC List and Adding Addresses on page 329
- Clearing an MGC List on page 330

Configuring an MGC List and Adding Addresses

In the following example, a TGM550 installed in slot 2 of a Services Router has the IP address 10.10.10.2. The TGM550 needs to have registered a primary MGC at address 172.16.0.0, and second and third MGC at addresses 10.10.10.30 and 10.10.10.40.

To configure the MGC list with the JUNOS CLI:

1. Enter operational mode on the JUNOS CLI.
2. To configure the IP addresses of the Media Gateway Controllers, enter the `set tgm fpc slot media-gateway-controller` command with the IP addresses of the primary, second, and third MGC:

```
user@host> set tgm fpc 2 media-gateway-controller [172.16.0.0 10.10.10.30
10.10.10.40]
```



NOTE: Running the `set tgm fpc slot media-gateway-controller` command updates the startup configuration on the TGM550. You do not need to run the `copy running-config start-config` command to save the configuration on the module.

3. Log in to the TGM550 with SSH, and verify that each MGC can be reached over the network.

```
user@host> ssh 10.10.10.2

password> root

TGM550-00<root># ping 172.16.0.0

...

TGM550-00<root># ping 10.10.10.30
```

...

TGM550-00<root># **ping 10.10.10.40**

...

4. Do one of the following:
 - To control bandwidth assignments for voice traffic, go on to “Configuring Dynamic Call Admission Control on WAN Interfaces (Optional)” on page 330.
 - To verify that VoIP is configured correctly on the router, see “Verifying the VoIP Configuration” on page 338.

Clearing an MGC List

In the following example, a TGM550 is installed in slot 2 of the router.

To remove all the IP addresses from the MGC list, with the JUNOS CLI:

1. Enter operational mode on the CLI.
2. Enter the **clear tgm fpc slot media-gateway-controller** command:

```
user@host> clear tgm fpc 2 media-gateway-controller
```

The **clear** command removes all the MGC IP addresses. You cannot clear the IP address of a single MGC with this command.

3. Add one or more new MGC IP addresses. (See “Configuring an MGC List and Adding Addresses” on page 329.)

Configuring Dynamic Call Admission Control on WAN Interfaces (Optional)

To configure dynamic call admission control (CAC), you define the bearer bandwidth limit (BBL) and activation priority on each WAN interface responsible for providing call bandwidth.

- The activation priority has a range from 1 through 255. The default value is 50.
- The BBL has a range from 0 Kbps through 9999 Kbps. The default BBL value of -1 Kbps indicates that the complete bandwidth of the interface is available for voice traffic. Use a BBL of 0, which indicates that no bandwidth is available for bearer traffic on the MGC, to use the interface for signaling only.

In this example, a Gigabit Ethernet, T1, and ISDN BRI interface are configured with the BBL and activation priority values shown in Table 95 on page 331.

Table 95: Dynamic CAC Configuration Example

Interface Providing Call Bandwidth	Bearer Bandwidth Limit (BBL) Value	Activation Priority Value
Gigabit Ethernet	3000 Kbps	200
T1	1000 Kbps	150
ISDN BRI	128 Kbps	100

The Gigabit Ethernet interface is used as the primary link for providing call bandwidth because it has the highest activation priority value. When the Gigabit Ethernet interface is active, the TGM550 reports its BBL value of 3000 Kbps to the MGC. If the Gigabit Ethernet interface fails, the TGM550 automatically switches over to the T1 interface because it has the next highest activation priority. The TGM550 now reports the BBL value of the T1 interface to the MGC. If the T1 interface also fails, the TGM550 switches over to the ISDN BRI interface and reports the BBL value of the ISDN BRI interface to the MGC. Configuring dynamic CAC on multiple WAN interfaces allows the MGC to automatically control the call bandwidth when interfaces responsible for providing call bandwidth are unavailable.

For more information about dynamic CAC, see “Dynamic Call Admission Control Overview” on page 318.

To configure dynamic CAC on Services Router WAN interfaces:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 96 on page 331.
3. If you are finished configuring the router, commit the configuration.
4. Configure Call Admission Control: Bandwidth Limitation (CAC-BL) on Avaya Communication Manager. For more information, see the *Administrator Guide for Avaya Communication Manager*.
5. To verify that dynamic CAC is configured correctly, see “Verifying the VoIP Configuration” on page 338.

Table 96: Configuring Dynamic CAC

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	From the [edit] hierarchy level, enter edit interfaces ge-0/0/3
Select the Gigabit Ethernet interface—for example, ge-0/0/3.	In the Interface name column, click ge-0/0/3 .	

Table 96: Configuring Dynamic CAC (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure dynamic CAC on logical unit 0 of the Gigabit Ethernet interface with the activation priority and BBL values given in Table 95 on page 331.	<ol style="list-style-type: none"> Under Unit, next to 0, click Edit. Next to Dynamic call admission control, click Configure or Edit. In the Activation priority box, type 200. In the Bearer bandwidth limit box, type 3000. Click OK until you return to the Interfaces page. 	<ol style="list-style-type: none"> Enter edit unit 0 Enter set dynamic-call-admission-control activation-priority 200 bearer-bandwidth-limit 3000
Select the T1 interface—for example, t1-6/0/0.	In the Interface name column, click t1-6/0/0 .	From the [edit] hierarchy level, enter edit interfaces t1-6/0/0
Configure dynamic CAC on logical unit 0 of the T1 interface with the activation priority and BBL values given in Table 95 on page 331.	<ol style="list-style-type: none"> Under Unit, next to 0, click Edit. Next to Dynamic call admission control, click Configure or Edit. In the Activation priority box, type 150. In the Bearer bandwidth limit box, type 1000. Click OK until you return to the Interfaces page. 	<ol style="list-style-type: none"> Enter edit unit 0 Enter set dynamic-call-admission-control activation-priority 150 bearer-bandwidth-limit 1000
Select the ISDN BRI interface—for example, br-1/0/3.	In the Interface name column, click br-1/0/3 .	From the [edit] hierarchy level, enter edit interfaces br-1/0/3
Configure dynamic CAC on logical unit 0 of the ISDN BRI interface with the activation priority and BBL values given in Table 95 on page 331.	<ol style="list-style-type: none"> Under Unit, next to 0, click Edit. Next to Dynamic call admission control, click Configure or Edit. In the Activation priority box, type 100. In the Bearer bandwidth limit box, type 128. Click OK. 	<ol style="list-style-type: none"> Enter edit unit 0 Enter set dynamic-call-admission-control activation-priority 100 bearer-bandwidth-limit 128

Modifying the IP Address of the TGM550



CAUTION: The TGM550 is reset when you commit the configuration after modifying the IP address. Before modifying the TGM550 IP address, take the following precautions:

- Log into the TGM550 and enter `copy running-config startup-config` to save the TGM550 configuration. (For login instructions, see “Accessing and Administering the TGM550 CLI” on page 333.)
- Ensure that the TGM550 is not currently handling voice traffic.

To modify the IP address of the TGM550:

1. Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 97 on page 333.
3. If you are finished configuring the router, commit the configuration.

Table 97: Modifying the IP Address of the TGM550

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none">1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration.2. Next to Interfaces, click Edit.	From the [edit] hierarchy level, enter edit interfaces vp-3/0/0 unit 0
Select the logical VoIP interface—for example, vp-3/0/0.0 .	<ol style="list-style-type: none">1. In the Interface name column, click the VoIP interface name vp-3/0/0.2. In the Interface unit number box, click 0.	
Modify the destination IP address for the TGM550 to a different address—for example, 10.10.10.80 . For guidelines, see “TGM550 IP Addressing Guidelines” on page 320. NOTE: You cannot specify more than one IP address.	<ol style="list-style-type: none">1. Under Family, next to Inet, click Edit.2. Under Address, in the Broadcast column, click Edit.3. In the Destination box, type 10.10.10.80.4. Click OK.	Enter set family inet address 10.10.10.1/32 destination 10.10.10.80

Accessing and Administering the TGM550 CLI

The CLI on the TGM550 allows you to configure, monitor, and diagnose the TGM550 and TIMs installed in a Services Router. You can access the TGM550 from a management device attached to the TGM550 console port or by opening a Telnet or secure shell (SSH) session from the JUNOS CLI on the Services Router.

You can also open a remote Telnet or SSH session directly to the TGM550 from a network location, or indirectly through the JUNOS CLI from a dial-up connection with a USB modem attached to the router. (For information about the modem connection, see the *J-series Services Router Administration Guide*.)

This section contains the following topics. For complete information about the TGM550 CLI, see the *Administration Guide and CLI Reference for the Avaya IG550 Integrated Gateway*.

- TGM550 Access Requirements on page 334
- Connecting Through the TGM550 Console Port on page 334
- Connecting to the TGM550 with SSH on page 335
- Accessing the TGM550 with Telnet on page 335
- Accessing the Services Router from the TGM550 on page 337
- Resetting the TGM550 on page 337
- Saving the TGM550 Configuration on page 338

TGM550 Access Requirements

Administrators can use the root password to access the TGM550 initially, but all users need a TGM550 user account (username and password) set up by the network administrator for regular access to the module. For information about user accounts on a TGM550, see the *Administration Guide and CLI Reference for the Avaya IG550 Integrated Gateway*.



NOTE: You cannot use a Services Router user account to access the TGM550 CLI.

- A console connection requires the Ethernet rollover cable and adapter provided with the TGM550. (See “Connecting Through the TGM550 Console Port” on page 334.)
- An SSH connection requires that the TGM550 have an IP address assigned.
- A Telnet connection to the TGM550 requires that the module have an IP address and that Telnet service be enabled on the module.

To assign an IP address to the TGM550, see “Configuring VoIP Interfaces with Quick Configuration” on page 325 or “Configuring the VoIP Interface (Required)” on page 327.

To enable Telnet, see “Accessing the TGM550 with Telnet” on page 335.

Connecting Through the TGM550 Console Port

To connect to the TGM550 through its console port:

1. Turn off the power to the management device, such as a PC or laptop computer, that you are using to access the TGM550.
2. Plug one end of an Ethernet rollover cable provided with the TGM550 into the RJ-45 to DB-9 serial port adapter provided with the TGM550.



CAUTION: Two different RJ-45 cables and RJ-45 to DB-9 adapters are provided. Do not use the RJ-45 cable and adapter for the Services Router console port to connect to the TGM550 console port.

3. Plug the RJ-45 to DB-9 serial port adapter provided with the TGM550 into the serial port on the management device.
4. Connect the other end of the Ethernet rollover cable to the console port (**CONSOLE**) on the TGM550.
5. Turn on power to the management device.
6. Start your asynchronous terminal emulation application (such as Microsoft Windows Hyperterminal), and select the appropriate **COM** port to use (for example, **COM1**).
7. Configure the port settings as follows:
 - Bits per second: 9600
 - Data bits: 8
 - Parity: None
 - Stop bits: 1
 - Flow control: Hardware
8. At the login prompt, type your username and press Enter.
9. At the password prompt, type your password and press Enter.

Connecting to the TGM550 with SSH

To connect to the TGM550 with SSH:

1. Ensure that the TGM550 has an IP address. (See “Configuring VoIP Interfaces with Quick Configuration” on page 325 or “Configuring the VoIP Interface (Required)” on page 327.)
2. From the JUNOS CLI or a remote connection, enter the following command:

```
ssh ip-address
```

Accessing the TGM550 with Telnet

By default, Telnet service is not enabled on the TGM550. You must enable Telnet service on the TGM550 before you can telnet to the TGM550 from other devices or from the TGM550 to other devices.



CAUTION: Telnet connections are not encrypted and therefore can be intercepted.

This section contains the following topics:

- Enabling Telnet Service on the TGM550 on page 336
- Connecting to the TGM550 with Telnet on page 336
- Disabling Telnet Service on the TGM550 on page 336

Enabling Telnet Service on the TGM550

To enable Telnet service on the TGM550:

1. Connect to the TGM550 through the console port. (See “Connecting Through the TGM550 Console Port” on page 334.
2. To enable incoming Telnet connections, enter the following command, replacing *port* with the Telnet port number:

```
TGM550-004(super)# ip telnet port port
```

3. To enable outgoing Telnet connections from the TGM550 to other devices, enter

```
TGM550-004(super)# ip telnet-client
```

4. Save the configuration:

```
TGM550-004(super)# copy running-config startup-config
```

Connecting to the TGM550 with Telnet

To connect to the TGM550 with Telnet:

1. Ensure that Telnet is enabled on the TGM550. (See “Enabling Telnet Service on the TGM550” on page 336.)
2. Ensure that the TGM550 has an IP address. (See “Configuring VoIP Interfaces with Quick Configuration” on page 325 or “Configuring the VoIP Interface (Required)” on page 327.)
3. From the JUNOS CLI or a remote connection, enter the following command:

```
telnet ip-address
```

Disabling Telnet Service on the TGM550

To disable Telnet service on the TGM550:

1. Connect to the TGM550 through the console port. For more information, see “Connecting Through the TGM550 Console Port” on page 334.
2. To disable incoming Telnet connections, enter the following command, replacing *port* with the Telnet port number:

```
TGM550-004(super)# no ip telnet
```

3. To disable outgoing Telnet connections from the TGM550 to other devices, enter

```
TGM550-004(super)# no ip telnet-client
```

4. Save the configuration:

```
TGM550-004(super)# copy running-config startup-config
```

Accessing the Services Router from the TGM550

You can access the Services Router from the CLI on its installed TGM550 in the following ways:

- Enter the **session chassis** command.
- Enter the **telnet** or **ssh** command.



NOTE: Before using the TGM550 CLI **telnet** command, ensure that Telnet service is enabled on the TGM550. For more information, see “Enabling Telnet Service on the TGM550” on page 336.

Resetting the TGM550



CAUTION: Before resetting the TGM550, take the following precautions:

- Log into the TGM550 and enter **copy running-config startup-config** to save the TGM550 configuration. (For login instructions, see “Accessing and Administering the TGM550 CLI” on page 333.)
 - Ensure that the TGM550 is not currently handling voice traffic.
-

You can reset the TGM550 from the module itself or from the Services Router.

To reset the TGM550 from the module itself, do one of the following:

- Press the **RST** button on the TGM550.
- Log into the TGM550, and enter the **reset** command. (For login instructions, see “Accessing and Administering the TGM550 CLI” on page 333.)

To reset the TGM550 from the Services Router:

1. Enter operational mode in the CLI.
2. Enter the **request chassis fpc slot slot-number restart** command.

For example, to reset a TGM550 installed in slot 2 on the router chassis, enter

```
user@host> request chassis fpc slot 2 restart
```



NOTE: You cannot reset the TIMs using the `request chassis fpc slot slot-number restart` command. TIMs are administered only from the TGM550.

Saving the TGM550 Configuration

To save the configuration on the TGM550:

1. Log in to the TGM550. (For login instructions, see “Accessing and Administering the TGM550 CLI” on page 333.)
2. To save the configuration, enter

```
TGM550-004(super)# copy running-config startup-config
```

For more information about saving a TGM550 configuration, see the *Administration Guide and CLI Reference for the Avaya IG550 Integrated Gateway*.

Verifying the VoIP Configuration

To verify the VoIP configuration, perform the following tasks:

- Verifying the VoIP Interface on page 339
- Verifying the Media Gateway Controller List on page 341
- Verifying Bandwidth Available for VoIP Traffic on page 341

Verifying the VoIP Interface

Purpose Verify that the VoIP interface is correctly configured.

Action From the CLI, enter the show interfaces extensive command.

```

user@host> show interfaces vp-3/0/0 extensive
Physical interface: vp-3/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 21, Generation: 142
  Type: VP-AV, Link-level type: VP-AV, MTU: 1518, Speed: 10mbps
  Device flags   : Present Running
  Link type      : Full-Duplex
  Link flags     : None
  Physical info  : Unspecified
  CoS queues     : 8 supported, 8 maximum usable queues
  Last flapped   : 2006-09-29 09:28:32 UTC (4d 18:35 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input bytes   :                8886912                0 bps
    Output bytes  :                6624354                0 bps
    Input packets :                90760                0 pps
    Output packets:                65099                0 pps
  Input errors:
    Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards:
0,
  Resource errors: 0
  Output errors:
    Carrier transitions: 0, Errors: 0, Drops: 0, MTU errors: 0, Resource errors:
0
  Egress queues: 8 supported, 8 in use
  Queue counters:      Queued packets  Transmitted packets      Dropped packets

    0 best-effort                65099                65099                0
    1 expedited-fo                 0                 0                0
    2 assured-forw                 0                 0                0
    3 network-cont                 0                 0                0

  Packet Forwarding Engine configuration:
    Destination slot: 2
    CoS transmit queue      Bandwidth      Buffer Priority
Limit
    0 best-effort           %      bps      %      usec
none
    3 network-control       5      500000    5      0      low
none

  Logical interface vp-3/0/0.0 (Index 71) (SNMP ifIndex 47) (Generation 137)
  Flags: Point-To-Point SNMP-Traps Encapsulation: VP-AV
  Protocol inet, MTU: 1500, Generation: 142, Route table: 0
  Flags: None
  Filters: Input: pcap, Output: pcap
  Addresses, Flags: Is-Preferred Is-Primary
    Destination: 10.10.10.2, Local: 10.10.10.1, Broadcast: Unspecified,
Generation: 144

```

- What It Means** The output shows a summary of interface information. Verify the following information:
- The physical interface is **Enabled**. If the interface is shown as **Disabled**, do either of the following:
 - In the CLI configuration editor, delete the **disable** statement at the [edit **interfaces** *interface-name*] level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the **Interfaces** > *interface-name* page.
 - The physical link is **Up**. A link state of **Down** indicates that the interface is disabled. Do either of the following:
 - In the CLI configuration editor, delete the **disable** statement at the [edit **interfaces** *interface-name*] level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the **Interfaces** > *interface-name* page.
 - The **Last Flapped** time is an expected value. The **Last Flapped** time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
 - The traffic statistics reflect expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics** *interface-name* command.
- See Also** For a complete description of **show interfaces** output, see the *JUNOS Interfaces Command Reference*.

Verifying the Media Gateway Controller List

Purpose Verify that the Media Gateway Controller (MGC) list is correctly configured and that the MGCs are reachable over the network.

Action From the operational mode in the CLI, enter `show tgm fpc slot-number media-gateway-controller`.

```
user@host> show tgm fpc 2 media-gateway-controller
Media gateway controller(s): 173.26.232.77
                             10.10.10.30
                             10.10.10.40
```

What It Means The output shows the configured MGC list. Verify the following:

- The IP addresses and the order of the IP addresses in the MGC list are correct. The first MGC on the list is the primary MGC. The TGM550 searches for the primary MGC first. If it cannot connect to the primary MGC or loses its connection to the primary MGC, it attempts to connect to the next MGC in the list, and so on.
- Use the JUNOS CLI `ping` command or the J-Web ping host tool (**Diagnose > Ping Host**) to verify that the configured MGCs can be reached over the network.

See Also For a complete description of `show tgm fpc` output, see the *JUNOS Interfaces Command Reference*.

Verifying Bandwidth Available for VoIP Traffic

Purpose Verify that the dynamic call admission control (CAC) configuration supports sufficient bandwidth for VoIP traffic.

Action From the operational mode in the CLI, enter `show tgm dynamic-call-admission-control`.

```
user@host> show tgm dynamic-call-admission-control
Reported bearer bandwidth limit: 3000 Kbps
Interface      State      Activation  Bearer bandwidth
               priority   priority    limit (Kbps)
ge-0/0/3.0     up         200         3000
tl-6/0/0.0     up         150         1000
br-1/0/3.0     up         50          128
```

What It Means The output shows the dynamic CAC configuration. Verify the following information:

- The activation priority and bearer bandwidth limit (BBL) configured on individual interfaces are correct.
- The **Reported bearer bandwidth limit** field displays the bandwidth available for VoIP traffic. Ensure that the bandwidth is sufficient for VoIP traffic.

See Also For a complete description of `show tgm dynamic-call-admission-control` output, see the *JUNOS Interfaces Command Reference*.

Frequently Asked Questions About the VoIP Interface

Use answers to the following question to solve configuration problems on a VoIP interface:

- TGM550 Is Installed But the VoIP Interface Is Unavailable on page 342

TGM550 Is Installed But the VoIP Interface Is Unavailable

Problem—I installed the TGM550 Telephony Gateway Module and configured the VoIP interface—for example, `vp-3/0/0`—but the interface is not accessible. The `show chassis hardware` command displays the TGM550 installed on slot 3. However, the `show interfaces terse` command does not display the `vp-3/0/0` interface, and the `show interfaces vp-3/0/0` command displays an error:

```
user@host> show interfaces vp-3/0/0
error: device vp-3/0/0 not found
```

Solution—The VoIP interface might be unavailable because the TGM550 firmware version is not compatible with the JUNOS software version installed on the Services Router. For more information, see “TGM550 Firmware Compatibility with JUNOS Internet Software” on page 320.

To correct the TGM550 firmware and JUNOS software version compatibility error:

1. Check the router’s system log messages for a version incompatibility error similar to the following:

```
Jan 5 11:07:03 host fwdd[2857]: TGMT: RE (1.0) - TGM (2.0) major protocol
version mismatch: not marking TGM slot ready
```

2. If the error exists, connect to the TGM550 through the console port. (See “Connecting Through the TGM550 Console Port” on page 334.)
3. To view the TGM550 firmware version, enter

```
TGM550-003(super)# show image version
```

Bank	Version
-----	-----
A (current)	26.23.0
B	26.22.0

In this example, the current TGM550 firmware version is 26.23.0.

4. Refer to the *Communication Manager Software & Firmware Compatibility Matrix* at <http://support.avaya.com> to identify the JUNOS software version that is compatible with the current TGM550 firmware version.
5. Upgrade the router with the compatible JUNOS software version. For information about upgrading JUNOS software, see the *J-series Services Router Administration Guide*.

Chapter 10

Configuring uPIMs as Ethernet Switches

The 6-port, 8-port, and 16-port Gigabit Ethernet uPIMs can function as Ethernet access switches that switch traffic at Layer 2, in addition to routing traffic at Layer 3.

This chapter contains the following topics:

- Gigabit Ethernet uPIM Switch Overview on page 343
- Joining uPIMs in a Daisy-Chain on page 344
- Configuring Gigabit Ethernet uPIM Switches on page 344
- Verifying Gigabit Ethernet uPIM Switch Configuration on page 345

Gigabit Ethernet uPIM Switch Overview

You can deploy a J-series router with multiport uPIMs in branch offices as an access or desktop switch with integrated routing capability, thus eliminating intermediate access switch devices from your topology. The Gigabit Ethernet uPIM provides Ethernet switching while the Routing Engine provides routing functionality, enabling you to use a single chassis to provide routing, access switching, and WAN interfaces.

You can set a multiport uPIM to either of two modes of operation: routing (the default) or switching. Routed traffic is forwarded from any port of the Gigabit Ethernet uPIM to the WAN interface. Switched traffic is forwarded from one port of the Gigabit Ethernet uPIM to another port on the same Gigabit Ethernet uPIM. Switched traffic is not forwarded from a port on one uPIM to a port on a different uPIM.

In routing mode, the multiport uPIM has the same configuration options as any other Gigabit Ethernet interface. To configure uPIM Gigabit Ethernet interfaces in routing mode, see “Configuring Gigabit Ethernet Interfaces with Quick Configuration” on page 115 and “Configuring Network Interfaces with a Configuration Editor” on page 131.

In switching mode, the uPIM appears in the list of interfaces as a single interface which is the first interface on the uPIM, for example **ge-2/0/0**. You can optionally configure each uPIM port only for autonegotiation, speed, and duplex mode. A uPIM in switching mode can perform the following functions:

- Layer 3 forwarding—Routes traffic destined for WAN interfaces and other PIMs present on the chassis.
- Layer 2 forwarding—Switches intra-LAN traffic from one host on the LAN to another LAN host (one port of uPIM to another port of same uPIM).

Gigabit Ethernet uPIMs in switching mode have the following limitations:

- Virtual LAN (VLAN) tagged traffic is switched transparently and is not limited to the VLAN.
- Layer 2 control plane protocols such as Spanning Tree Protocol (STP) and Link Aggregation Control Protocol (LACP) are not supported, limiting usage in switching mode for connection to other switches.

Joining uPIMs in a Daisy-Chain

You cannot combine multiple uPIMs to act as a single integrated switch. However, you can connect uPIMs on the same chassis externally by physically connecting a port on one uPIM to a port on another uPIM in a daisy-chain fashion.

Two or more uPIMs daisy-chained together create a single switch with a higher port count than either individual uPIM. One port on each uPIM is used solely for the connection. For example, if you daisy-chain a 6-port uPIM and an 8-port uPIM, the result operates as a 12-port uPIM. Any port of a uPIM can be used for daisy-chaining.

Configure the IP address for only one of the daisy-chained uPIMs, making it the primary uPIM. The secondary uPIM routes traffic to the primary uPIM, which forwards it to the Routing Engine. This results in some increase in latency and packet drops due to oversubscription of the external link.

Only one link between the two uPIMs is supported. Connecting more than one link between uPIMs creates a loop topology, which is not supported.

Configuring Gigabit Ethernet uPIM Switches

When you set a multiport uPIM to switching mode, the uPIM appears as a single entity for monitoring purposes. The only physical port settings that you can configure are autonegotiation, speed, and duplex mode on each uPIM port, and these settings are optional.



NOTE: You cannot configure switch ports from J-Web Quick Configuration pages. You must use the J-Web or CLI configuration editor.

To configure a multiport Gigabit Ethernet uPIM as a switch:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 98 on page 345.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying Gigabit Ethernet uPIM Switch Configuration” on page 345.

Table 98: Configuring uPIMs as Switches

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Chassis level of the configuration hierarchy	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Chassis, click Configure. 	From the [edit] hierarchy level, enter edit chassis
Set the uPIM mode of operation to switching. NOTE: Routing mode is the default setting.	<ol style="list-style-type: none"> 1. Next to Fpc, click Add new entry. 2. In the Slot field, enter the number of the slot of the chassis in which the uPIM is inserted, and click OK. 3. Next to Pic, click Add new Entry. 4. Enter 0 in the Slot field. (This number is always 0 on a J-series Services Router.) 5. Next to Ethernet, click Configure. 6. From the Pic mode list, choose Switching and click OK. 	Enter set pim <i>pim-number</i> pic 0 ethernet pic-mode switching
(Optional) Set the physical port parameters for each port on the uPIM.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 3. Click the name of the uPIM interface—for example ge-2/0/0. 4. Next to Switch options, click Configure. 5. Next to Switch port, click Add new entry. 6. In the Port field, enter the number of the port you want to configure. 7. Choose the settings for Autonegotiation, Link mode, and Speed, and click OK. 	<ol style="list-style-type: none"> 1. From the [edit] hierarchy level, enter edit interfaces 2. Configure parameters for each uPIM port that you want to specify: <i>ge-pim/0/0</i> switch-port <i>port-number</i> (auto-negotiation no-auto-negotiation) speed (10m 100m 1g) link-mode (full-duplex half-duplex) For example: set <i>ge-2/0/0</i> switch-port 1 autonegotiation

Verifying Gigabit Ethernet uPIM Switch Configuration

The operational mode command for checking the status and statistics for multiport uPIMs switching mode is different from that of routing mode. For uPIMs in routing mode, the operational commands are the same as for other Gigabit Ethernet interfaces, such as the 1-port Gigabit Ethernet ePIM and built-in Gigabit Ethernet ports.

Not all operational mode commands are supported for ports of a uPIM in switching mode. For example, the operational mode command for monitoring port statistics is not supported.



NOTE: To clear the statistics for the individual switch ports using the `clear interfaces statistics ge-pim/0/0 switch-port port-number` command.

Verifying Status of uPIM Switch Ports

Purpose To verify the status and view statistics for a port on a uPIM in switching mode.

Action From the CLI, enter the `show interfaces ge-pim/0/0 switch-port port-number` command.

```
user@host show interfaces ge-pim/0/0 switch-port port-number
Port 0, Physical link is Up
Speed: 100mbps, Auto-negotiation: Enabled
Statistics:
  Receive      Transmit
  Total bytes   28437086     21792250
  Total packets 409145       88008
  Unicast packets 9987        83817
  Multicast packets 145002       0
  Broadcast packets 254156      4191
  Multiple collisions 23          10
  FIFO/CRC/Align errors 0           0
  MAC pause frames 0            0
  Oversized frames 0
  Runt frames 0
  Jabber frames 0
  Fragment frames 0
  Discarded frames 0
Autonegotiation information:
Negotiation status: Complete
Link partner:
  Link mode: Full-duplex, Flow control: None, Remote fault: OK, Link
partner Speed: 100 Mbps
Local resolution:
  Flow control: None, Remote fault: Link OK
```

Part 3

Configuring Routing Protocols

- Routing Overview on page 349
- Configuring Static Routes on page 383
- Configuring a RIP Network on page 395
- Configuring an OSPF Network on page 411
- Configuring the IS-IS Protocol on page 431
- Configuring BGP Sessions on page 441

Chapter 11

Routing Overview

Routing is the process of delivering a message across a network or networks. This process has two primary components: the exchange of routing information to forward packets accurately from source to destination and the packet-forwarding procedure.

To use the routing capabilities of a J-series Services Router, you must understand the fundamentals of IP routing and the routing protocols that are primarily responsible for the transmission of unicast traffic. To read this chapter, you need a basic understanding of IP addressing and TCP/IP.



NOTE: Unless otherwise specified, J-series Services Routers support IPv6 addressing and routing. For information about IPv6, see the *JUNOS Routing Protocols Configuration Guide*.

This chapter includes the following topics. For more information, see the *JUNOS Routing Protocols Configuration Guide*.

- Routing Terms on page 349
- Routing Overview on page 354
- RIP Overview on page 360
- OSPF Overview on page 364
- IS-IS Overview on page 368
- BGP Overview on page 371

Routing Terms

To understand routing, become familiar with the terms defined in Table 99 on page 349.

Table 99: Routing Terms

Term	Definition
adjacency	Portion of the local routing information that pertains to the reachability of a single neighbor over a single circuit or interface.

Table 99: Routing Terms (*continued*)

Term	Definition
area	Administrative group of OSPF networks within an autonomous system (AS) that operates independently from other areas in the AS. Multiple areas within an AS reduce the amount of link-state advertisement (LSA) traffic on the network and the size of topology databases.
area border router (ABR)	In OSPF, a router having interfaces in multiple areas of an autonomous system (AS) so that it can link the areas to each other. An area border router maintains a separate topological database for each area it is connected to and shares topology information between areas.
AS path	In BGP, the list of autonomous system (ASs) that a packet must traverse to reach a given set of destinations within a single AS.
autonomous system (AS)	Network, collection of routers, or portion of a large internetwork under a single administrative authority.
backbone area	In OSPF, the central area in an autonomous system (AS) to which all other areas are connected by area border routers (ABRs). The backbone area always has the area ID 0.0.0.0.
bidirectional connectivity	Ability of directly connected devices to communicate with each other over the same link.
Border Gateway Protocol (BGP)	Exterior gateway protocol used to exchange routing information among routers in different autonomous systems.
broadcast	Operation of sending network traffic from one network node to all other network nodes.
cluster	In BGP, a set of routers that have been grouped together. A cluster consists of one system that acts as a route reflector, along with any number of client peers. The client peers receive their route information only from the route reflector system. Routers in a cluster do not need to be fully meshed.
confederation	In BGP, a group of autonomous systems (ASs) that appears to external ASs to be a single AS.
confederation sequence	Ordered set of autonomous systems (ASs) for a confederation. The closest AS in the path is first in the sequence.
convergence	After a topology change, the time all the routers in a network take to receive the information and update their routing tables.
cost	Unitless number assigned to a path between neighbors, based on throughput, round-trip time, and reliability. The sum of path costs between source and destination hosts determines the overall path cost. OSPF uses the lowest cost to determine the best path.
designated router (DR)	In OSPF, a node designated to process link-state advertisements (LSAs) and distribute topology updates for an autonomous system (AS).
distance vector	Number of hops to a routing destination.
dynamic routing	Routing method that enables the route of a message through a network to change as network conditions change. Compare <i>static routing</i> .
end systems	Network entities that send and receive packets.
exterior gateway protocol (EGP)	Protocol that exchanges routing information between autonomous systems (ASs). BGP is an EGP. Compare <i>interior gateway protocol (IGP)</i> .

Table 99: Routing Terms (*continued*)

Term	Definition
external BGP (EBGP)	BGP configuration in which sessions are established between routers in different autonomous systems (ASs).
external peer	In BGP, a peer that resides in a different autonomous system (AS) from the Services Router.
external route	Route to an area outside the network.
flooding	Technique by which a router forwards traffic to every node attached to the router, except the node from which the traffic arrived. Flooding is a simple but sometimes inefficient way to distribute routing information quickly to every node in a network. RIP and OSPF are flooding protocols, but BGP is not.
forwarding table	JUNOS software forwarding information base (FIB). The JUNOS routing protocol process installs active routes from its routing tables into the Routing Engine forwarding table. The kernel copies this forwarding table into the Packet Forwarding Engine, which is responsible for determining which interface transmits the packets.
full mesh	Network in which devices are organized in a mesh topology, with each node connected to every other network node.
gateway router	Node on a network that serves as an entrance to another network.
global AS	Global autonomous system (AS). An AS consisting of multiple subautonomous systems (sub-ASs).
handshake	Process of exchanging signaling information between two communications devices to establish the method and transmission speed of a connection.
hello packet	In OSPF, a packet sent periodically by a router to first establish and then maintain network adjacency, and to discover neighbor routers.
hold time	Maximum number of seconds allowed to elapse between the time a BGP system receives successive keepalive or update messages from a peer.
hop	Trip a data packet takes from one router to another in the network. The number of routers through which a packet passes to get from its source to its destination is known as the hop count. In general, the best route is the one with the shortest hop count.
intermediate systems	Network entities that relay (forward) packets as well as send and receive them on the network. Intermediate systems are also known as routers.
Intermediate System-to-Intermediate System (IS-IS)	Link-state, interior gateway routing protocol for IP networks that also uses the shortest-path-first (SPF) algorithm to determine routes.
interior gateway protocol (IGP)	Protocol that exchanges routing information within autonomous systems (ASs). IS-IS, OSPF, and RIP are IGPs. Compare <i>exterior gateway protocol (EGP)</i> .
Internal BGP (IBGP)	BGP configuration in which sessions are established between routers in the same autonomous systems (ASs).
internal peer	In BGP, a peer that resides in the same autonomous system (AS) as the Services Router.
keepalive message	Periodic message sent by one BGP peer to another to verify that the session between them is still active.

Table 99: Routing Terms (*continued*)

Term	Definition
latency	Delay that occurs when a packet or signal is transmitted over a communications system.
link-state advertisement (LSA)	Messages that announce the presence of OSPF-enabled interfaces to adjacent OSPF interfaces (neighbors). The exchange of LSAs establishes bidirectional connectivity between neighbors.
local preference	Optional BGP path attribute carried in internal BGP update packets that indicates the degree of preference for an external route.
mesh	Network topology in which devices are organized in a manageable, segmented manner with many, often redundant, interconnections between network nodes. See also <i>full mesh</i> .
metric	Numerical value that determines how quickly a packet can reach its destination. See also <i>cost</i> .
multiple exit discriminator (MED)	Optional BGP path attribute consisting of a metric value that is used to determine the exit point to a destination when all other factors in determining the exit point are equal.
neighbor	Adjacent router interface. A node can directly route packets to its neighbors only. See also <i>peer</i> .
network	Series of nodes interconnected by communication paths.
network diameter	Maximum hop count in a network.
network topology	Arrangement of nodes and connections in a network.
node	Connection point that operates as a redistribution point or an end point in a network, recognizing data transmissions and either forwarding or processing them.
notification message	Message sent between BGP peers to inform the receiving peer that the sending peer is terminating the session because an error occurred, and explaining the error.
not-so-stubby area (NSSA)	In OSPF, a type of stub area in which external route advertisements can be flooded.
open message	Message sent between BGP peers to establish communication.
Open Shortest Path First protocol (OSPF)	A link-state interior gateway protocol (IGP) that makes routing decisions based on the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm).
origin	Value assigned to a BGP route to indicate whether the first router to advertise the route learned it from an external, internal, or unknown source.
path-vector protocol	Protocol that uses the path between autonomous systems (ASs) to select the best route, rather than the shortest distance or the characteristics of the route (link state). BGP is a path-vector protocol. In contrast, RIP is a distance-vector protocol, and OSPF and IS-IS are link-state protocols.
peer	Immediately adjacent router with which a protocol relationship has been established. See also <i>neighbor</i> .
peering	The practice of exchanging Internet traffic with directly connected peers according to commercial and contractual agreements.
point of presence (POP)	Access point to the Internet, having a unique IP address, where telecommunications equipment is located. POPs usually belong to Internet service providers (ISPs) or telephone companies.

Table 99: Routing Terms (*continued*)

Term	Definition
poison reverse	An efficiency technique in a RIP network. By setting the number of hops to an unavailable router to 16 hops or more, a router informs all the other routers in the network. Because RIP allows only up to 15 hops to another router, this technique reduces RIP updates and helps defeat large routing loops. See also <i>split horizon</i> .
propagation	Process of translating and forwarding route information discovered by one routing protocol in the update messages of another routing protocol. Route propagation is also called route redistribution.
reachability	In BGP, the feasibility of a route.
round-robin	Scheduling algorithm in which items have the same priority and are handled in a fixed cyclic order.
route advertisement	Distribution of routing information at specified intervals throughout a network, to establish adjacencies with neighbors and communicate usable routes to active destinations. See also <i>link-state advertisement (LSA)</i> .
route aggregation	Combining groups of routes with common addresses into a single entry in the routing table, to decrease routing table size and the number of route advertisements sent by a router.
route reflection	In BGP, configuring a group of routers into a cluster and having one system act as a route reflector, redistributing routes from outside the cluster to all routers in the cluster. Routers in a cluster do not need to be fully meshed.
Routing Information Protocol (RIP)	Distance-vector routing protocol that keeps a database of routing information gathered from periodic broadcasts by each router in a network.
routing table	Table stored on a router that keeps track of all possible paths (routes) between sources and destinations in a network and, in some cases, metrics associated with the routes.
split horizon	An efficiency technique in a RIP network. A router reduces the number of RIP updates in the network by not retransmitting a route advertisement out the interface through which it was received. Split-horizon updates also help prevent routing loops. See also <i>poison reverse</i> .
static routing	Routing method in which routes are manually entered in the routing table and do not change unless you explicitly update them. Unlike dynamic routes, which must be imported into the routing table each time a host comes online, static routes are available immediately. Static routes are generally preferred over other types of routes. Compare <i>dynamic routing</i> .
stub area	In OSPF, an area through which or into which autonomous system (AS) external route advertisements are not flooded.
subautonomous system (sub-AS)	Autonomous system (AS) members of a BGP confederation.
subnetwork	Subdivision of a network, which functions exactly like a network except that it has a more specific address and subnet mask (destination prefix).
three-way handshake	Process by which two routers synchronize protocols and establish a bidirectional connection.
topology database	Map of connections between the nodes in a network. The topology database is stored in each node.

Table 99: Routing Terms *(continued)*

Term	Definition
triggered update	In a network that uses RIP, a routing update that is automatically sent whenever routing information changes.
virtual link	In OSPF, a link you create between two area border routers (ABRs) that have an interface to a common nonbackbone area, to connect a third area to the backbone area. One of the area border routers must be directly connected to the backbone area.

Routing Overview

Routing is the transmission of data packets from a source to a destination address. For packets to be correctly forwarded to the appropriate host address, the host must have a unique numeric identifier or IP address. The unique IP address of the destination host forms entries in the routing table. These entries are primarily responsible for determining the path that a packet traverses when transmitted from source to destination.

This overview contains the following topics:

- Networks and Subnetworks on page 354
- Autonomous Systems on page 355
- Interior and Exterior Gateway Protocols on page 355
- Routing Tables on page 355
- Forwarding Tables on page 356
- Dynamic and Static Routing on page 357
- Route Advertisements on page 357
- Route Aggregation on page 358

Networks and Subnetworks

Large groups of machines that are interconnected and can communicate with one another form networks. Typically, networks identify large systems of computers and devices that are owned or operated by a single entity. Traffic is routed between or through the networks as data is passed from host to host.

As networks grow large, the ability to maintain the network and effectively route traffic between hosts within the network becomes increasingly difficult. To accommodate growth, networks are divided into subnetworks. Fundamentally, subnetworks behave exactly like networks, except that they are identified by a more specific network address and subnet mask (destination prefix). Subnetworks have routing gateways and share routing information in exactly the same way as large networks.

Autonomous Systems

A large network or collection of routers under a single administrative authority is termed an autonomous system (AS). Autonomous systems are identified by a unique numeric identifier that is assigned by the Internet Assigned Numbers Authority (IANA). Typically, the hosts within an AS are treated as internal peers, and hosts in a peer AS are treated as external peers. The status of the relationship between hosts—internal or external—governs the protocol used to exchange routing information.


Interior and Exterior Gateway Protocols

Routing information that is shared within an AS is transmitted by an interior gateway protocol (IGP). Of the different IGPs, the most common are RIP, OSPF, and IS-IS. IGPs are designed to be fast acting and light duty. They typically incorporate only a moderate security system, because trusted internal peers do not require the stringent security measures that untrusted peers require. As a result, you can usually begin routing within an AS by enabling the IGP on all internal interfaces and performing minimal additional configuration. You do not need to establish individual adjacencies.

Routing information that is shared with a peer AS is transmitted by an exterior gateway protocol (EGP). The primary EGP in use in almost all networks is the Border Gateway Protocol (BGP). BGP is designed to be very secure. Individual connections must be explicitly configured on each side of the link. As a result, although large numbers of connections are difficult to configure and maintain, each connection is secure.

Routing Tables

To route traffic from a source host to a destination host, the routers through which the traffic will pass must learn the path that the packet is to take. Once learned, the information is stored in routing tables. The routing table maintains a list of all the possible paths from point A to point B. Figure 45 on page 356 shows a simple network of routers.

Figure 45: Simple Network Topology

This simple network provides multiple ways to get from Host San Francisco to Host Miami. The packet can follow the path through Denver and Cleveland. Alternatively, the packet can be routed through Phoenix and directly to Miami. The routing table includes all the possible paths and combinations—an exhaustive list of all the ways to get from the source to the destination.

The routing table must include every possible path from a source to a destination. Routing tables for the network in Figure 45 on page 356 must include entries for San Francisco-Denver, San Francisco-Cleveland, San Francisco-Miami, Denver-Cleveland, and so on. As the number of sources and destinations increases, the routing table quickly becomes large. The unwieldy size of routing tables is the primary reason for the division of networks into subnetworks.

Forwarding Tables

If the routing table is a list of all the possible paths a packet can take, the forwarding table is a list of only the best routes to a particular destination. The best path is determined according to the particular routing protocol being used, but generally the number of hops between the source and destination determines the best possible route.

In the network shown in Figure 45 on page 356, because the path with the fewest number of hops from San Francisco to Miami is through Phoenix, the forwarding table distills all the possible San Francisco-Miami routes into the single route through Phoenix. All traffic with a destination address of Miami is sent directly to the next hop, Phoenix.

After it receives a packet, the Phoenix router performs another route lookup, using the same destination address. The Phoenix router then routes the packet appropriately. Although it considers the entire path, the router at any individual hop along the way is responsible only for transmitting the packet to the next hop in the path. If the Phoenix router is managing its traffic in a particular way, it might send the packet through Houston on its route to Miami. This scenario is likely if specific

customer traffic is treated as priority traffic and routed through a faster or more direct route, while all other traffic is treated as nonpriority traffic.


Dynamic and Static Routing

Entries are imported into a router's routing table from dynamic routing protocols or by manual inclusion as static routes. Dynamic routing protocols allow routers to learn the network topology from the network. The routers within the network send out routing information in the form of route advertisements. These advertisements establish and communicate active destinations, which are then shared with other routers in the network.

Although dynamic routing protocols are extremely useful, they have associated costs. Because they use the network to advertise routes, dynamic routing protocols consume bandwidth. Additionally, because they rely on the transmission and receipt of route advertisements to build a routing table, dynamic routing protocols create a delay (latency) between the time a router is powered on and the time during which routes are imported into the routing table. Some routes are therefore effectively unavailable until the routing table is completely updated, when the router first comes online or when routes change within the network (due to a host going offline, for example).

Static routing avoids the bandwidth cost and route import latency of dynamic routing. Static routes are manually included in the routing table, and never change unless you explicitly update them. Static routes are automatically imported into the routing table when a router first comes online. Additionally, all traffic destined for a static address is routed through the same router. This feature is particularly useful for networks with customers whose traffic must always flow through the same routers. Figure 46 on page 357 shows a network that uses static routes.

Figure 46: Static Routing Example




In Figure 46 on page 357, the customer routes in the **192.176.14/24** subnetwork are static routes. These are hard links to specific customer hosts that never change. Because all traffic destined for any of these routes is forwarded through Router A, these routes are included as static routes in Router A's routing table. Router A then advertises these routes to other hosts so that traffic can be routed to and from them.

Route Advertisements

The routing table and forwarding table contain the routes for the routers within a network. These routes are learned through the exchange of route advertisements. Route advertisements are exchanged according to the particular protocol being employed within the network.

Generally, a router transmits hello packets out each of its interfaces. Neighboring routers detect these packets and establish adjacencies with the router. The adjacencies are then shared with other neighboring routers, which allows the routers to build up the entire network topology in a topology database, as shown in Figure 47 on page 358.

Figure 47: Route Advertisement



In Figure 47 on page 358, Router A sends out hello packets to each of its neighbors. Routers B and C detect these packets and establish an adjacent relationship with Router A. Router B and C then share this information with their neighbors, Routers D and E, respectively. By sharing information throughout the network, the routers create a network topology, which they use to determine the paths to all possible destinations within the network. The routes are then distilled into the forwarding table of best routes according to the route selection criteria of the protocol in use.

Route Aggregation

As the number of hosts in a network increases, the routing and forwarding tables must establish and maintain more routes. As these tables become larger, the time routers require to look up particular routes so that packets can be forwarded becomes prohibitive. The solution to the problem of growing routing tables is to group (aggregate) the routers by subnetwork, as shown in Figure 48 on page 359.


Figure 48: Route Aggregation

Figure 48 on page 359 shows three different ASs. Each AS contains multiple subnetworks with thousands of host addresses. To allow traffic to be sent from any host to any host, the routing tables for each host must include a route for each destination. For the routing tables to include every combination of hosts, the flooding of route advertisements for each possible route becomes prohibitive. In a network of hosts numbering in the thousands or even millions, simple route advertisement is not only impractical but impossible.

By employing route aggregation, instead of advertising a route for each host in AS 3, the gateway router advertises only a single route that includes all the routes to all the hosts within the AS. For example, instead of advertising the particular route `170.16.124.17`, the AS 3 gateway router advertises only `170.16/16`. This single route advertisement encompasses all the hosts within the `170.16/16` subnetwork, which reduces the number of routes in the routing table from 2^{16} (one for every possible IP address within the subnetwork) to 1. Any traffic destined for a host within the AS is forwarded to the gateway router, which is then responsible for forwarding the packet to the appropriate host.

Similarly, in this example, the gateway router is responsible for maintaining 2^{16} routes within the AS (in addition to any external routes). The division of this AS into subnetworks allows for further route aggregation to reduce this number. In the

subnetwork in the example, the subnetwork gateway router advertises only a single route (170.16.124/24), which reduces the number of routes from 2^8 to 1.

RIP Overview

In a Routing Information Protocol (RIP) network, each router's forwarding table is distributed among the nodes through the flooding of routing table information. Because topology changes are flooded throughout the network, every node maintains the same list of destinations. Packets are then routed to these destinations based on path-cost calculations done at each node in the network.

This overview contains the following topics:

- Distance-Vector Routing Protocols on page 360
- Maximizing Hop Count on page 361
- RIP Packets on page 361
- Split Horizon and Poison Reverse Efficiency Techniques on page 362
- Limitations of Unidirectional Connectivity on page 363

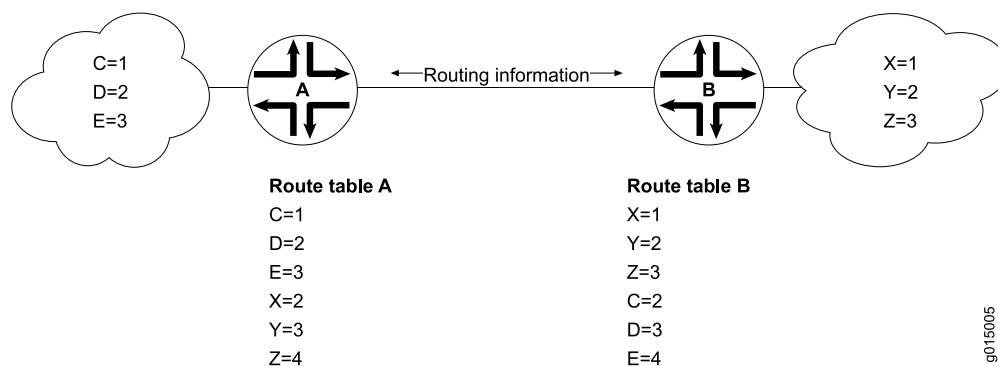


NOTE: The J-series Services Router supports both RIP version 1 and RIP version 2. In this guide, the term *RIP* refers to both versions of the protocol.

Distance-Vector Routing Protocols

Distance-vector routing protocols transmit routing information that includes a distance vector, typically expressed as the number of hops to the destination. This information is flooded out all protocol-enabled interfaces at regular intervals (every 30 seconds in the case of RIP) to create a network map that is stored in each node's local topology database. Figure 49 on page 360 shows how distance-vector routing works.

Figure 49: Distance-Vector Protocol



In Figure 49 on page 360, Routers A and B have RIP enabled on adjacent interfaces. Router A has known RIP neighbors Routers C, D, and E, which are 1, 2, and 3 hops

away, respectively. Router B has known RIP neighbors Routers X, Y, and Z, which are 1, 2, and 3 hops away, respectively. Every 30 seconds, each router floods its entire routing table information out all RIP-enabled interfaces. In this case, flooding exchanges routing table information across the RIP link.

When Router A receives routing information from Router B, it adds 1 to the hop count to determine the new hop count. For example, Router X has a hop count of 1, but when Router A imports the route to X, the new hop count is 2. The imported route also includes information about where the route was learned, so that the original route is imported as a route to Router X through Router B with a hop count of 2.

When multiple routes to the same host are received, RIP uses the distance-vector algorithm to determine which path to import into the forwarding table. The route with the smallest hop count is imported. If there are multiple routes with the same hop count, all are imported into the forwarding table, and traffic is sent along the paths in round-robin fashion.

Maximizing Hop Count

The successful routing of traffic across a RIP network requires that every node in the network maintain the same view of the topology. Topology information is broadcast between RIP neighbors every 30 seconds. If Router A is many hops away from a new host, Router B, the route to B might take significant time to propagate through the network and be imported into Router A's routing table. If the two routers are 5 hops away from each other, Router A cannot import the route to Router B until 2.5 minutes after Router B is online. For large numbers of hops, the delay becomes prohibitive. To help prevent this delay from growing arbitrarily large, RIP enforces a maximum hop count of 15 hops. Any prefix that is more than 15 hops away is treated as unreachable and assigned a hop count equal to infinity. This maximum hop count is called the network diameter.

RIP Packets

Routing information is exchanged in a RIP network by RIP request and RIP response packets. A router that has just booted can broadcast a RIP request on all RIP-enabled interfaces. Any routers running RIP on those links receive the request and respond by sending a RIP response packet immediately to the router. The response packet contains the routing table information required to build the local copy of the network topology map.

In the absence of RIP request packets, all RIP routers broadcast a RIP response packet every 30 seconds on all RIP-enabled interfaces. The RIP broadcast is the primary way in which topology information is flooded throughout the network.

Once a router learns about a particular destination through RIP, it starts a timer. Every time it receives a new response packet with information about the destination, the router resets the timer to zero. However, if the router receives no updates about a particular destination for 180 seconds, it removes the destination from its RIP routing table.

In addition to the regular transmission of RIP packets every 30 seconds, if a router detects a new neighbor or detects that an interface is unavailable, it generates a triggered update. The new routing information is immediately broadcast out all


RIP-enabled interfaces, and the change is reflected in all subsequent RIP response packets.

Split Horizon and Poison Reverse Efficiency Techniques

Because RIP functions by periodically flooding the entire routing table out to the network, it generates a lot of traffic. The split horizon and poison reverse techniques can help reduce the amount of network traffic originated by RIP hosts and make the transmission of routing information more efficient.


If a router receives a set of route advertisements on a particular interface, RIP determines that those advertisements do not need to be retransmitted out the same interface. This technique, known as split horizon, helps limit the amount of RIP routing traffic by eliminating information that other neighbors on that interface have already learned. Figure 50 on page 362 shows an example of the split horizon technique.

Figure 50: Split Horizon Example



In Figure 50 on page 362, Router A advertises routes to Routers C, D, and E to Router B. In this example, Router A can reach Router C in 2 hops. When Router A advertises the route to Router B, B imports it as a route to Router C through Router A in 3 hops. If Router B then readvertised this route to Router A, A would import it as a route to Router C through Router B in 4 hops. However, the advertisement from Router B to Router A is unnecessary, because Router A can already reach the route in 2 hops. The split horizon technique helps reduce extra traffic by eliminating this type of route advertisement.


Similarly, the poison reverse technique helps to optimize the transmission of routing information and improve the time to reach network convergence. If Router A learns about unreachable routes through one of its interfaces, it advertises those routes as unreachable (hop count of 16) out the same interface. Figure 51 on page 363 shows an example of the poison reverse technique.

Figure 51: Poison Reverse Example

In Figure 51 on page 363, Router A learns through one of its interfaces that routes to Routers C, D, and E are unreachable. Router A readvertises those routes out the same interface as unreachable. The advertisement informs Router B that Hosts C, D, and E are definitely not reachable through Router A.

Limitations of Unidirectional Connectivity

Because RIP processes routing information based solely on the receipt of routing table updates, it cannot ensure bidirectional connectivity. As Figure 52 on page 363 shows, RIP networks are limited by their unidirectional connectivity.

Figure 52: Limitations of Unidirectional Connectivity

In Figure 52 on page 363, Routers A and D flood their routing table information to Router B. Because the path to Router E has the fewest hops when routed through Router A, that route is imported into Router B's forwarding table. However, suppose that Router A can transmit traffic but is not receiving traffic from Router B due to an unavailable link or invalid routing policy. If the only route to Router E is through Router A, any traffic destined for Router A is lost, because bidirectional connectivity was never established.

OSPF establishes bidirectional connectivity with a three-way handshake. For more information, see “Link-State Advertisements” on page 364.

OSPF Overview

In an Open Shortest Path First (OSPF) network, the network topology is distributed among the nodes of the autonomous system (AS) and is regularly updated through the exchange of link-state advertisements (LSAs). As a result, OSPF is known as a link-state protocol. Because topology changes are flooded throughout the network, every node maintains the same copy of the network map in its local topological database. Packets are then routed based on the shared topology using the shortest path first (SPF) algorithm.

This overview contains the following topics:

- Link-State Advertisements on page 364
- Role of the Designated Router on page 365
- Path Cost Metrics on page 365
- Areas and Area Border Routers on page 365
- Role of the Backbone Area on page 366
- Stub Areas and Not-So-Stubby Areas on page 367




NOTE: The J-series Services Router supports both OSPF version 2 and OSPF version 3. In this guide, the term *OSPF* refers to both versions of the protocol.

Link-State Advertisements

OSPF creates a topology map by flooding link-state advertisements (LSAs) across OSPF-enabled links. LSAs announce the presence of OSPF-enabled interfaces to adjacent OSPF interfaces. The exchange of LSAs establishes bidirectional connectivity between all adjacent OSPF interfaces (neighbors) using a three-way handshake, as shown in Figure 53 on page 364.

Figure 53: OSPF Three-Way Handshake



In Figure 53 on page 364, Router A sends hello packets out all its OSPF-enabled interfaces when it comes online. Router B receives the packet, which establishes that Router B can receive traffic from Router A. Router B generates a response to Router A to acknowledge receipt of the hello packet. When Router A receives the response, it establishes that Router B can receive traffic from Router A. Router A then generates a final response packet to inform Router B that Router A can receive traffic from Router B. This three-way handshake ensures bidirectional connectivity.

As new neighbors are added to the network or existing neighbors lose connectivity, the adjacencies in the topology map are modified accordingly through the exchange (or absence) of LSAs. These LSAs advertise only the incremental changes in the network, which helps minimize the amount of OSPF traffic on the network. The adjacencies are shared and used to create the network topology in the topological database.

Role of the Designated Router

Large local area networks (LANs) that have many routers and therefore many OSPF adjacencies can produce heavy control-packet traffic as LSAs are flooded across the network. To alleviate the potential traffic problem, OSPF uses designated routers (DRs). Rather than broadcasting LSAs to all their OSPF neighbors, the routers send their LSAs to the designated router, which processes the LSAs, generates responses, and multicasts topology updates to all OSPF routers.

In LANs, the election of the designated router takes place when the OSPF network is initially established. When the first OSPF links are active, the router with the highest router identifier (defined by the `router-id` configuration value or the loopback address) is elected designated router. The router with the second highest router identifier is elected the backup designated router (BDR). If the designated router fails or loses connectivity, the BDR assumes its role and a new BDR election takes place between all the routers in the OSPF network.

Path Cost Metrics

Once the topology is shared across the network, OSPF uses it to route packets between network nodes. Each path between neighbors is assigned a cost based on the throughput, round-trip time, and reliability of the link. The sum of the costs across a particular path between hosts determines the overall cost of the path. Packets are then routed along the shortest path using the shortest path first (SPF) algorithm. If multiple equal-cost paths exist between a source and destination address, OSPF routes packets along each path alternately, in round-robin fashion.

OSPF allows you to manually assign a cost (or metric) to a particular path segment to control the flow of packets across the network.


Areas and Area Border Routers

The OSPF networks in an AS are administratively grouped into areas. Each area within an AS operates like an independent network and has a unique 32-bit area ID, which functions like a network address. Within an area, the topology database contains only information about the area, LSAs are flooded only to nodes within the area, and routes are computed only within the area. Subnetworks are divided into other areas, which are connected to form the whole of the main network.

The central area of an AS, called the backbone area, has a special function and is always assigned the area ID 0.0.0.0. Area IDs are unique numeric identifiers, in dotted decimal notation, but they are not IP addresses. Area IDs need only be unique within an AS. All other networks or areas in the AS must be directly connected to the backbone area by a router that has interfaces in more than one area. These

connecting routers are called area border routers (ABRs). Figure 54 on page 366 shows an OSPF topology of three areas connected by two area border routers.

Figure 54: Multiarea OSPF Topology




Area border routers are responsible for sharing topology information between areas. They summarize the link-state records of each area and advertise destination address summaries to neighboring areas. The advertisements contain the ID of the area in which each destination lies, so that packets are routed to the appropriate area border router. For example, in the OSPF areas shown in Figure 54 on page 366, packets sent from Router A to Router C are automatically routed through Area Border Router B.

Role of the Backbone Area

An OSPF restriction requires all areas to be directly connected to the backbone area so that packets can be properly routed. All packets are routed first to the backbone area by default. Packets that are destined for an area other than the backbone area are then routed to the appropriate area border router and on to the remote host within the destination area.


In large networks with many areas, in which direct connectivity between all areas and the backbone area is physically difficult or impossible, you can configure virtual links to connect noncontiguous areas. For example, Figure 55 on page 367 shows a virtual link between a noncontiguous area and the backbone area through an area connected to both.

Figure 55: OSPF Topology with a Virtual Link

In the topology shown in Figure 55 on page 367, a virtual link is established between area 0.0.0.3 and the backbone area through area 0.0.0.2. All outbound traffic destined for other areas is routed through area 0.0.0.2 to the backbone area and then to the appropriate area border router. All inbound traffic destined for area 0.0.0.3 is routed to the backbone area and then through area 0.0.0.2.

Stub Areas and Not-So-Stubby Areas

Figure 56 on page 367 shows an AS across which many external routes are advertised. If external routes make up a significant portion of a topology database, you can suppress the advertisements in areas that do not have links outside the network. By doing so, you can reduce the amount of memory the nodes use to maintain the topology database and free it for other uses.

Figure 56: OSPF AS Network with Stub Areas and NSSAs

To control the advertisement of external routes into an area, OSPF uses stub areas. By designating an area border router interface to the area as a stub interface, you

suppress external route advertisements through the area border router. Instead, the area border router automatically advertises a default route (through itself) in place of the external routes. Packets destined for external routes are automatically sent to the area border router, which acts as a gateway for outbound traffic and routes them appropriately.

For example, area 0.0.0.3 in Figure 56 on page 367 is not directly connected to the outside network. All outbound traffic is routed through the area border router to the backbone and then to the destination addresses. By designating area 0.0.0.3 a stub area, you reduce the size of the topology database for that area by limiting the route entries to only those routes internal to the area.

Like area 0.0.0.3 in Figure 56 on page 367, area 0.0.0.4 has no external connections. However, area 0.0.0.4 has static customer routes that are not internal OSPF routes. You can limit the external route advertisements to the area and advertise the static customer routes by designating it a not-so-stubby area (NSSA). External routes are flooded into the NSSA and then leaked to the other areas, but external routes from other areas are not advertised within the NSSA.

IS-IS Overview

The Intermediate System-to-Intermediate System (IS-IS) protocol is a classless interior routing protocol developed by the International Organization for Standardization (ISO) as part of the development of the Open Systems Interconnection (OSI) protocol suite. Like OSPF routing, IS-IS uses hello packets that allow network convergence to occur quickly when network changes are detected.

This overview contains the following topics:

- IS-IS Areas on page 368
- Network Entity Titles and System Identifiers on page 369
- IS-IS Path Selection on page 369
- Protocol Data Units on page 369

IS-IS Areas

An IS-IS network is a single autonomous system (AS), also called a routing domain, that consists of end systems and intermediate systems. End systems are network entities that send and receive packets. Intermediate systems (routers) send, receive, and relay (forward) packets.

IS-IS does not force the network to use a hierarchical physical topology. Instead, a single AS can be divided into two types of areas: Level 1 areas and Level 2 areas. A Level 1 area is similar to an OSPF stub area, and a Level 2 area interconnects all Level 1 areas. The router and its interfaces reside within one area, and Level 2 routers share link-state information. No IS-IS area functions strictly as a backbone.

Level 1 routers share intra-area routing information, and Level 2 routers share interarea information about IP addresses available within each area. Uniquely, IS-IS

routers can act as both Level 1 and Level 2 routers, sharing intra-area routes with other Level 1 routers and interarea routes with other Level 2 routers.

The propagation of link-state updates is determined by the level boundaries. All routers within a level maintain a complete link-state database of all other routers in the same level. Each router then uses the Dijkstra algorithm to determine the shortest path from the local router to other routers in the link-state database.

Network Entity Titles and System Identifiers

In IS-IS, special network addresses are called network entity titles (NETs) and take several forms, depending on your network requirements. NET addresses are hexadecimal and range from 8 octets to 20 octets in length. Generally, the format consists of an authority and format Identifier (AFI), a domain ID, an area ID, a system identifier, and a selector. The simplest format omits the domain ID and is 10 octets long. For example, the NET address 49.0001.1921.6800.1001.00 consists of the following parts:

- 49—AFI
- 0001—Area ID
- 1921.6800.1001—System identifier
- 00—Selector

The system identifier must be unique within the network. For an IP-only network, we recommend using the IP address of an interface on the router. Configuring a loopback NET address with the IP address is helpful when troubleshooting is required on the network.

IS-IS Path Selection

Level 1 routers store information about all the subnets within an area, and choose intranetwork paths over internetwork paths. Using the area ID portion of the NET address, Level 1 routers determine which neighboring routers are Level 1 routers within the same area.

If the destination address is not within the area, Level 1 routers forward the packet to the nearest router configured as both a Level 1 and Level 2 router within the area. The Level 1 and Level 2 router forwards the packet, using the Level 2 topology, to the proper area. The destination router, which is configured as a Level 1 and Level 2 router, then determines the best path through the destination area.

Protocol Data Units

IS-IS routers use protocol data units (PDUs) to exchange information. Each protocol data unit (PDU) shares a common header.

IS-IS Hello PDU

IS-IS hello PDUs establish adjacencies with other routers and have three different formats: one for point-to-point hello packets, one for Level 1 broadcast links, and

one for Level 2 broadcast links. Level 1 routers must share the same area address to form an adjacency, while Level 2 routers do not have this limitation. The request for adjacency is encoded in the Circuit type field of the PDU.

Hello PDUs have a preset length assigned to them. The IS-IS router does not resize any PDU to match the maximum transmission unit (MTU) on a router interface. Each interface supports the maximum IS-IS PDU of 1492 bytes, and hello PDUs are padded to meet the maximum value. When the hello is sent to a neighboring router, the connecting interface supports the maximum PDU size.

Link-State PDU

A link-state PDU (LSP) contains information about each router in the network and the connected interfaces. Also included is metric and IS-IS neighbor information. Each LSP must be refreshed periodically on the network and is acknowledged by information within a sequence number packet.

On point-to-point links, each LSP is acknowledged by a partial sequence number PDU (PSNP), but on broadcast links, a complete sequence number PDU (CSNP) is sent out over the network. Any router that finds newer LSP information in the CSNP then purges the out-of-date entry and updates the link-state database.

LSPs support variable-length subnet mask addressing.

Complete Sequence Number PDU

The complete sequence number PDU (CSNP) lists all the link-state PDUs (LSPs) in the link-state database of the local router. Contained within the CSNP is an LSP identifier, a lifetime, a sequence number, and a checksum for each entry in the database. Periodically, a CSNP is sent on both broadcast and point-to-point links to maintain a correct database. Also, the advertisement of CSNPs occurs when an adjacency is formed with another router. Like IS-IS hello PDUs, CSNPs come in two types: Level 1 and Level 2.

When a Services Router receives a CSNP, it checks the database entries against its own local link-state database. If it detects missing information, the router requests specific LSP details using a partial sequence number PDU (PSNP).

Partial Sequence Number PDU

A partial sequence number PDU (PSNP) is used by an IS-IS router to request LSP information from a neighboring router. A PSNP can also explicitly acknowledge the receipt of an LSP on a point-to-point link. On a broadcast link, a CSNP is used as implicit knowledge. Like hello PDUs and CSNPs, the PSNP also has two types: Level 1 and Level 2.

When a Services Router compares a CSNP to its local database and determines that an LSP is missing, the router issues a PSNP for the missing LSP, which is returned in a link-state PDU from the router sending the CSNP. The received LSP is then stored in the local database, and an acknowledgement is sent back to the originating router.

BGP Overview

The Border Gateway Protocol (BGP) is an exterior gateway protocol (EGP) used primarily to establish point-to-point connections and transmit data between peer ASs. Unlike the IGP's RIP, OSPF and IS-IS, BGP must explicitly advertise the routes between its peers. The route advertisements determine prefix reachability and the way packets are routed between BGP neighbors. Because BGP uses the packet path to determine route selection, it is considered a path-vector protocol.


This overview contains the following topics:

- Point-to-Point Connections on page 371
- BGP Messages for Session Establishment on page 372
- BGP Messages for Session Maintenance on page 372
- IBGP and EBGP on page 372
- Route Selection on page 373
- Local Preference on page 374
- AS Path on page 375
- Origin on page 376
- Multiple Exit Discriminator on page 376
- Scaling BGP for Large Networks on page 378

Point-to-Point Connections

To establish point-to-point connections between peer ASs, you configure a BGP session on each interface of a point-to-point link. Figure 57 on page 371 shows an example of a BGP peering session.

Figure 57: BGP Peering Session



In Figure 57 on page 371, Router A is a gateway router for AS 3, and Router B is a gateway router for AS 10. For traffic internal to either AS, an IGP (OSPF, for instance) is used. To route traffic between peer ASs, a BGP session is used.

BGP Messages for Session Establishment

When the routers on either end of a BGP session first boot, the session between them is in the **Idle** state. The BGP session remains idle until a start event is detected. Typically, the start event is the configuration of a new BGP session or the resetting of an existing BGP session. At boot time, the start event is generated by the router as the BGP session is initiated.

After it detects a start event, the BGP host sends TCP request packets to its configured BGP neighbors. These packets are directed only to neighboring interfaces that have been explicitly configured as BGP neighbors. Upon receipt of the TCP request packet, the neighboring host generates a TCP response to complete the three-way handshake and establish a TCP connection between the peers. While this handshake is taking place, the BGP state for the connection is **Connect**. If a TCP timeout occurs while the originating host is waiting for a TCP response packet, the BGP state for the connection is **Active**. The **Active** state indicates that the router is actively listening for a TCP response and the TCP retry timer has been initiated.

Once a TCP connection has been established between both ends of a BGP session, the BGP session state is **OpenSent**, indicating that the originating router has generated an open message. The open message is an initial BGP handshake that must occur before any route advertisement can take place. Upon receipt of the open message, the neighboring router generates a keepalive message. Receipt of the keepalive message establishes a point-to-point connection, and the BGP session state transitions to **Established**. While the originating host waits for the keepalive response packet, the BGP session state is **OpenConfirm**.

BGP Messages for Session Maintenance

Once a BGP session has been established, the BGP peers exchange route advertisements by means of update messages. Update messages contain a one or more route advertisements, and they can contain one or more prefixes that are to be removed from the BGP routing table. If the peers need to advertise multiple routes, they generate and send multiple update messages as they detect changes to the network. In the absence of changes to the routing table, no update messages are generated.

While a BGP session is active, each router on the BGP session generates keepalive messages periodically. The timing of these messages is determined by the hold time on the session. The hold time is a negotiated value specifying the number of seconds that can elapse without keepalive messages before BGP designates the link inactive. Three messages are sent during every hold time interval.

When a peer connection is closed (either by error or if the BGP session is closed), a notification message is generated and sent to the peer router that did not experience the error or did not terminate the BGP session.

IBGP and EBGp

BGP uses two primary modes of information exchange, internal BGP (IBGP) and external BGP (EBGP), to communicate with internal and external peers, respectively.

Peer ASs establish links through an external peer BGP session. As a result, all route advertisement between the external peers takes place by means of the EBGP mode of information exchange. To propagate the routes through the AS and advertise them to internal peers, BGP uses IBGP. To advertise the routes to a different peer AS, BGP again uses EBGP.

To avoid routing loops, IBGP does not advertise routes learned from an internal BGP peer to other internal BGP peers. For this reason, BGP cannot propagate routes throughout an AS by passing them from one router to another. Instead, BGP requires that all internal peers be fully meshed so that any route advertised by one router is advertised to all peers within the AS.

As a network grows, the full mesh requirement becomes increasingly difficult to manage. In a network with 1000 routers, the addition of a single router requires that all the routers in the network be modified to account for the new addition. To combat these scaling problems, BGP uses route reflection and BGP confederations.

For information about route reflection, see “Scaling BGP for Large Networks” on page 378. For information about routing confederations, see “Scaling BGP for Large Networks” on page 378.

Route Selection

The BGP route selection process compares BGP attributes to select a single best path or active route for each prefix in the routing table. The attributes are compared in a particular order. A local BGP router uses the following criteria, in the order presented, to select a route from the routing table for the forwarding table:


1. Next-hop accessibility—If the next hop is inaccessible, the local router does not consider the route. The router must verify that it has a route to the BGP next-hop address. If a local route to the next hop does not exist, the local route does not include the router in its forwarding table. If such a route exists, route selection continues.
2. Highest local preference—The local router selects the route with the highest local preference value. If multiple routes have the same preference, route selection continues. (For more information, see “Local Preference” on page 374.)
3. Shortest AS path—The local router selects the route with the fewest entries in the AS path. If multiple routes have the same AS path length, route selection continues. (For more information, see “AS Path” on page 375.)
4. Lowest origin—The local router selects the route with the lowest origin value. If multiple routes have the same origin value, route selection continues. (For more information, see “Origin” on page 376.)
5. Lowest MED value—The local router selects the route with the lowest multiple exit discriminator (MED) value, comparing the routes from the same AS only. If multiple routes have the same MED value, route selection continues. For more information, see “Multiple Exit Discriminator” on page 376.
6. Strictly external paths—The local router prefers strictly external (EBGP) paths over external paths learned through interior sessions (IBGP). If multiple routes have the same strictly external paths, route selection continues.

7. Lowest IGP route metric—The local router selects the path for which the next hop is resolved through the IGP route with the lowest metric. If multiple routes have the same IGP route metric, route selection continues.
8. Maximum IGP next hops—The local router selects the path for which the BGP next hop is resolved through the IGP route with the largest number of next hops. If multiple routes have the same number of next hops, route selection continues.
9. Shortest route reflection cluster list—The local router selects the path with the shortest route reflection cluster list. Routes without a cluster list are considered to have a cluster list of length 0. If multiple routes have the same route reflection cluster list, route selection continues.
10. Lowest router ID—The local router selects the route with the lowest IP address value for the BGP router ID. By default, the router IDs of routes received from different ASs are not compared. You can change this default behavior. For more information, see the *JUNOS Routing Protocols Configuration Guide*.
11. Lowest peer IP address—The local router selects the path that was learned from the neighbor with the lowest peer IP address.

You can change the default behavior of some attributes (such as MED and router ID) used in the route selection process. For more information, see the *JUNOS Routing Protocols Configuration Guide*.

Local Preference

The local preference is typically used to direct all outbound AS traffic to a certain peer. When you configure a local preference, all routes that are advertised through that peer are assigned the preference value. The preference is a numeric value, and higher values are preferred during BGP route selection. Figure 58 on page 375 illustrates how to use local preference to determine BGP route selection.

Figure 58: Local Preference


The network in Figure 58 on page 375 shows two possible routes to the prefixes accessible through Host E. The first route, through Router A, uses an OC3 link to Router C and is then forwarded to Host E. The second route, through Router B, uses an OC48 link to Router D and is then forwarded to Host E. Although the number of hops to Host E is identical regardless of the route selected, the route through Router B is more desirable because of the increased bandwidth. To force traffic through Router B, you can set the local preference on Router A to **100** and the local preference on Router B to **300**. During BGP route selection, the route with the higher local preference is selected.



NOTE: In contrast to almost every other metric associated with dynamic routing protocols, the local preference gives higher precedence to the larger value.

AS Path

Routes advertised by BGP maintain a list of the ASs through which the route travels. This information is stored in the route advertisement as the AS path, and it is one of the primary criteria that a local router uses to evaluate BGP routes for inclusion in its forwarding table. Figure 59 on page 376 shows how BGP creates an AS path.

Figure 59: BGP AS Path

In the network shown in Figure 59 on page 376, the route from Host A to Host B travels through two intermediate ASs. As the route advertisement is propagated through the BGP network, it accumulates an AS path number each time it exits one AS and enters another. Each AS number is prepended to the AS path, which is stored as part of the route advertisement. When the route advertisement first leaves Host B's AS, the AS path is 17. When the route is advertised between intermediate ASs, the AS number 7 is prepended to the AS path, which becomes 7 17. When the route advertisement exits the third AS, the AS path becomes 4 7 17. The route with the shortest AS path is preferred for inclusion into the BGP forwarding table.

Origin

The BGP router that first advertises a route assigns it of the following values to identify its origin. During route selection, the lowest origin value is preferred.

- 0—The router originally learned the route through an IGP (OSPF, IS-IS, or a static route).
- 1—The router originally learned the route through an EGP (most likely BGP).
- 2—The route's origin is unknown.

Multiple Exit Discriminator

A multiple exit discriminator (MED) is an arbitrary metric assigned to a route to determine the exit point to a destination when all other factors are equal. By default, MED metrics are compared only for routes to the same peer AS, but you can also configure routing table path selection options for different ways of comparing MEDs.

Default MED Usage

Because the AS path rather than the number of hops between hosts is the primary criterion for BGP route selection, an AS with multiple connections to a peer AS can have multiple equivalent AS paths. When the routing table contains two routes to the same host in a neighboring AS, a multiple exit discriminator (MED) metric assigned to each route can determine which to include in the forwarding table. The MED metric you assign can force traffic through a particular exit point in an AS.

Figure 60 on page 377 illustrates how MED metrics are used to determine route selection.


Figure 60: Default MED Example

Figure 60 on page 377 shows AS 1 and AS 2 connected by two separate BGP links to Routers C and D. Host E in AS 1 is located nearer Router C. Host F, also in AS 1, is located nearer Router D. Because the AS paths are equivalent, two routes exist for each host, one through Router C and one through Router D. To force all traffic destined for Host E through Router C, network administrator for AS 2 assigns an MED metric for each router to Host E at its exit point. An MED metric of 10 is assigned to the route to Host E through Router C, and an MED metric of 20 is assigned to the route to Host E through Router D. BGP routers in AS 2 then select the route with the lower MED metric for the forwarding table.

Additional MED Options for Path Selection

By default, only the MEDs of routes that have the same peer ASs are compared. However, you can configure the routing table path selection options listed in Table 100 on page 377 to compare MEDs in different ways. The MED options are not mutually exclusive and can be configured in combination or independently. For the MED options to take effect, you must configure them uniformly all through your network. The MED option or options you configure determine the route selected. Thus we recommend that you carefully evaluate your network for preferred routes before configuring the MED options. For information about configuring the MED options, see the *JUNOS Routing Protocols Configuration Guide*.

Table 100: MED Options for Routing Table Path Selection

Option (Name)	Function	Use
Always comparing MEDs (always-compare-med)	Ensures that the MEDs for paths from peers in different ASs are always compared in the route selection process	Useful when all enterprises participating in a network agree on a uniform policy for setting MEDs. For example, in a network shared by two ISPs, both must agree that a certain path is the better path to configure the MED values correctly.

Table 100: MED Options for Routing Table Path Selection *(continued)*

Option (Name)	Function	Use
Adding IGP cost to MED (med-plus-igp)	<p>Before comparing MED values for path selection, adds to the MED the cost of the IGP route to the BGP next-hop destination.</p> <p>This option replaces the MED value for the router, but does not affect the IGP metric comparison. As a result, when multiple routes have the same value after the MED-plus-IGP comparison, and route selection continues, the IGP route metric is also compared, even though it was added to the MED value and compared earlier in the selection process.</p>	Useful when the downstream AS requires the complete cost of a certain route that is received across multiple ASs.
Applying Cisco IOS nondeterministic behavior (cisco-non-deterministic)	<p>Specifies the nondeterministic behavior of the Cisco IOS software:</p> <ul style="list-style-type: none"> ■ The active path is always first. All nonactive but eligible paths follow the active path and are maintained in the order in which they were received. Ineligible paths remain at the end of the list. ■ When a new path is added to the routing table, path comparisons are made among all routes, including those paths that must never be selected because they lose the MED tie-breaking rule. 	We recommend that you do not configure this option, because the nondeterministic behavior sometimes prevents the system from properly comparing the MEDs between paths.

Scaling BGP for Large Networks

BGP is not a flooding protocol like RIP or OSPF. Instead, it is a peering protocol that exchanges routes with fully meshed peers only. However, in large networks, the full mesh requirement causes scaling problems. BGP combats scaling problems with the following methods:

- Route Reflectors—for Added Hierarchy on page 378
- Confederations—for Subdivision on page 381

Route Reflectors—for Added Hierarchy

To use route reflection in an AS, you designate one or more routers as a route reflector—typically, one per point of presence (POP). Route reflectors have the special BGP ability to readvertise routes learned from an internal peer to other internal peers. So rather than requiring all internal peers to be fully meshed with each other, route reflection requires only that the route reflector be fully meshed with all internal peers. The route reflector and all its internal peers form a cluster, as shown in Figure 61 on page 379.



NOTE: You must have an Advanced BGP Feature license installed on each Services Router that uses a route reflector. For license details, see the *J-series Services Router Administration Guide*.

Figure 61: Simple Route Reflector Topology (One Cluster)




Figure 61 on page 379 shows Router RR configured as the route reflector for Cluster 127. The other routers are designated internal peers within the cluster. BGP routes are advertised to Router RR by any of the internal peers. RR then readvertises those routes to all other peers within the cluster.

You can configure multiple clusters and link them by configuring a full mesh of route reflectors (see Figure 62 on page 380).


Figure 62: Basic Route Reflection (Multiple Clusters)

Figure 62 on page 380 shows Route Reflectors RR1, RR2, RR3, and RR4 as fully meshed internal peers. When a router advertises a route to Reflector RR1, RR1 readvertises the route to the other route reflectors, which, in turn, readvertise the route to the remaining routers within the AS. Route reflection allows the route to be propagated throughout the AS without the scaling problems created by the full mesh requirement.

However, as clusters become large, a full mesh with a route reflector becomes difficult to scale, as does a full mesh between route reflectors. To help offset this problem, you can group clusters of routers together into clusters of clusters for hierarchical route reflection (see Figure 63 on page 380).

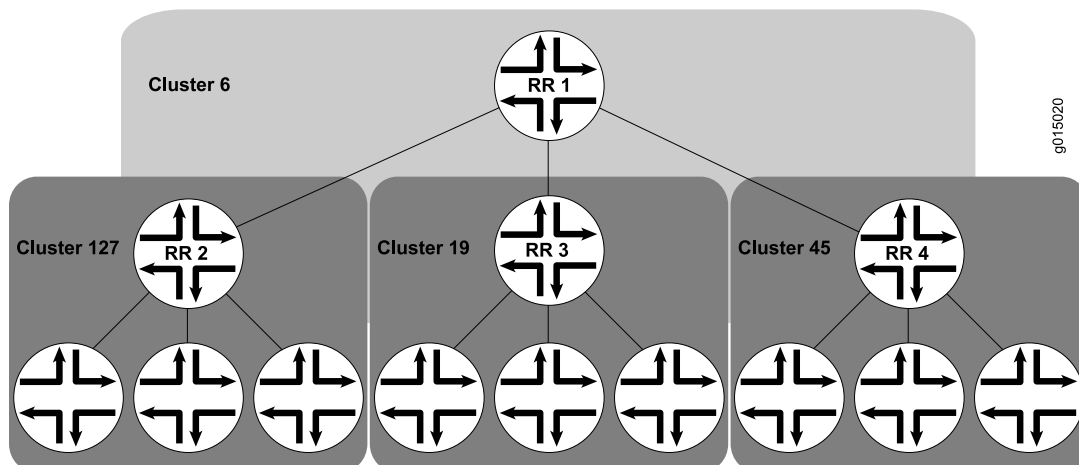
Figure 63: Hierarchical Route Reflection (Clusters of Clusters)

Figure 63 on page 380 shows RR2, RR3, and RR4 as the route reflectors for Clusters 127, 19, and 45, respectively. Rather than fully mesh those route reflectors, the network administrator has configured them as part of another cluster (Cluster 6) for which RR1 is the route reflector. When a router advertises a route to RR2, RR2 readvertises the route to all the routers within its own cluster, and then readvertises the route to RR1. RR1 readvertises the route to the routers in its cluster, and those routers propagate the route down through their clusters.

Confederations—for Subdivision

BGP confederations are another way to solve the scaling problems created by the BGP full mesh requirement. BGP confederations effectively break up a large AS into subautonomous systems (sub-ASs). Each sub-AS must be uniquely identified within the confederation AS by a sub-AS number. Typically, sub-AS numbers are taken from the private AS numbers between 64512 and 65535.

Within a sub-AS, the same IBGP full mesh requirement exists. Connections to other confederations are made with standard EBGP, and peers outside the sub-AS are treated as external. To avoid routing loops, a sub-AS uses a confederation sequence, which operates like an AS path but uses only the privately assigned sub-AS numbers.

The confederation AS appears whole to other confederation ASs. The AS path received by other ASs shows only the globally assigned AS number. It does not include the confederation sequence or the privately assigned sub-AS numbers. The sub-AS numbers are removed when the route is advertised out of the confederation AS. Figure 64 on page 381 shows an AS divided into four confederations.

Figure 64: BGP Confederations




Figure 64 on page 381 shows AS 3 divided into four sub-ASs, 64517, 64550, 65300, and 65410, which are linked through EBGP sessions. Because the confederations are connected by EBGP, they do not need to be fully meshed. EBGP routes are readvertised to other sub-ASs.

Chapter 12

Configuring Static Routes

Static routes are routes that you explicitly enter into the routing table as permanent additions. Traffic through static routes is always routed the same way.

You can use either J-Web Quick Configuration or a configuration editor to configure static routes.

This chapter contains the following topics. For more information about static routes, see the *JUNOS Routing Protocols Configuration Guide*.

- Static Routing Overview on page 383
- Before You Begin on page 385
- Configuring Static Routes with Quick Configuration on page 386
- Configuring Static Routes with a Configuration Editor on page 387
- Verifying the Static Route Configuration on page 393

Static Routing Overview

Routes that are permanent fixtures in the routing and forwarding tables are often configured as static routes. These routes generally do not change, and often include only one or very few paths to the destination.

To create a static route in the routing table, you must, at minimum, define the route as static and associate a next-hop address with it. The static route in the routing table is inserted into the forwarding table when the next-hop address is reachable. All traffic destined for the static route is transmitted to the next-hop address for transit.

This overview contains the following topics:

- Static Route Preferences on page 384
- Qualified Next Hops on page 384
- Control of Static Routes on page 384
- Default Properties on page 385

Static Route Preferences

A static route destination address can have multiple next hops associated with it. In this case, multiple routes are inserted into the routing table, and route selection must occur. Because the primary criterion for route selection is the route preference, you can control the routes that are used as the primary route for a particular destination by setting the route preference associated with a particular next hop. The routes with a higher preference are always used to route traffic. When you do not set a preferred route, traffic is alternated between routes in round-robin fashion.

Qualified Next Hops

In general, the default properties assigned to a static route apply to all the next-hop addresses configured for the static route. If, however, you want to configure two possible next-hop addresses for a particular route and have them treated differently, you can define one as a qualified next hop.

Qualified next hops allow you to associate one or more properties with a particular next-hop address. You can set an overall preference for a particular static route and then specify a different preference for the qualified next hop. For example, suppose two next-hop addresses (10.10.10.10 and 10.10.10.7) are associated with the static route 192.168.47.5/32. A general preference is assigned to the entire static route, and then a different preference is assigned to only the qualified next-hop address 10.10.10.7. For example:

```
route 192.168.47.5/32 {
  next-hop 10.10.10.10;
  qualified-next-hop 10.10.10.7 {
    preference 2;
  }
  preference 6;
}
```

In this example, the qualified next hop 10.10.10.7 is assigned the preference 2, and the next-hop 10.10.10.10 is assigned the preference 6.

Control of Static Routes

You can control the importation of static routes into the routing and forwarding tables in a number of ways. Primary ways include assigning one or more of the following attributes to the route:

- **retain**—Keeps the route in the forwarding table after the routing process shuts down or the Services Router reboots. For more information, see “Route Retention” on page 385.
- **no-readvertise**—Prevents the route from being readvertised to other routing protocols. For more information, see “Readvertisement Prevention” on page 385.
- **passive**—Rejects traffic destined for the route. For more information, see “Forced Rejection of Passive Route Traffic” on page 385.

Route Retention

By default, static routes are not retained in the forwarding table when the routing process shuts down. When the routing process starts up again, any routes configured as static routes must be added to the forwarding table again. To avoid this latency, routes can be flagged as **retain**, so that they are kept in the forwarding table even after the routing process shuts down. Retention ensures that the routes are always in the forwarding table, even immediately after a system reboot.

Readvertisement Prevention

Static routes are eligible for readvertisement by other routing protocols by default. In a stub area where you might not want to readvertise these static routes under any circumstances, you can flag the static routes as **no-readvertise**.

Forced Rejection of Passive Route Traffic

Generally, only active routes are included in the routing and forwarding tables. If a static route's next-hop address is unreachable, the route is marked **passive**, and it is not included in the routing or forwarding tables. To force a route to be included in the routing tables regardless of next-hop reachability, you can flag the route as **passive**. If a route is flagged **passive** and its next-hop address is unreachable, the route is included in the routing table and all traffic destined for the route is rejected.

Default Properties

The basic configuration of static routes defines properties for a particular route. To define a set of properties to be used as defaults on all static routes, set those properties as default values. For example:

```
defaults {
    retain;
    no-readvertise;
    passive;
}
route 0.0.0.0/0 next-hop 192.168.1.1;
route 192.168.47.5/32 {
    next-hop 10.10.10.10;
    qualified-next-hop 10.10.10.7 {
        preference 6;
    }
    preference 2;
}
```

In this example, the **retain**, **no-readvertise**, and **passive** attributes are set as defaults for all static routes. If any local setting for a particular route conflicts with the default values, the local setting supersedes the default.

Before You Begin

Before you begin configuring static routes, complete the following tasks:

- Establish basic connectivity. See the Getting Started Guide for your router.
- Configure network interfaces. See “Configuring Ethernet, DS1, DS3, and Serial Interfaces” on page 103.

Configuring Static Routes with Quick Configuration

J-Web Quick Configuration allows you to configure static routes. Figure 65 on page 386 shows the Quick Configuration Routing page for static routing.

Figure 65: Quick Configuration Routing Page for Static Routing

Juniper
NETWORKS

ROUTER - J6300

Monitor **Configuration** Diagnose Manage Events Alarms Logged in as: regress Help About Logout

[Configuration](#) > [Quick Configuration](#) > [Routing and Protocols](#)

Quick Configuration

Routing and Protocols

Default Route

Default Route

Static Routes

	Static Route Address	Next Hop
<input type="checkbox"/>	172.16.0.0/12	10.209.63.254
<input type="checkbox"/>	192.168.0.0/16	10.209.63.254
<input type="checkbox"/>	207.17.136.192/32	10.209.63.254
<input type="checkbox"/>	10.10.0.0/16	10.209.63.254
<input type="checkbox"/>	10.5.0.0/16	10.209.63.254
<input type="checkbox"/>	192.168.102.0/23	10.209.63.254
<input type="checkbox"/>	207.17.136.0/24	10.209.63.254
<input type="checkbox"/>	10.209.0.0/16	10.209.63.254
<input type="checkbox"/>	10.150.0.0/16	10.209.63.254
<input type="checkbox"/>	10.157.64.0/19	10.209.63.254

Add... Delete

OK Cancel Apply

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. Trademark Notice. Privacy. Juniper your Net.

To configure static routes with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Quick Configuration > Routing and Protocols**.
2. Enter information into the Static Routing Quick Configuration page, as described in Table 101 on page 387.

3. From the main static routing Quick Configuration page, click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration Routing page for static routing, click **Apply**.
 - To apply the configuration and return to the Quick Configuration Routing page, click **OK**.
 - To cancel your entries and return to the Quick Configuration Routing page, click **Cancel**.
4. To check the configuration, see “Verifying the Static Route Configuration” on page 393.

Table 101: Static Routing Quick Configuration Summary

Field	Function	Your Action
Default Route		
Default Route	Specifies the default gateway for the router.	Type the 32-bit IP address of the Services Router's default route in dotted decimal notation.
Static Routes		
Static Route Address (required)	Specifies the static route to add to the routing table.	<ol style="list-style-type: none"> 1. On the main static routing Quick Configuration page, click Add. 2. In the Static Route Address box, type the 32-bit IP address of the static route in dotted decimal notation.
Next-Hop Addresses	Specifies the next-hop address or addresses to be used when routing traffic to the static route.	<ol style="list-style-type: none"> 1. In the Add box, type the 32-bit IP address of the next-hop host. 2. Click Add. 3. Add more next-hop addresses as necessary. <p>NOTE: If a route has multiple next-hop addresses, traffic is routed across each address in round-robin fashion.</p> <ol style="list-style-type: none"> 4. When you have finished adding next-hop addresses, click OK.

Configuring Static Routes with a Configuration Editor

To configure static routes on the Services Router, you must perform the following tasks marked *(Required)*.

- Configuring a Basic Set of Static Routes (Required) on page 388
- Controlling Static Route Selection (Optional) on page 389


- Controlling Static Routes in the Routing and Forwarding Tables (Optional) on page 391
- Defining Default Behavior for All Static Routes (Optional) on page 392

For information about using the J-Web and CLI configuration editors, see “Using Services Router Configuration Tools” on page 3.

Configuring a Basic Set of Static Routes (Required)

Customer routes that are connected to stub networks are often configured as static routes. Figure 66 on page 388 shows a sample network.

Figure 66: Customer Routes Connected to a Stub Network



To configure customer routes as static routes, like the ones in Figure 66 on page 388, follow these steps on the Services Router to which the customer routes are connected:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 102 on page 389.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To manually control static route selection, see “Controlling Static Route Selection (Optional)” on page 389.
 - To determine how static routes are imported into the routing and forwarding tables, see “Controlling Static Routes in the Routing and Forwarding Tables (Optional)” on page 391.
 - To define default properties for static routes, see “Defining Default Behavior for All Static Routes (Optional)” on page 392.
 - To check the configuration, see “Verifying the Static Route Configuration” on page 393.


Table 102: Configuring Basic Static Routes

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Static level in the configuration hierarchy.	<ol style="list-style-type: none">1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration.2. Next to Routing options, click Configure or Edit.3. Next to Static, click Configure or Edit.	From the [edit] hierarchy level, enter edit routing-options static
Add the static route 192.168.47.5/32, and define the next-hop address 10.10.10.10.	<ol style="list-style-type: none">1. Next to Route, click Add new entry.2. In the Destination box, type 192.168.47.5/32.3. From the Next hop list, select Next hop.4. Next to Next hop, click Add new entry.5. In the Value box, type 10.10.10.10.6. Click OK.	Define the static route and set the next-hop address: set route 192.168.47.5 next-hop 10.10.10.10

Controlling Static Route Selection (Optional)

When multiple next hops exist for a single static route (see Figure 67 on page 389), you can specify how traffic is to be routed to the destination.

Figure 67: Controlling Static Route Selection



In this example, the static route 192.168.47.5/32 has two possible next hops. Because of the links between those next-hop hosts, host 10.10.10.7 is the preferred path. To

configure the static route **192.168.47.5/32** with two next hops and give preference to host **10.10.10.7**, follow these steps:

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 103 on page 390.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To determine how static routes are imported into the routing and forwarding tables, see “Controlling Static Routes in the Routing and Forwarding Tables (Optional)” on page 391.
 - To define default properties for static routes, see “Defining Default Behavior for All Static Routes (Optional)” on page 392.
 - To check the configuration, see “Verifying the Static Route Configuration” on page 393.

Table 103: Controlling Static Route Selection

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Static level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Routing options, click Configure or Edit. 3. Next to Static, click Configure or Edit. 	From the [edit] hierarchy level, enter edit routing-options static
Add the static route 192.168.47.5/32 , and define the next-hop address 10.10.10.10 .	<ol style="list-style-type: none"> 1. Next to Route, click Add new entry. 2. In the Destination box, type 192.168.47.5/32. 3. From the Next hop list, select Next hop. 4. In the Next hop box, click Add new entry. 5. In the Value box, type 10.10.10.10. 6. Click OK. 	Define the static route and set the next-hop address: set route 192.168.47.5 next-hop 10.10.10.10
Set the preference for the 10.10.10.10 next hop to 7 .	<ol style="list-style-type: none"> 1. Next to Preference, select the Yes check box. 2. Click Configure. 3. In the Metric value box, type 7. 4. Click OK. 	Set the preference to 7: set route 192.168.47.5 next-hop 10.10.10.10 preference 7
Define the qualified next-hop address 10.10.10.7 .	<ol style="list-style-type: none"> 1. Next to Qualified next hop, click Add new entry. 2. In the Nexthop box, type 10.10.10.7. 	Set the qualified-next-hop address: set route 192.168.47.5 qualified-next-hop 10.10.10.7

Table 103: Controlling Static Route Selection (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
Set the preference for the 10.10.10.7 qualified next hop to 6.	<ol style="list-style-type: none"> 1. In the Preference box, type 6. 2. Click OK. 	Set the preference to 6: set route 192.168.47.5 qualified-next-hop 10.10.10.7 preference 6

Controlling Static Routes in the Routing and Forwarding Tables (Optional)

Static routes have a number of attributes that define how they are inserted and maintained in the routing and forwarding tables. To customize this behavior for the static route 192.168.47.5/32, perform these steps:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 104 on page 391.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To define default properties for static routes, see “Defining Default Behavior for All Static Routes (Optional)” on page 392.
 - To check the configuration, see “Verifying the Static Route Configuration” on page 393.

Table 104: Controlling Static Routes in the Routing and Forwarding Tables

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the 192.168.47.5/32 level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Routing options, click Edit. 3. Next to Static, click Edit. 4. Under Route and Destination, click 192.168.47.5/32. 	From the [edit] hierarchy level, enter edit routing-options static route 192.168.47.5/32
Specify that the route is to be retained in the forwarding table after the routing process shuts down. By default, static routes are not retained.	Next to Retain, select the Yes check box.	Set the retain attribute: set retain
Specify that the static route is not to be readvertised. By default, static routes are eligible to be readvertised.	Next to Readvertise, select the No check box.	Set the no-readvertise attribute: set no-readvertise

Table 104: Controlling Static Routes in the Routing and Forwarding Tables *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
Specify that the static route is to be included in the routing table whether the route is active or not. By default, passive routes are not included in the routing table.	<ol style="list-style-type: none"> From the Passive flag list, select Passive. Click OK. 	Set the passive attribute: set passive

Defining Default Behavior for All Static Routes (Optional)

Attributes that define static route behavior can be configured either at the individual route level or as a default behavior that applies to all static routes. In the case of conflicting configuration, the configuration at the individual route level overrides static route defaults. To configure static route defaults, perform these steps:

- Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- Perform the configuration tasks described in Table 105 on page 392.
- If you are finished configuring the router, commit the configuration.
- To check the configuration, see “Verifying the Static Route Configuration” on page 393.

Table 105: Defining Static Route Defaults

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Defaults level in the configuration hierarchy.	<ol style="list-style-type: none"> In the J-Web interface, select Configuration > View and Edit > Edit Configuration. Next to Routing options, click Edit. Next to Static, click Edit. Next to Defaults, click Configure. 	From the [edit] hierarchy level, enter edit routing-options static defaults
Specify that the route is to be retained in the forwarding table after the routing process shuts down. By default, static routes are not retained.	<ol style="list-style-type: none"> Next to Retain, select the Yes check box. Click OK. 	Set the retain attribute: set retain
Specify that the static route is not to be readvertised. By default, static routes are eligible to be readvertised.	<ol style="list-style-type: none"> Next to Readvertise, select the No check box. Click OK. 	Set the no-readvertise attribute: set no-readvertise
Specify that the static route is to be included in the routing table whether the route is active or not. By default, passive routes are not included in the routing table.	<ol style="list-style-type: none"> From the Passive flag list, select Passive. Click OK. 	Set the passive attribute: set passive

Verifying the Static Route Configuration

Verify that the static routes are in the routing table and that those routes are active.

Displaying the Routing Table

Purpose Verify static route configuration as follows by displaying the routing table and checking its contents.

Action From the CLI, enter the `show route terse` command.

```
user@host> show route terse
inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination      P Prf  Metric 1  Metric 2  Next hop      AS path
* 192.168.47.5/32   S   5          Reject
* 172.16.0.0/12     S   5          >192.168.71.254
* 192.168.0.0/18     S   5          >192.168.71.254
* 192.168.40.0/22    S   5          >192.168.71.254
* 192.168.64.0/18    S   5          >192.168.71.254
* 192.168.64.0/21    D   0          >fxp0.0
* 192.168.71.246/32  L   0          Local
* 192.168.220.4/30   D   0          >ge-0/0/1.0
* 192.168.220.5/32   L   0          Local
* 192.168.220.8/30   D   0          >ge-0/0/2.0
* 192.168.220.9/32   L   0          Local
* 192.168.220.12/30  D   0          >ge-0/0/3.0
* 192.168.220.13/32  L   0          Local
* 192.168.220.17/32  L   0          Reject
* 192.168.220.21/32  L   0          Reject
* 192.168.220.24/30  D   0          >at-1/0/0.0
* 192.168.220.25/32  L   0          Local
* 192.168.220.28/30  D   0          >at-1/0/1.0
* 192.168.220.29/32  L   0          Local
* 224.0.0.9/32       R 100         1          MultiRecv
```

- What It Means** The output shows a list of the routes that are currently in the `inet.0` routing table. Verify the following information:
- Each configured static route is present. Routes are listed in ascending order by IP address. Static routes are identified with an **S** in the protocol (**P**) column of the output.
 - Each static route is active. Routes that are active show the next-hop IP address in the **Next hop** column. If a route's next-hop address is unreachable, the next-hop address is identified as **Reject**. These routes are not active routes, but they appear in the routing table because the **passive** attribute is set.
 - The preference for each static route is correct. The preference for a particular route is listed in the **Prf** column of the output.

See Also For a complete description of `show route terse` output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Chapter 13

Configuring a RIP Network

The Routing Information Protocol (RIP) is an interior gateway protocol that routes packets within a single autonomous system (AS). To use RIP, you must understand the basic components of a RIP network and configure the J-series Services Router to act as a node in the network.



NOTE: The J-series Services Router supports RIP version 1 (RIPv1) and RIP version 2 (RIPv2) only. Unless otherwise specified, the term *RIP* in this chapter refers to these versions of the protocol.

You can use either J-Web Quick Configuration or a configuration editor to configure a RIP network.

This chapter contains the following topics. For more information about RIP, see the *JUNOS Routing Protocols Configuration Guide*.

- RIP Overview on page 395
- Before You Begin on page 396
- Configuring a RIP Network with Quick Configuration on page 396
- Configuring a RIP Network with a Configuration Editor on page 398
- Verifying the RIP Configuration on page 406

RIP Overview

To achieve basic connectivity between all RIP hosts in a RIP network, you enable RIP on every interface that is expected to transmit and receive RIP traffic. To enable RIP on an interface, you define RIP groups, which are logical groupings of interfaces, and add interfaces to the groups. Additionally, you must configure a routing policy to export directly connected routes and routes learned through RIP routing exchanges.

RIP Traffic Control with Metrics

To tune a RIP network and control traffic flowing through the network, you increase or decrease the cost of the paths through the network. RIP provides two ways to modify the path cost: an incoming metric and an outgoing metric, which are each set to 1 by default. These metrics are attributes that manually specify the cost of any

route advertised through a host. By increasing or decreasing the metrics—and thus the cost—of links throughout the network, you can control packet transmission across the network.

The incoming metric modifies the cost of an individual segment when a route across the segment is imported into the routing table. For example, if you set the incoming metric on the segment to **3**, the individual segment cost along the link is changed from **1** to **3**. The increased cost affects all route calculations through that link. Other routes that were previously excluded because of a high hop count might now be selected into the router's forwarding table.

The outgoing metric modifies the path cost for all the routes advertised out a particular interface. Unlike the incoming metric, the outgoing metric modifies the routes that other routers are learning and thereby controls the way they send traffic.

Authentication

RIPv2 provides authentication support so that RIP links can require authentication keys (passwords) before they become active. These authentication keys can be specified in either plain-text or MD5 form. Authentication provides an additional layer of security on the network beyond the other security features.

This type of authentication is not supported on RIPv1 networks.

Before You Begin


Before you begin configuring a RIP network, complete the following tasks:

- Establish basic connectivity. See the Getting Started Guide for your router.
- Configure network interfaces. See “Configuring Ethernet, DS1, DS3, and Serial Interfaces” on page 103.

Configuring a RIP Network with Quick Configuration

J-Web Quick Configuration allows you to create RIP networks. Figure 68 on page 397 shows the Quick Configuration Routing page for RIP.

Figure 68: Quick Configuration Routing Page for RIP



To configure a RIP network with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Quick Configuration > Routing and Protocols**.
2. Enter information into the Quick Configuration page for RIP, as described in Table 106 on page 397.
3. From the main RIP routing Quick Configuration page, click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration Routing page for RIP, click **Apply**.
 - To apply the configuration and return to the Quick Configuration Routing page, click **OK**.
 - To cancel your entries and return to the Quick Configuration Routing page, click **Cancel**.
4. To check the configuration, see “Verifying the RIP Configuration” on page 406.

Table 106: RIP Routing Quick Configuration Summary

Field	Function	Your Action
RIP		
Enable RIP	Enables or disables RIP.	<ul style="list-style-type: none">■ To enable RIP, select the check box.■ To disable RIP, clear the check box.

Table 106: RIP Routing Quick Configuration Summary (*continued*)

Field	Function	Your Action
Advertise Default Route	Advertises the default route using RIPv2.	<ul style="list-style-type: none"> ■ To advertise the default route using RIPv2, select the check box. ■ To disable the default route advertisement, clear the check box.
RIP-Enabled Interfaces	<p>Designates one or more Services Router interfaces on which RIP is enabled.</p> <p>For information about interface names, see “Network Interface Naming” on page 47.</p>	<p>The first time you configure RIP, the Logical Interfaces box displays a list of all the logical interfaces configured on the Services Router. Do any of the following:</p> <ul style="list-style-type: none"> ■ To enable RIP on an interface, click the interface name to highlight it, and click the left arrow to add the interface to the RIP interfaces list. ■ To enable RIP on multiple interfaces at once, press Ctrl while you click multiple interface names to highlight them. Then click the left arrow to add the interfaces to the RIP interfaces list. ■ To disable RIP on one or more interfaces, highlight the interface or interfaces in the RIP interfaces box and click the right arrow to move them back to the Logical Interfaces list.

Configuring a RIP Network with a Configuration Editor


To configure the Services Router as a node in a RIP network, you must perform the following task marked *(Required)*.

- Configuring a Basic RIP Network (Required) on page 398
- Controlling Traffic in a RIP Network (Optional) on page 401
- Enabling Authentication for RIP Exchanges (Optional) on page 404

For information about using the J-Web and CLI configuration editors, see “Using Services Router Configuration Tools” on page 3.

Configuring a Basic RIP Network (Required)

To use RIP on the Services Router, you must configure RIP on all the RIP interfaces within a network like the one shown in Figure 69 on page 399.

Figure 69: Typical RIP Network Topology

By default, RIP does not advertise the subnets that are directly connected through the Services Router's interfaces. For traffic to pass through a RIP network, you must create a routing policy to export these routes. Advertising only the direct routes propagates the routes to the immediately adjacent RIP-enabled router only. To propagate all routes through the entire RIP network, you must configure the routing policy to export the routes learned through RIP.

To configure a RIP network like the one in Figure 69 on page 399, with a routing policy, perform these steps on each Services Router in the network:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 107 on page 400.
3. If you are finished configuring the router, commit the configuration.

After you add the appropriate interfaces to the RIP group, RIP begins sending routing information. No additional configuration is required to enable RIP traffic on the network.

4. Go on to one of the following procedures:
 - To control RIP traffic on the network, see “Controlling Traffic in a RIP Network (Optional)” on page 401.
 - To authenticate RIP exchanges, see “Enabling Authentication for RIP Exchanges (Optional)” on page 404.
 - To check the configuration, see “Verifying the RIP Configuration” on page 406.

Table 107: Configuring a RIP Network

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Rip level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Configure or Edit. 3. Next to Rip, click Configure or Edit. 	<p>From the [edit] hierarchy level, enter</p> <pre>edit protocols rip</pre>
Create the RIP group alpha1 .	<ol style="list-style-type: none"> 1. Next to Group, click Add new entry. 2. In the Group name box, type alpha1. 	<ol style="list-style-type: none"> 1. Create the RIP group alpha1, and add an interface: <pre>set group alpha1 neighbor ge-0/0/0.0</pre>
<p>Add interfaces to the RIP group alpha1.</p> <p>For information about interface names, see “Network Interface Naming” on page 47.</p>	<ol style="list-style-type: none"> 1. Next to Neighbor, click Add new entry. 2. In the Neighbor name box, type the name of an interface on the Services Router—for example, ge-0/0/0.0—and click OK. 3. Repeat Step 2 for each interface on this Services Router that you are adding to the RIP group. Only one interface is required. 	<ol style="list-style-type: none"> 2. Repeat Step 1 for each interface on this Services Router that you are adding to the RIP group. Only one interface is required.
Configure a routing policy to advertise directly connected routes.	<ol style="list-style-type: none"> 1. On the main Configuration page next to Policy options, click Configure or Edit. 2. Next to Policy statement, click Add new entry. 3. In the Policy name box, type the name of the policy statement—for example, advertise-rip-routes. 4. Next to Term, click Add new entry. 5. In the Term name box, type the name of the policy statement—for example, from-direct. 6. Next to From, click Configure. 7. Next to Protocol, click Add new entry. 8. From the Value list, select Direct. 9. Click OK until you return to the Policy statement page. 10. Next to Then, click Configure. 11. From the Accept reject list, select Accept. 12. Click OK. 	<ol style="list-style-type: none"> 1. From the [edit] hierarchy level, enter <pre>edit policy-options</pre> 2. Set the match condition to match on direct routes: <pre>set policy-statement advertise-rip-routes term from-direct from protocol direct</pre> 3. Set the match action to accept these routes: <pre>set policy-statement advertise-rip-routes term from-direct then accept</pre>

Table 107: Configuring a RIP Network (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the previous routing policy to advertise routes learned from RIP.	<ol style="list-style-type: none"> 1. On the main Configuration page next to Policy options, click Configure or Edit. 2. Next to Policy statement, click advertise-rip-routes. 3. Next to Term, click Add new entry. 4. In the Term name box, type the name of the policy statement—for example, from-rip. 5. Next to From, click Configure. 6. Next to Protocol, click Add new entry. 7. From the Value list, select rip. 8. Click OK until you return to the Policy statement page. 9. Next to Then, click Configure. 10. From the Accept reject list, select Accept. 11. Click OK. 	<ol style="list-style-type: none"> 1. From the [edit] hierarchy level, enter edit policy-options 2. Set the match condition to match on direct routes: set policy-statement advertise-rip-routes term from-rip from protocol rip 3. Set the match action to accept these routes: set policy-statement advertise-rip-routes term from-rip then accept


Controlling Traffic in a RIP Network (Optional)

There are two primary means for controlling traffic in a RIP network: the incoming metric and the outgoing metric. To modify these attributes, see the following sections:

- Controlling Traffic with the Incoming Metric on page 401
- Controlling Traffic with the Outgoing Metric on page 403

Controlling Traffic with the Incoming Metric

Depending on the RIP network topology and the links between nodes in the network, you might want to control traffic flow through the network to maximize flow across higher-bandwidth links. Figure 70 on page 401 shows a network with alternate routes between Routers A and D.

Figure 70: Controlling Traffic in a RIP Network with the Incoming Metric

In this example, routes to Router D are received by Router A across both of its RIP-enabled interfaces. Because the route through Router B and the route through Router C have the same number of hops, both routes are imported into the forwarding table. However, because the T3 link from Router B to Router D has a higher bandwidth than the T1 link from Router C to Router D, you want traffic to flow from A through B to D.

To force this flow, you can modify the route metrics as they are imported into Router A's routing table. By setting the incoming metric on the interface from Router A to Router C, you modify the metric on all routes received through that interface. Setting the incoming route metric on Router A changes only the routes in Router A's routing table, and affects only how Router A sends traffic to Router D. Router D's route selection is based on its own routing table, which, by default, includes no adjusted metric values.

In the example, Router C receives a route advertisement from Router D and readvertises the route to Router A. When Router A receives the route, it applies the incoming metric on the interface. Instead of incrementing the metric by 1 (the default), Router A increments it by 3 (the configured incoming metric), giving the route from Router A to Router D through Router C a total path metric of 4. Because the route through Router B has a metric of 2, it becomes the preferred route for all traffic from Router A to Router D.

To modify the incoming metric on all routes learned on the link between Router A and Router C and force traffic through Router B:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 108 on page 402.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To authenticate RIP exchanges, see “Enabling Authentication for RIP Exchanges (Optional)” on page 404.
 - To check the configuration, see “Verifying the RIP Configuration” on page 406.

Table 108: Modifying the Incoming Metric

Task	J-Web Configuration Editor	CLI Configuration Editor
In the configuration hierarchy, navigate to the level of an interface in the alpha1 RIP group.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to Rip, click Edit. 4. Under Group name, click alpha1. 5. Under Neighbor name, click the interface name—for example, ge-0/0/0.0. 	<p>From the [edit] hierarchy level, enter</p> <pre>edit protocols rip group alpha1 neighbor ge-0/0/0</pre>


Table 108: Modifying the Incoming Metric (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Increase the incoming metric to 3.	In the Metric in box, type 3, and click OK .	Set the incoming metric to 3: set metric-in 3

Controlling Traffic with the Outgoing Metric

If an exported route was learned from a member of the same RIP group, the metric associated with that route is the normal RIP metric. For example, a RIP route with a metric of 5 learned from a neighbor configured with an incoming metric of 2 is advertised with a combined metric of 7 when advertised to neighbors in the same group. However, if this route was learned from a RIP neighbor in a different group or from a different protocol, the route is advertised with the metric value configured in the outgoing metric for that group. Figure 71 on page 403 shows a network with alternate routes between Routers A and D.

Figure 71: Controlling Traffic in a RIP Network with the Outgoing Metric



In this example, each route from Router A to Router D has two hops. However, because the link from Router A to Router B in RIP group Beta 1 has a higher bandwidth than the link from Router A to Router C in RIP group Alpha 1, you want traffic from Router D to Router A to flow through Router B. To control the way Router D sends traffic to Router A, you can alter the routes that Router D receives by configuring the outgoing metric on Router A's interfaces in the Alpha 1 RIP group.

If the outgoing metric for the Alpha 1 RIP group—the A-to-C link—is changed to 3, Router D calculates the total path metric from to A through C as 4. In contrast, the unchanged default total path metric to A through B in the Beta 1 RIP group is 2. The fact that Router A's interfaces belong to two different RIP groups allows you to configure two different outgoing metrics on its interfaces, because you configure path metrics at the group level.

By configuring the *incoming* metric, you control the way Router A sends traffic to Router D. By configuring the *outgoing* metric on the same router, you control the way Router D sends traffic to Router A.

To modify the outgoing metric on Router A and force traffic through Router B:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 109 on page 404.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To authenticate RIP exchanges, see “Enabling Authentication for RIP Exchanges (Optional)” on page 404.
 - To check the configuration, see “Verifying the RIP Configuration” on page 406.

Table 109: Modifying the Outgoing Metric

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the alpha1 level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to Rip, click Edit. 4. Under Group name, click alpha1. 	From the [edit] hierarchy level, enter edit protocols rip group alpha1
Increase the outgoing metric to 3.	In the Metric out box, type 3, and click OK .	Set the outgoing metric to 3: set metric-out 3

Enabling Authentication for RIP Exchanges (Optional)

All RIPv2 protocol exchanges can be authenticated to guarantee that only trusted routers participate in the AS's routing. By default, this authentication is disabled. Authentication requires all routers within the RIP network or subnetwork to have the same authentication type and key (password) configured.

You can enable RIP authentication exchanges by either of the following methods:

- Enabling Authentication with Plain-Text Passwords on page 404
- Enabling Authentication with MD5 Authentication on page 405

Enabling Authentication with Plain-Text Passwords

To configure authentication that requires a plain-text password to be included in the transmitted packet, enable simple authentication by performing these steps on all RIP Services Routers in the network:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 110 on page 405.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying the RIP Configuration” on page 406.

Table 110: Configuring Simple RIP Authentication

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to Rip level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to Rip, click Edit. 	From the [edit] hierarchy level, enter edit protocols rip
Set the authentication type to simple .	From the Authentication type list, select simple .	Set the authentication type to simple : set authentication-type simple
Set the authentication key to a simple-text password. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.	In the Authentication key box, type a simple-text password, and click OK .	Set the authentication key to a simple-text password: set authentication-key <i>password</i>

Enabling Authentication with MD5 Authentication

To configure authentication that requires an MD5 password to be included in the transmitted packet, enable MD5 authentication by performing these steps on all RIP Services Routers in the network:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 111 on page 406.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying the RIP Configuration” on page 406.

Table 111: Configuring MD5 RIP Authentication

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to Rip level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to Rip, click Edit. 	From the [edit] hierarchy level, enter edit protocols rip
Set the authentication type to MD5 .	From the Authentication type list, select md5 .	Set the authentication type to md5: set authentication-type md5
Set the MD5 authentication key (password). The key can be from 1 through 16 contiguous characters long and can include any ASCII strings.	In the Authentication key box, type an MD5 authentication key, and click OK .	Set the MD5 authentication key: set authentication-key <i>password</i>

Verifying the RIP Configuration

To verify the RIP configuration, perform these tasks:

- Verifying the RIP-Enabled Interfaces on page 407
- Verifying the Exchange of RIP Messages on page 408
- Verifying Reachability of All Hosts in the RIP Network on page 409

Verifying the RIP-Enabled Interfaces

Purpose Verify that all the RIP-enabled interfaces are available and active.

Action From the CLI, enter the show rip neighbor command.

```
user@host> show rip neighbor
Source      Destination      Send  Receive  In
Neighbor    State  Address         Address
-----
ge-0/0/0.0   Dn  (null)         (null)
ge-0/0/1.0   Up  192.168.220.5  224.0.0.9  mcast  both    1
ge-0/0/1.0   Up  192.168.220.5  224.0.0.9  mcast  both    1
```

What It Means The output shows a list of the RIP neighbors that are configured on the Services Router. Verify the following information:

- Each configured interface is present. Interfaces are listed in alphabetical order.
- Each configured interface is up. The state of the interface is listed in the **Destination State** column. A state of **Up** indicates that the link is passing RIP traffic. A state of **Dn** indicates that the link is not passing RIP traffic. In a point-to-point link, this state generally means that either the end point is not configured for RIP or the link is unavailable.

See Also For a complete description of show rip neighbor output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying the Exchange of RIP Messages

Purpose Verify that RIP messages are being sent and received on all RIP-enabled interfaces.

Action From the CLI, enter the show rip statistics command.

```
user@host> show rip statistics
RIPv2 info: port 520; holddown 120s.
      rts learned  rts held down  rqsts dropped  resps dropped
           10           0           0           0

t1-0/0/2.0: 0 routes learned; 13 routes advertised; timeout 120s; update interval
45s
Counter                Total    Last 5 min  Last minute
-----
Updates Sent            2855         11         2
Triggered Updates Sent    5          0          0
Responses Sent           0          0          0
Bad Messages             0          0          0
RIPv1 Updates Received    0          0          0
RIPv1 Bad Route Entries   0          0          0
RIPv1 Updates Ignored     0          0          0
RIPv2 Updates Received    41         0          0
RIPv2 Bad Route Entries   0          0          0
RIPv2 Updates Ignored     0          0          0
Authentication Failures   0          0          0
RIP Requests Received     0          0          0
RIP Requests Ignored      0          0          0

ge-0/0/1.0: 10 routes learned; 3 routes advertised; timeout 180s; update interval
30s
Counter                Total    Last 5 min  Last minute
-----
Updates Sent            2855         11         2
Triggered Updates Sent    3          0          0
Responses Sent           0          0          0
Bad Messages             1          0          0
RIPv1 Updates Received    0          0          0
RIPv1 Bad Route Entries   0          0          0
RIPv1 Updates Ignored     0          0          0
RIPv2 Updates Received    2864        11         2
RIPv2 Bad Route Entries   14         0          0
RIPv2 Updates Ignored     0          0          0
Authentication Failures   0          0          0
RIP Requests Received     0          0          0
RIP Requests Ignored      0          0          0
```

What It Means The output shows the number of RIP routes learned. It also shows the number of RIP updates sent and received on the RIP-enabled interfaces. Verify the following information:

- The number of RIP routes learned matches the number of expected routes learned. Subnets learned by direct connectivity through an outgoing interface are not listed as RIP routes.
- RIP updates are being sent on each RIP-enabled interface. If no updates are being sent, the routing policy might not be configured to export routes.
- RIP updates are being received on each RIP-enabled interface. If no updates are being received, the routing policy might not be configured to export routes on

the host connected to that subnet. The lack of updates might also indicate an authentication error.

See Also For a complete description of `show rip statistics` output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying Reachability of All Hosts in the RIP Network

Purpose By using the `traceroute` tool on each loopback address in the network, verify that all hosts in the RIP network are reachable from each Services Router.

Action For each Services Router in the RIP network:

1. In the J-Web interface, select **Diagnose > Traceroute**.
2. In the Remote Host box, type the name of a host for which you want to verify reachability from the Services Router.
3. Click **Start**. Output appears on a separate page.

```
1 172.17.40.254 (172.17.40.254) 0.362 ms 0.284 ms 0.251 ms
2 routera-fxp0.englab.mycompany.net (192.168.71.246) 0.251 ms 0.235 ms 0.200 ms
```

What It Means Each numbered row in the output indicates a router (“hop”) in the path to the host. The three time increments indicate the round-trip time (RTT) between the Services Router and the hop, for each traceroute packet.

To ensure that the RIP network is healthy, verify the following information:

- The final hop in the list is the host you want to reach.
- The number of expected hops to the host matches the number of hops in the traceroute output. The appearance of more hops than expected in the output indicates that a network segment is probably unreachable. It might also indicate that the incoming or outgoing metric on one or more hosts has been set unexpectedly.

See Also For information about the `traceroute` command and its output, see the *JUNOS System Basics and Services Command Reference*.

Chapter 14

Configuring an OSPF Network

The Open Shortest Path First protocol (OSPF) is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). To use OSPF, you must understand the basic components of an OSPF network and configure the J-series Services Router to act as a node in the network.



NOTE: The J-series Services Router supports both OSPF version 2 and OSPF version 3. In this chapter, the term *OSPF* refers to both versions of the protocol.

You can use either J-Web Quick Configuration or a configuration editor to configure an OSPF network.

This chapter contains the following topics. For more information about OSPF, see the *JUNOS Routing Protocols Configuration Guide*.

- OSPF Overview on page 411
- Before You Begin on page 412
- Configuring an OSPF Network with Quick Configuration on page 413
- Configuring an OSPF Network with a Configuration Editor on page 414
- Tuning an OSPF Network for Efficient Operation on page 422
- Verifying an OSPF Configuration on page 426

OSPF Overview

In an OSPF network, the network topology is distributed among the nodes of the autonomous system (AS) and is regularly updated. Because topology changes are flooded throughout the network, every node maintains the same copy of the network map in its local topological database. Packets are then routed based on the shared topology.

Enabling OSPF

To activate OSPF on a network, you must enable the protocol on all interfaces within the network on which OSPF traffic is to travel. To enable OSPF on one or more interfaces, you must configure one or more interfaces on the Services Router within an OSPF area. Once the interfaces are configured, OSPF link-state advertisements

(LSAs) are transmitted on all OSPF-enabled interfaces, and the network topology is shared throughout the network.

OSPF Areas

OSPF is enabled on a per-interface basis. Those interfaces are configured as OSPF enabled, and are assigned to an area. In a simple, single-area network, the area has the numeric identifier 0.0.0.0, which designates it as the backbone area. As the network grows, it is divided into multiple subnetworks or areas that are identified by numeric identifiers unique to the AS.

In a multiarea network, all areas must be directly connected to the backbone area by area border routers (ABRs). Because all areas are adjacent to the backbone area, OSPF routers send all traffic not destined for their own area through the backbone area. The ABRs in the backbone area are then responsible for transmitting the traffic through the appropriate ABR to the destination area.

Path Cost Metrics

Once the topology is shared across the network, OSPF uses it to route packets between network nodes. Each path between neighbors is assigned a cost based on the throughput, round-trip time, and reliability of the link. The sum of the costs across a particular path between hosts determines the overall cost of the path. Packets are then routed along the shortest path using the shortest path first (SPF) algorithm. If multiple equal-cost paths exist between a source and destination address, OSPF routes packets along each path alternately, in round-robin fashion.

OSPF allows you to manually assign a cost (or metric) to a particular path segment to control the flow of packets across the network.

OSPF Dial-on-Demand Circuits

If you are configuring OSPF across a demand circuit such as an ISDN link, you must enable dial-on-demand routing backup on the OSPF-enabled interface. Because demand circuits do not pass all traffic required to maintain an OSPF adjacency (hello packets, for example), you configure dial-on-demand routing so individual nodes in an OSPF network can maintain adjacencies despite the lack of LSA exchanges.

To configure an ISDN link, see “Configuring ISDN” on page 211. For information about configuring OSPF demand circuits, see “Configuring Dial-on-Demand Routing Backup with OSPF Support (Optional)” on page 233.

Before You Begin

Before you begin configuring an OSPF network, complete the following tasks:

- Establish basic connectivity. See the Getting Started Guide for your router.
- Configure network interfaces. See “Configuring Ethernet, DS1, DS3, and Serial Interfaces” on page 103.

Configuring an OSPF Network with Quick Configuration

J-Web Quick Configuration allows you to create single-area OSPF networks. Figure 72 on page 413 shows the Quick Configuration Routing page for OSPF.

Figure 72: Quick Configuration Routing Page for OSPF

The screenshot displays the Juniper J-Web interface for configuring OSPF. The top navigation bar includes links for Monitor, Configuration, Diagnose, Manage, Events, and Alarms. The left sidebar shows the Quick Configuration menu with options for View and Edit, History, and Rescue. The main content area is titled 'Quick Configuration' and 'Routing and Protocols'. Under 'Router Identification', the 'Router Identifier' is set to 10.255.0.10. The 'OSPF' section includes an 'Enable OSPF' checkbox, 'OSPF Area ID' set to 0.0.0.0, and 'Area Type' set to regular. There is also an 'Enable OSPF on All Interfaces' checkbox. Below these are two lists: 'OSPF-Enabled Interfaces' (currently empty) and 'OSPF-Disabled Interfaces' (containing fe-0/0/0.0, lo0.0, and tl-4/0/0.0). Navigation arrows are between the lists. At the bottom are 'OK', 'Cancel', and 'Apply' buttons. The footer contains copyright information and the Juniper logo.

To configure a single-area OSPF network with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Quick Configuration > Routing > OSPF Routing**.
2. Enter information into the Quick Configuration Routing page for OSPF, as described in Table 112 on page 414.
3. Click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration Routing page for OSPF, click **Apply**.
 - To apply the configuration and return to the Quick Configuration Routing page, click **OK**.
 - To cancel your entries and return to the Quick Configuration Routing page, click **Cancel**.
4. To check the configuration, see “Verifying an OSPF Configuration” on page 426.

Table 112: OSPF Routing Quick Configuration Summary

Field	Function	Your Action
Router Identification		
Router Identifier (required)	Uniquely identifies the router.	Type the Services Router's 32-bit IP address, in dotted decimal notation.
OSPF		
Enable OSPF	Enables or disables OSPF.	<ul style="list-style-type: none"> ■ To enable OSPF, select the check box. ■ To disable OSPF, clear the check box.
OSPF Area ID	Uniquely identifies the area within its AS.	<p>Type a 32-bit numeric identifier for the area, or an integer.</p> <p>If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter 3, the value assigned to the area is 0.0.0.3.</p>
Area Type	Designates the type of OSPF area.	<p>From the list, select the type of OSPF area you are creating:</p> <ul style="list-style-type: none"> ■ regular—A regular OSPF area, including the backbone area ■ stub—A stub area ■ nssa—A not-so-stubby area (NSSA)
OSPF-Enabled Interfaces	<p>Designates one or more Services Router interfaces on which OSPF is enabled.</p> <p>For information about interface names, see “Network Interface Naming” on page 47.</p>	<p>The first time you configure OSPF, the Logical Interfaces box displays a list of all the logical interfaces configured on the Services Router. Do any of the following:</p> <ul style="list-style-type: none"> ■ To enable OSPF on an interface, click the interface name to highlight it, and click the left arrow to add the interface to the OSPF interfaces list. ■ To enable OSPF on multiple interfaces at once, press Ctrl while you click multiple interface names to highlight them. Then click the left arrow to add the interfaces to the OSPF interfaces list. ■ To enable OSPF on all logical interfaces except the special fxp0 management interface, select All Interfaces in the Logical Interfaces list and click the left arrow. ■ To enable OSPF on all the interfaces displayed in the Logical Interfaces list, click All to highlight every interface. Then click the left arrow to add the interfaces to the OSPF interfaces list. ■ To disable OSPF on one or more interfaces, highlight the interface or interfaces in the OSPF interfaces box and click the right arrow to move them back to the Logical Interfaces list.

Configuring an OSPF Network with a Configuration Editor

To configure the Services Router as a node in an OSPF network, you must perform the following tasks marked *(Required)*.

- Configuring the Router Identifier (Required) on page 415
- Configuring a Single-Area OSPF Network (Required) on page 415
- Configuring a Multiarea OSPF Network (Optional) on page 417
- Configuring Stub and Not-So-Stubby Areas (Optional) on page 420

To configure an OSPF demand circuit, see “Configuring Dial-on-Demand Routing Backup with OSPF Support (Optional)” on page 233. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 211.)

For information about using the J-Web and CLI configuration editors, see “Using Services Router Configuration Tools” on page 3.

Configuring the Router Identifier (Required)

The router identifier is the IP address that uniquely identifies the J-series Services Router.

OSPF uses the router identifier to elect a designated router, unless you manually specify a priority value. When the OSPF network first becomes active, by default, the router with the highest router identifier is elected the designated router.

To configure the router identifier for the Services Router:


- 1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
- 2. Perform the configuration tasks described in Table 113 on page 415.
- 3. Go on to “Configuring a Single-Area OSPF Network (Required)” on page 415.

Table 113: Configuring the Router Identifier

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Routing-options level in the configuration hierarchy.	<div>1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration.</div> <div>2. Next to Routing options, click Configure or Edit.</div>	From the [edit] hierarchy level, enter edit routing-options
Set the router ID value to the IP address of the Services Router—for example, 177.162.4.24.	<div>1. In the Router Id box, type 177.162.4.24.</div> <div>2. Click OK.</div>	<div>Enter</div> <div>set router-id 177.162.4.24</div>

Configuring a Single-Area OSPF Network (Required)

To use OSPF on the Services Router, you must configure at least one OSPF area, like the one shown in Figure 73 on page 416.

Figure 73: Typical Single-Area OSPF Network Topology

To configure a single-area OSPF network with a backbone area, like the one in Figure 73 on page 416, perform these steps on each Services Router in the network:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 114 on page 417.
3. If you are finished configuring the router, commit the configuration.

After you create the backbone area and add the appropriate interfaces to the area, OSPF begins sending LSAs. No additional configuration is required to enable OSPF traffic on the network.

4. Go on to one of the following procedures:
 - To add more areas to the AS, see “Configuring a Multiarea OSPF Network (Optional)” on page 417.
 - To control external route advertisement in the AS, see “Configuring Stub and Not-So-Stubby Areas (Optional)” on page 420.
 - To configure an OSPF demand circuit, see “Configuring Dial-on-Demand Routing Backup with OSPF Support (Optional)” on page 233. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 211.)
 - To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 422.
 - To check the configuration, see “Verifying an OSPF Configuration” on page 426.


Table 114: Configuring a Single-Area OSPF Network

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Ospf level in the configuration hierarchy.	<div>1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration.</div> <div>2. Next to Protocols, click Configure or Edit.</div> <div>3. Next to Ospf, click Configure or Edit.</div>	<div>From the [edit] hierarchy level, enter</div> <div> edit protocols ospf</div>
Create the backbone area with area ID 0.0.0.0.	<div>1. In the Area box, click Add new entry.</div> <div>2. In the Area ID box, type 0.0.0.0.</div>	<div>1. Set the backbone area ID to 0.0.0.0 and add an interface:</div>
Add interfaces as needed to the OSPF area—for example, ge-0/0/0.	<div>1. In the Interface box, click Add new entry.</div> <div>2. In the Interface name box, type ge-0/0/0.</div> <div>3. Click OK.</div> <div>4. Repeat Step 1 through Step 3 for each interface on this Services Router that you are adding to the backbone area. Only one interface is required.</div>	<div>set area 0.0.0.0 interface ge-0/0/0</div> <div>2. Repeat Step 1 for each interface on this Services Router that you are adding to the backbone area. Only one interface is required.</div>

Configuring a Multiarea OSPF Network (Optional)

To reduce traffic and topology maintenance for the Services Routers in an OSPF autonomous system (AS), you can group them into multiple areas, as shown in Figure 74 on page 417.

Figure 74: Typical Multiarea OSPF Network Topology



To configure a multiarea OSPF network shown in Figure 74 on page 417, perform the following tasks on the appropriate Services Routers in the network. You must create

a backbone area. To link each additional area to the backbone area, you must configure one of the Services Routers as an area border router (ABR).

- Creating the Backbone Area on page 418
- Creating Additional OSPF Areas on page 418
- Configuring Area Border Routers on page 419

Creating the Backbone Area

On each Services Router that is to operate as an ABR in the network, create backbone area 0.0.0.0 with at least one interface enabled for OSPF.

For instruction, see “Configuring a Single-Area OSPF Network (Required)” on page 415.

Creating Additional OSPF Areas

To create additional OSPF areas:

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 115 on page 418.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To configure this Services Router as an area border router, see “Configuring Area Border Routers” on page 419.
 - To control external route advertisement in the AS, see “Configuring Stub and Not-So-Stubby Areas (Optional)” on page 420.
 - To configure an OSPF demand circuit, see “Configuring Dial-on-Demand Routing Backup with OSPF Support (Optional)” on page 233. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 211.)
 - To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 422.
 - To check the configuration, see “Verifying an OSPF Configuration” on page 426.

Table 115: Configuring a Multiarea OSPF Network

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Ospf level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to Ospf, click Edit. 	<p>From the [edit] hierarchy level, enter</p> <p>edit protocols ospf</p>

Table 115: Configuring a Multiarea OSPF Network (*continued*)

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the additional area with a unique area ID, in dotted decimal notation—for example, 0.0.0.2.	<ol style="list-style-type: none"> 1. In the Area box, click Add new entry. 2. In the Area ID box, type 0.0.0.2. 	<ol style="list-style-type: none"> 1. Set the area ID to 0.0.0.2 and add an interface: set area 0.0.0.2 interface ge-0/0/0
Add interfaces as needed to the OSPF area—for example, ge-0/0/0.	<ol style="list-style-type: none"> 1. In the Interface box, click Add new entry. 2. In the Interface name box, type ge-0/0/0. 3. Click OK. 4. Repeat Step 1 through Step 3 for each interface on this Services Router that you are adding to the backbone area. Only one interface is required. 	<ol style="list-style-type: none"> 2. Repeat Step 1 for each interface on this Services Router that you are adding to the area. Only one interface is required.

Configuring Area Border Routers

A Services Router operating as an area border router (ABR) has interfaces enabled for OSPF in the backbone area and in the area you are linking to the backbone. For example, Services Router B acts as the ABR in Figure 74 on page 417 and has interfaces in both the backbone area and area 0.0.0.3.

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 116 on page 420.
3. If you are finished configuring the router, commit the configuration.

After you create the areas on the appropriate Services Routers and add and enable the appropriate interfaces to the areas, no additional configuration is required to enable OSPF traffic within or across the areas.

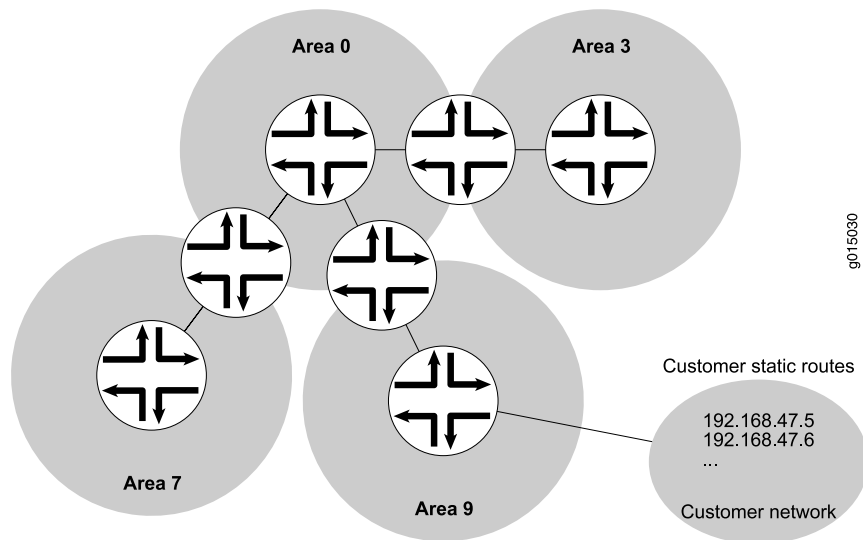
4. Go on to one of the following procedures:
 - To control external route advertisement in the AS, see “Configuring Stub and Not-So-Stubby Areas (Optional)” on page 420.
 - To configure an OSPF demand circuit, see “Configuring Dial-on-Demand Routing Backup with OSPF Support (Optional)” on page 233. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 211.)
 - To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 422.
 - To check the configuration, see “Verifying an OSPF Configuration” on page 426.

Table 116: Configuring Area Border Routers

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Ospf level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to Ospf, click Edit. 	From the [edit] hierarchy level, enter edit protocols ospf
Verify that the backbone area has at least one interface enabled for OSPF.	<p>Click 0.0.0.0 to display the Area ID 0.0.0.0 page, and verify that the backbone area has at least one interface enabled for OSPF.</p> <p>For example, Services Router B in Figure 74 on page 417 has the following interfaces enabled for OSPF in the backbone area:</p> <ul style="list-style-type: none"> ■ Interface ge-0/0/0.0 ■ Interface ge-0/0/1.0 <p>To enable an interface on the backbone area, see “Configuring a Single-Area OSPF Network (Required)” on page 415.</p>	<p>View the configuration using the show command:</p> <p>show</p> <p>For example, Services Router B in Figure 74 on page 417 has the following interfaces enabled for OSPF in the backbone area:</p> <pre>area 0.0.0.0 { interface ge-0/0/0.0; interface ge-0/0/1.0; }</pre> <p>To enable an interface on the backbone area, see “Configuring a Single-Area OSPF Network (Required)” on page 415.</p>
Create the additional area with a unique area ID—for example, 0.0.0.2.	<ol style="list-style-type: none"> 1. In the Area box, click Add new entry. 2. In the Area ID box, type 0.0.0.2. 	<ol style="list-style-type: none"> 1. Set the area ID to 0.0.0.2 and add an interface: <pre>set area 0.0.0.2 interface ge-0/0/0</pre>
Add interfaces as needed to the OSPF area—for example, ge-0/0/0.	<ol style="list-style-type: none"> 1. In the Interface box, click Add new entry. 2. In the Interface name box, type ge-0/0/0. 3. Click OK. 4. Repeat Step 1 through Step 3 for each interface on this Services Router that you are adding to the backbone area. Only one interface is required. 	<ol style="list-style-type: none"> 2. Repeat Step 1 for each interface on this Services Router that you are adding to the area. Only one interface is required.

Configuring Stub and Not-So-Stubby Areas (Optional)

To control the advertisement of external routes into an area, you can create stub areas and not-so-stubby areas (NSSAs) in an OSPF network. In the network shown in Figure 75 on page 421, area 0.0.0.7 has no external connections and can be configured as a stub area. Area 0.0.0.9 only has external connections to static routes and can be configured as an NSSA.

Figure 75: OSPF Network Topology with Stub Areas and NSSAs

To configure stub areas and NSSAs in an OSPF network like the one shown in Figure 75 on page 421:

1. Create the area and enable OSPF on the interfaces within that area.
For instructions, see “Creating Additional OSPF Areas” on page 418.
2. Configure an area border router to bridge the areas.
For instructions, see “Configuring Area Border Routers” on page 419.
3. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
4. To configure each Services Router in area 0.0.0.7 as a stub area router, perform the configuration tasks described in Table 117 on page 422.
5. If you are finished configuring the router, commit the configuration.
6. Go on to one of the following procedures:
 - To configure an OSPF demand circuit, see “Configuring Dial-on-Demand Routing Backup with OSPF Support (Optional)” on page 233. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 211.)
 - To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 422.
 - To check the configuration, see “Verifying an OSPF Configuration” on page 426.

Table 117: Configuring Stub Area and Not-So-Stubby Area Routers

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the 0.0.0.7 level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to Ospf, click Edit. 4. Under Area id, click 0.0.0.7. 	From the [edit] hierarchy level, enter edit protocols ospf area 0.0.0.7
Configure each Services Router in area 0.0.0.7 as a stub router.	<ol style="list-style-type: none"> 1. In the Stub option list, select Stub and click OK. 2. Repeat Step 1 for every Services Router in the stub area to configure them with the stub parameter for the area. 	<ol style="list-style-type: none"> 1. Set the stub attribute: set stub 2. Repeat Step 1 for every Services Router in the stub area to configure them with the stub parameter for the area.
Navigate to the 0.0.0.9 level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. On the main Configuration page next to Protocols, click Edit. 2. Next to Ospf, click Edit. 3. Under Area id, click 0.0.0.9. 	From the [edit] hierarchy level, enter edit protocols ospf area 0.0.0.9
Configure each Services Router in area 0.0.0.9 as an NSSA router.	<ol style="list-style-type: none"> 1. In the Stub option list, select Nssa and click OK. 2. Repeat Step 1 for every Services Router in the NSSA to configure them with the nssa parameter for the area. 	<ol style="list-style-type: none"> 1. Set the nssa attribute: set nssa 2. Repeat Step 1 for every Services Router in the NSSA to configure them with the nssa parameter for the area.

Tuning an OSPF Network for Efficient Operation

To make your OSPF network operate more efficiently, you can change some default settings on the Services Router by performing the following tasks:

- Controlling Route Selection in the Forwarding Table on page 422
- Controlling the Cost of Individual Network Segments on page 423
- Enabling Authentication for OSPF Exchanges on page 424
- Controlling Designated Router Election on page 425

Controlling Route Selection in the Forwarding Table

OSPF uses route preferences to select the route that is installed in the forwarding table when several routes have the same shortest path first (SFP) calculation. To

evaluate a route, OSPF calculates the sum of the individual preferences of every router along the path and selects the route with the lowest total preference.

By default, internal OSPF routes have a preference value of **10**, and external OSPF routes have a preference value of **150**. Suppose all routers in your OSPF network use the default preference values. By setting the internal preference to **7** and the external preference to **130**, you can ensure that the path through a particular Services Router is selected for the forwarding table any time multiple equal-cost paths to a destination exist.

To modify the default preferences on a Services Router:

- 1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
- 2. Perform the configuration tasks described in Table 118 on page 423.

Table 118: Controlling Route Selection in the Forwarding Table by Setting Preferences

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Ospf level in the configuration hierarchy.	<ul style="list-style-type: none">1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration.2. Next to Protocols, click Edit.3. Next to Ospf, click Edit.	From the [edit] hierarchy level, enter edit protocols ospf
Set the external and internal route preferences.	<ul style="list-style-type: none">1. In the External preference box, type 130.2. In the Preference box, type the internal preference value of 7.3. Click OK.	<ul style="list-style-type: none">1. Set the external preference: set external-preference 1302. Set the internal preference: set preference 7

Controlling the Cost of Individual Network Segments

When evaluating the cost of individual network segments, OSPF evaluates the reference bandwidth. For any link faster than 100 Mbps, the default cost metric is **1**. When OSPF calculates the SPF algorithm, it sums the metrics of all interfaces along a path to determine the overall cost of the path. The path with the lowest metric is selected for the forwarding table.

To control the cost of the network segment, you can modify the metric value on an individual interface. Suppose all routers in the OSPF network use default metric values. If you increase the metric on an interface to **5**, all paths through this interface have a calculated metric higher than the default and are *not* preferred.

To manually set the cost of a network segment on the stub area's Fast Ethernet interface by modifying the interface metric:

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 119 on page 424.

Table 119: Controlling the Cost of Individual Network Segments by Modifying the Metric

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the ge-0/0/0.0 level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to Ospf, click Edit. 4. Under Area id, click 0.0.0.0. 5. Under Interface name, click ge-0/0/0.0. 	From the [edit] hierarchy level, enter edit protocols ospf area 0.0.0.0 interface ge-0/0/0.0
Set the interface metric.	<ol style="list-style-type: none"> 1. In the Metric box, type the interface metric value 5. 2. Click OK. 	Set the interface metric: set metric 5

Enabling Authentication for OSPF Exchanges

All OSPFv2 protocol exchanges can be authenticated to guarantee that only trusted routers participate in the AS's routing. By default, OSPF authentication is disabled.



NOTE: OSPFv3 does not support authentication.

You can enable either of two authentication types:

- Simple authentication—Authenticates by means of a plain-text password (key) included in the transmitted packet.
- MD5 authentication—Authenticates by means of an MD5 checksum included in the transmitted packet.

Because OSPF performs authentication at the area level, all routers within the area must have the same authentication and corresponding password (key) configured. For MD5 authentication to work, both the receiving and transmitting routers must have the same MD5 key.

To enable OSPF authentication on the stub area:

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 120 on page 425.

Table 120: Enabling OSPF Authentication

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the 0.0.0.0 level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to Ospf, click Edit. 4. Under Area id, click 0.0.0.0. 	<p>From the [edit] hierarchy level, enter</p> <pre>edit protocols ospf area 0.0.0.0</pre>
Set the authentication type for the stub area to either simple or MD5—for example, MD5.	<ol style="list-style-type: none"> 1. From the Authentication type list, select md5. 2. Click OK. 	<p>Set the authentication type:</p> <pre>set authentication-type md5</pre>
Navigate to the <i>interface-name</i> level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. On the main Configuration page next to Protocols, click Edit. 2. Next to Ospf, click Edit. 3. Under Area id, click 0.0.0.0. 4. Under Interface name, click an interface name. 	<p>From the [edit] hierarchy level, enter</p> <pre>edit protocols ospf area 0.0.0.0 interface interface-name</pre>
<p>Set the authentication password (key) and, for MD5 authentication only, the key identifier to associate with the MD5 password:</p> <ul style="list-style-type: none"> ■ For simple authentication, set a password of from 1 through 8 ASCII characters—for example, Chey3nne. ■ For MD5 authentication: <ul style="list-style-type: none"> ■ Set a password of from 1 through 16 ASCII characters—for example, Chey3nne. ■ Set a key identifier between 0 (the default) and 255—for example, 2. 	<ol style="list-style-type: none"> 1. In the Key name box, type Chey3nne. 2. For MD5 authentication only, in the Key ID box, type 2. 3. Click OK. 4. Repeat Step 1 through Step 3 for each interface in the stub area for which you are enabling authentication. 	<ol style="list-style-type: none"> 1. Set the authentication password and, for MD5 authentication only, set the key identifier: <pre>set authentication-key Chey3nne key-id 2</pre> 2. Repeat Step 1 for each interface in the stub area for which you are enabling authentication.

Controlling Designated Router Election

At designated router election, the router priorities are evaluated first, and the router with the highest priority is elected designated router.

By default, routers have a priority of **128**. A priority of **0** marks the router as ineligible to become the designated router. To configure a router so it is always the designated router, set its priority to **255**.

To change the priority of a Services Router to control designated router election:

1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 121 on page 426.

Table 121: Controlling Designated Router Election

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the OSPF interface address for the Services Router. For example, navigate to the <code>ge-0/0/1</code> level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to Ospf, click Edit. 4. Under Area id, click 0.0.0.3. 5. Under Interface name, click ge-0/0/1. 	<p>From the [edit] hierarchy level, enter</p> <p><code>edit protocols ospf area 0.0.0.3 interface ge-0/0/1</code></p>
Set the Services Router priority to a value between 0 and 255—for example, 200. The default value is 128.	<ol style="list-style-type: none"> 1. In the Priority box, type 200. 2. Click OK. 	<p>Set the priority value:</p> <p><code>set priority 200</code></p>

Verifying an OSPF Configuration

To verify an OSPF configuration, perform these tasks:

- Verifying OSPF-Enabled Interfaces on page 427
- Verifying OSPF Neighbors on page 428
- Verifying the Number of OSPF Routes on page 429
- Verifying Reachability of All Hosts in an OSPF Network on page 430

Verifying OSPF-Enabled Interfaces

- Purpose

Verify that OSPF is running on a particular interface and that the interface is in the desired area.
- Action

From the CLI, enter the `show ospf interface` command.

```
user@host> show ospf interface
Intf      State   Area      DR ID      BDR ID      Nbrs
at-5/1/0.0 PtToPt  0.0.0.0    0.0.0.0    0.0.0.0      1
ge-2/3/0.0 DR       0.0.0.0    192.168.4.16 192.168.4.15  1
lo0.0     DR       0.0.0.0    192.168.4.16 0.0.0.0       0
so-0/0/0.0 Down    0.0.0.0    0.0.0.0    0.0.0.0       0
so-6/0/1.0 PtToPt  0.0.0.0    0.0.0.0    0.0.0.0       1
so-6/0/2.0 Down    0.0.0.0    0.0.0.0    0.0.0.0       0
so-6/0/3.0 PtToPt  0.0.0.0    0.0.0.0    0.0.0.0       1
```

- What It Means

The output shows a list of the Services Router interfaces that are configured for OSPF. Verify the following information:
- Each interface on which OSPF is enabled is listed.
- Under **Area**, each interface shows the area for which it was configured.
- Under **Intf** and **State**, the Services Router loopback (**lo0.0**) interface and LAN interface that are linked to the OSPF network's designated router (**DR**) are identified.
- Under **DR ID**, the IP address of the OSPF network's designated router appears.
- Under **State**, each interface shows a state of **PtToPt** to indicate a point-to-point connection. If the state is **Waiting**, check the output again after several seconds. A state of **Down** indicates a problem.
- The designated router addresses always show a state of **DR**.

See Also

For a complete description of `show ospf interface` output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying OSPF Neighbors

Purpose OSPF neighbors are interfaces that have an immediate adjacency. On a point-to-point connection between the Services Router and another router running OSPF, verify that each router has a single OSPF neighbor.

Action From the CLI, enter the `show ospf neighbor` command.

```
user@host> show ospf neighbor
```

Address	Intf	State	ID	Pri	Dead
192.168.254.225	fxp3.0	2Way	10.250.240.32	128	36
192.168.254.230	fxp3.0	Full	10.250.240.8	128	38
192.168.254.229	fxp3.0	Full	10.250.240.35	128	33
10.1.1.129	fxp2.0	Full	10.250.240.12	128	37
10.1.1.131	fxp2.0	Full	10.250.240.11	128	38
10.1.2.1	fxp1.0	Full	10.250.240.9	128	32
10.1.2.81	fxp0.0	Full	10.250.240.10	128	33

What It Means The output shows a list of the Services Router's OSPF neighbors and their addresses, interfaces, states, router IDs, priorities, and number of seconds allowed for inactivity ("dead" time). Verify the following information:

- Each interface that is immediately adjacent to the Services Router is listed.
- The Services Router's own loopback address and the loopback addresses of any routers with which the Services Router has an immediate adjacency are listed.
- Under **State**, each neighbor shows a state of **Full**. Because full OSPF connectivity is established over a series of packet exchanges between clients, the OSPF link might take several seconds to establish. During that time, the state might be displayed as **Attempt**, **Init**, or **2way**, depending on the stage of negotiation.

If, after 30 seconds, the state is not **Full**, the OSPF configuration between the neighbors is not functioning correctly.


See Also For a complete description of `show ospf neighbor` output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying the Number of OSPF Routes

- Purpose** Verify that the OSPF routing table has entries for the following:
- Each subnetwork reachable through an OSPF link
 - Each loopback address reachable on the network

For example, Figure 76 on page 429 shows a sample network with an OSPF topology.

Figure 76: Sample OSPF Network Topology



In this topology, OSPF is being run on all interfaces. Each segment in the network is identified by an address with a /24 prefix, with interfaces on either end of the segment being identified by unique IP addresses.

Action From the CLI, enter the show ospf route command.

user@host> show ospf route							
Prefix	Path	Route	NH	Metric	NextHop	Nexthop	
	Type	Type	Type		Interface	addr/label	
10.10.10.1/24	Intra	Network	IP	1	ge-0/0/2.0	10.0.21.1	
10.10.10.2/24	Intra	Network	IP	1	ge-0/0/2.0	10.0.21.1	
10.10.10.4/24	Intra	Network	IP	1	ge-0/0/1.0	10.0.13.1	
10.10.10.5/24	Intra	Network	IP	1	ge-0/0/2.0	10.0.21.1	
10.10.10.6/24	Intra	Network	IP	1	ge-0/0/1.0	10.0.13.1	
10.10.10.10/24	Intra	Network	IP	1	ge-0/0/2.0	10.0.21.1	
10.10.10.11/24	Intra	Network	IP	1	ge-0/0/1.0	10.0.13.1	
10.10.10.13/24	Intra	Network	IP	1	ge-0/0/1.0	10.0.13.1	
10.10.10.16/24	Intra	Network	IP	1	ge-0/0/1.0	10.0.13.1	
10.10.10.19/24	Intra	Network	IP	1	ge-0/0/1.0	10.0.13.1	
10.10.10.20/24	Intra	Network	IP	1	ge-0/0/2.0	10.0.21.1	
10.10.10.21/24	Intra	Network	IP	1	ge-0/0/2.0	10.0.21.1	
192.168.5.1	Intra	Router	IP	1	ge-0/0/2.0	10.0.21.1	
192.168.5.2	Intra	Router	IP	1	lo0		
192.168.5.3	Intra	Router	IP	1	ge-0/0/1.0	10.0.13.1	
192.168.5.4	Intra	Router	IP	1	ge-0/0/1.0	10.0.13.1	
192.168.5.5	Intra	Router	IP	1	ge-0/0/1.0	10.0.13.1	
192.168.5.6	Intra	Router	IP	1	ge-0/0/2.0	10.0.21.1	
192.168.5.7	Intra	Router	IP	1	ge-0/0/2.0	10.0.21.1	

```

192.168.5.8          Intra Router   IP 1      ge-0/0/2.0  10.0.21.1
192.168.5.9          Intra Router   IP 1      ge-0/0/1.0  10.0.13.1

```

What It Means The output lists each route, sorted by IP address. Routes are shown with a route type of **Network**, and loopback addresses are shown with a route type of **Router**.

For the example shown in Figure 76 on page 429, verify that the OSPF routing table has 21 entries, one for each network segment and one for each router's loopback address.

See Also For a complete description of `show ospf route` output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying Reachability of All Hosts in an OSPF Network

Purpose By using the traceroute tool on each loopback address in the network, verify that all hosts in the network are reachable from each Services Router.

Action For each Services Router in the OSPF network:

1. In the J-Web interface, select **Diagnose > Traceroute**.
2. In the Host Name box, type the name of a host for which you want to verify reachability from the Services Router.
3. Click **Start**. Output appears on a separate page.

```

1 172.17.40.254 (172.17.40.254) 0.362 ms 0.284 ms 0.251 ms
2 routera-fxp0.englab.mycompany.net (192.168.71.246) 0.251 ms 0.235 ms 0.200 ms

```

What It Means Each numbered row in the output indicates a router (“hop”) in the path to the host. The three time increments indicate the round-trip time (RTT) between the Services Router and the hop, for each traceroute packet. To ensure that the OSPF network is healthy, verify the following information:

- The final hop in the list is the host you want to reach.
- The number of expected hops to the host matches the number of hops in the traceroute output. The appearance of more hops than expected in the output indicates that a network segment is likely not reachable. In this case, verify the routes with the `show ospf route` command.

For information about `show ospf route`, see Verifying the Number of OSPF Routes on page 429.

See Also For information about the `traceroute` command and its output, see the *JUNOS System Basics and Services Command Reference*.

Chapter 15

Configuring the IS-IS Protocol

The Services Router supports the Intermediate System-to-Intermediate System (IS-IS) protocol.

You use either the J-Web configuration editor or CLI configuration editor to configure IS-IS.

This chapter contains the following topics. For more information about IS-IS, see the *JUNOS Routing Protocols Configuration Guide*.

- IS-IS Overview on page 431
- Before You Begin on page 432
- Configuring IS-IS with a Configuration Editor on page 433
- Verifying IS-IS on a Services Router on page 434

IS-IS Overview

On the Services Router, Intermediate System-to-Intermediate System (IS-IS) protocol is an interior gateway routing protocol (IGP) that uses link-state information for routing network traffic. IS-IS uses the shortest path first (SPF) algorithm to determine routes. Using SPF, IS-IS evaluates network topology changes and determines if a full or partial route calculation is required. The protocol was originally developed for routing International Organization for Standards (ISO) connectionless network protocol (CLNP) packets.

This overview contains the following topics:

- ISO Network Addresses on page 431
- System Identifier Mapping on page 432

ISO Network Addresses

IS-IS uses ISO network addresses. Each address identifies a point of connection to the network, such as a router interface, which is called a network service access point (NSAP). NSAP addresses are supported on the loopback (lo0) interface. (For information about interface names, see “Network Interface Naming” on page 47.)

An end system can have multiple NSAP addresses, which differ by the last byte called an n-selector. Each NSAP represents a service that is available at the node. In addition to multiple services, a single node can belong to multiple areas.

Each network entity also has a special address called a network entity title (NET) with an identical structure to an NSAP address but an n-selector of 00. Most end systems and intermediate systems have one NET address, while intermediate systems participating in more than one area can have more than one NET address.

The following ISO addresses are examples of the IS-IS address format:

49.0001.00a0.c96b.c490.00

49.0001.2081.9716.9018.00

The first part of the address is the area number, which is a variable number from 1 to 13 bytes. The first byte of the area number, **49**, is the authority and format indicator (AFI). The next bytes are the assigned area identifier and can be from 0 to 12 bytes. In the examples, **0001** is the area identifier.

The next 6 bytes are the system identifier and can be any 6 bytes unique throughout the entire domain. The system identifier is commonly the media access control (MAC) address, as shown in the first example, **00a0.c96b.c490**. Otherwise, the system identifier is the IP address expressed in binary-coded decimal (BCD) format, as shown in the second example, **2081.9716.9018**, which corresponds to **208.197.169.18**. The last byte, **00**, is the n-selector.



NOTE: The system identifier cannot be configured as 0000.0000.0000. Using all zeros as an identifier is not supported and does not form an adjacency.

System Identifier Mapping

To provide assistance with debugging IS-IS, the Services Router supports dynamic mapping of ISO system identifiers to the hostname. Each router can be configured with a hostname that allows the system identifier-to-hostname mapping to be sent in a dynamic hostname type length value (TLV) in the IS-IS link-state PDU (LSP). The mapping permits an intermediate system in the routing domain to learn the ISO system identifier of another intermediate system.

Before You Begin

Before you begin configuring IS-IS, complete the following tasks:

- Establish basic connectivity. See the Getting Started Guide for your router.
- If you do not already have an understanding of IS-IS, read “IS-IS Overview” on page 368 or the *JUNOS Routing Protocols Configuration Guide*.
- Obtain ISO addresses for participating routers in the AS.

Configuring IS-IS with a Configuration Editor

To configure IS-IS with a configuration editor, you do the following:

- Enable IS-IS on the router.
- Configure a network entity title (NET) on one of the router interfaces, preferably the loopback interface, lo0.
- Configure the ISO family on all interfaces that are supporting the IS-IS protocol.

To configure IS-IS:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 122 on page 433.
3. Commit the configuration on the Services Router.
4. Repeat the configuration tasks on each Services Router in the IS-IS autonomous system (AS).

Table 122: Configuring the IS-IS Protocol

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none">1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration.2. Next to Interfaces, click Edit.	From the [edit] hierarchy level, enter edit interfaces.
Configure the loopback interface lo0.	<ol style="list-style-type: none">1. Next to Interface, click Add new entry.2. In the Interface name box, type lo0.3. Click OK.	Enter edit interfaces lo0
Configure the logical unit on the loopback interface—for example 0. Add the NET address to the loopback interface—for example, 49.0001.00a0.c96b.c490.00.	<ol style="list-style-type: none">1. Next to lo0, click Edit under Encapsulation.2. Next to Unit, click Add new entry.3. In the Interface unit number box, type 0.4. Under Family, select Iso.5. Next to Address, click Add new entry.6. In the Source box, type 49.0001.00a0.c96b.c490.00.7. Click OK until you return to the Interfaces page.	<ol style="list-style-type: none">1. Enter edit unit 02. Enter set family iso address 49.0001.00a0.c96b.c490.00

Table 122: Configuring the IS-IS Protocol *(continued)*

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a physical interface—for example, ge-0/0/1 —with the NET address, and add the Family type iso .	<ol style="list-style-type: none"> Next to ge-0/0/1, click Edit under Encapsulation. Next to Unit, click Add new entry. In the Interface unit number box, type 0. Under Family, select Iso. Next to Iso, click Configure. Next to Address, click Add new entry. In the Source box, type 49.0001.00a0.c96b.c490.00. Click OK until you return to the Edit Configuration page. 	<p>Enter</p> <p>edit interfaces ge-0/0/1</p> <p>Enter</p> <p>set unit 0</p> <p>Enter</p> <p>set family iso address 49.0001.00a0.c96b.c490.00</p>
Navigate to the Protocols level in the configuration hierarchy.	On the main Configuration page next to Protocols, click Edit .	From the [edit] hierarchy level, enter edit protocols
Add the IS-IS protocol to all interfaces on the Services Router.	<ol style="list-style-type: none"> Next to Isis, click Edit. In the Interface name box, type all. Click OK. 	<p>Enter</p> <p>set isis interface all</p>

Verifying IS-IS on a Services Router

To verify IS-IS, perform these tasks:

- Displaying IS-IS Interface Configuration on page 435
- Displaying IS-IS Interface Configuration Detail on page 436
- Displaying IS-IS Adjacencies on page 437
- Displaying IS-IS Adjacencies in Detail on page 438

Displaying IS-IS Interface Configuration

Purpose Verify the status of IS-IS-enabled interfaces.

Action From the CLI, enter the `show isis interface brief` command.

```
user@host> show isis interface brief
IS-IS interface database:
Interface  L CirID Level 1 DR Level 2 DR
lo0.0      3 0x1  router1 router.01
ge-0/0/1.0 2 0x9  Disabled router.03
ge-1/0/0.0 2 0x7  Disabled router.05
```

What It Means Verify that the output shows the intended configuration of the interfaces on which IS-IS is enabled.

See Also For a complete description of `show isis interface` output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Displaying IS-IS Interface Configuration Detail

Purpose Verify the details of IS-IS-enabled interfaces.

Action From the CLI, enter the show isis interface detail command.

```
user@host> show isis interface detail
lo0.0
  Index:3, State:0x7, Circuit id: 0x1, Circuit type:3
  LSP interval: 100 ms, Sysid: router1
  Level Adjacencies Priority Metric Hello(s) Hold(s)
    1      0      64      0      9      27
    2      0      64      0      9      27
ge-0/0/1.0
  Index:3, State:0x106, Circuit id: 0x9, Circuit type:2
  LSP interval: 100 ms, Sysid: router1
  Level Adjacencies Priority Metric Hello(s) Hold(s)
    1      0      64      0      9      27
    2      0      64      0      9      27
```

What It Means Check the following output fields and verify that the output shows the intended configuration of IS-IS-enabled interfaces:

- **Interface**—Interface configured for IS-IS
- **State**—Internal implementation information
- **Circuit id**—Circuit identifier
- **Circuit type**—Configured level of IS-IS:
 - 1—Level 1 only
 - 2—Level 2 only
 - 3—Level 1 and Level 2
- **LSP interval**—Time between IS-IS information messages
- **Sysid**—System identifier
- **L or Level**—Type of adjacency:
 - 1—Level 1 only
 - 2—Level 2 only
 - 3—Level 1 and Level 2
- **Adjacencies**—Adjacencies established on the interface
- **Priority**—Priority value established on the interface
- **Metric**—Metric value for the interface
- **Hello(s)**—Intervals between hello PDUs
- **Hold(s)**—Hold time on the interface

See Also For a complete description of show isis interface detail output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Displaying IS-IS Adjacencies

Purpose Display brief information about IS-IS neighbors.

Action From the CLI, enter the `show isis adjacency brief` command.

```
user@host> show isis adjacency brief
IS-IS adjacency database:
  Interface System L State Hold (secs) SNPA
  ge-0/0/0.0 1921.6800.5067 2 Up 13
  ge-0/0/1.0 1921.6800.5067 2 Up 25
  ge-0/0/2.0 1921.6800.5067 2 Up 19
```

What It Means Verify adjacent routers in the IS-IS database.

See Also For a complete description of `show isis adjacency brief` output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Displaying IS-IS Adjacencies in Detail

Purpose Display extensive information about IS-IS neighbors.

Action From the CLI, enter the show isis adjacency extensive command.

```
user@host> show isis adjacency extensive
R1
  Interface: so-0/0/0.0, Level: 2, State: Up, Expires in 25 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 4w6d 19:38:52 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.12.1
  Transition log:
  When                State      Reason
  Wed Jul 13 16:26:11  Up        SeenseIf

R3
  Interface: so-0/0/1.0, Level: 2, State: Up, Expires in 23 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 6w5d 19:07:16 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.23.2
  Transition log:
  When                State      Reason
  Thu Jun 30 16:57:46  Up        SeenseIf

R6
  Interface: so-0/0/2.0, Level: 2, State: Up, Expires in 25 secs
  Priority: 0, Up/Down transitions: 1, Last transition: 6w0d 18:01:18 ago
  Circuit type: 2, Speaks: IP, IPv6
  Topologies: Unicast
  Restart capable: Yes
  IP addresses: 10.1.26.2
  Transition log:
  When                State      Reason
  Tue Jul  5 18:03:45  Up        SeenseIf
```

What It Means Check the following fields and verify adjacency information about IS-IS neighbors:

- **Interface**—Interface through which the neighbor is reachable
- **L or Level**—Configured level of IS-IS:
 - 1—Level 1 only
 - 2—Level 2 only
 - 3—Level 1 and Level 2

An exclamation point before the level number indicates that the adjacency is missing an IP address.

- **State**—Status of the adjacency: Up, Down, New, One-way, Initializing, or Rejected
- **Event**—Message that identifies the cause of a state
- **Down reason**—Reason the adjacency is down

- **Restart capable**—Denotes a neighbor configured for graceful restart
- **Transition log**—List of transitions including **When**, **State**, and **Reason**

See Also For a complete description of `show isis adjacency extensive` output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Chapter 16

Configuring BGP Sessions

Connections between peering networks are typically made through an exterior gateway protocol, most commonly the Border Gateway Protocol (BGP).

You can use either J-Web Quick Configuration or a configuration editor to configure BGP sessions.

This chapter contains the following topics. For more information about BGP, see the *JUNOS Routing Protocols Configuration Guide*.

- BGP Overview on page 441
- Before You Begin on page 442
- Configuring BGP Sessions with Quick Configuration on page 443
- Configuring BGP Sessions with a Configuration Editor on page 444
- Verifying a BGP Configuration on page 452

BGP Overview

BGP is a heavy-duty, secure protocol that must be configured on a per-peer basis. Once a peering session has been configured, BGP uses a TCP connection to establish a session. After a BGP session is established, traffic is passed along the BGP-enabled link.

Although BGP requires a full-mesh topology to share route information, you can use route reflectors and confederations in a large autonomous system (AS) to reduce scaling problems.

BGP Peering Sessions

Unlike RIP and OSPF links, BGP peering sessions must be explicitly configured at both ends. To establish a session between BGP peers, you must manually specify the interface address to which you are establishing a connection. Once this configuration is complete on both ends of a link, a TCP negotiation takes place and a BGP session is established.

The type of the BGP peering session depends on whether the peer is outside or inside the host's autonomous system (AS):

- Peering sessions established with hosts outside the local AS are external sessions. Traffic that passes along such links uses external BGP (EBGP) as its protocol.
- Peering sessions established with hosts within the local AS are internal sessions. Traffic that passes along such links uses internal BGP (IBGP) as its protocol.

To monitor BGP neighbors, see the information about real-time performance monitoring (RPM) in the *J-series Services Router Administration Guide*.

IBGP Full Mesh Requirement

By default, BGP does not readvertise routes that are learned from BGP. To share route information throughout the network, BGP requires a full mesh of internal peering sessions within an AS. To achieve an IBGP full mesh, you configure a direct peering session every host to every other host within the network. These sessions are configured on every router within the network, as type **internal**.

Route Reflectors and Clusters

In larger networks, the overhead needed to implement the IBGP full-mesh requirement is prohibitive. Many networks use route reflectors to avoid having to configure an internal connection to each node for every new router.



NOTE: You must have an Advanced BGP Feature license installed on each Services Router that uses a route reflector. For license details, see the *J-series Services Router Administration Guide*.

A route reflector can readvertise routes learned through BGP to its BGP neighbors. If you define clusters of routers and configure a single router as a route reflector within each cluster, a full mesh is required only between the route reflectors and all their internal peers within the network. The route reflector is responsible for propagating BGP routes throughout the cluster.

For more information about route reflectors, see “Route Reflectors—for Added Hierarchy” on page 378

BGP Confederations

Large ASs can be divided into smaller sub-ASs, which are groups of routers known as confederations. You configure EBGP peering sessions between confederations, and IBGP peering sessions within confederations. Within a confederation, the IBGP full mesh is required. For more information about confederations, see “Confederations—for Subdivision” on page 381

Before You Begin

Before you begin configuring a BGP network, complete the following tasks:

- Establish basic connectivity. See the Getting Started Guide for your router.
- Configure network interfaces. See “Configuring Ethernet, DS1, DS3, and Serial Interfaces” on page 103.

Configuring BGP Sessions with Quick Configuration

J-Web Quick Configuration allows you to create BGP peering sessions. Figure 77 on page 443 shows the Quick Configuration Routing page for BGP.

Figure 77: Quick Configuration Routing Page for BGP

The screenshot shows the Juniper J-Web interface for a J6300 router. The top navigation bar includes Monitor, Configuration, Diagnose, Manage, Events, Alarms, and a user login status. The left sidebar has links for Quick Configuration, View and Edit, History, and Rescue. The main content area is titled 'Quick Configuration' and 'Routing and Protocols'. Under 'Router Identification', the 'Router Identifier' is set to 10.255.0.10. The 'BGP' section has 'Enable BGP' checked. Below this are input fields for 'Autonomous System Number', 'Peer Autonomous System Number', 'Peer Address', and 'Local Address'. At the bottom of the form are 'OK', 'Cancel', and 'Apply' buttons. The footer contains copyright information and the Juniper logo.

To configure a BGP peering session with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Quick Configuration > Routing and Protocols**.
2. Enter information into the Quick Configuration page for BGP, as described in Table 123 on page 444.
3. From the main BGP routing Quick Configuration page, click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration Routing page for BGP, click **Apply**.
 - To apply the configuration and return to the Quick Configuration Routing page, click **OK**.

- To cancel your entries and return to the Quick Configuration Routing page, click **Cancel**.
4. To check the configuration, see “Verifying a BGP Configuration” on page 452.

Table 123: BGP Routing Quick Configuration Summary

Field	Function	Your Action
Router Identification		
Router Identifier (required)	Uniquely identifies the router	Type the Services Router's 32-bit IP address, in dotted decimal notation.
BGP		
Enable BGP	Enables or disables BGP.	<ul style="list-style-type: none"> ■ To enable BGP, select the check box. ■ To disable BGP, clear the check box.
Autonomous System Number	Sets the unique numeric identifier of the AS in which the Services Router is configured.	<p>Type the Services Router's 32-bit AS number, in dotted decimal notation.</p> <p>If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter 3, the value assigned to the AS is 0.0.0.3.</p>
Peer Autonomous System Number	Sets the unique numeric identifier of the AS in which the peer host resides.	<p>Type the peer host's 32-bit AS number, in dotted decimal notation.</p> <p>If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter 3, the value assigned to the AS is 0.0.0.3.</p>
Peer Address	Specifies the IP address of the peer host's interface to which the BGP session is being established.	Type the IP address of the peer host's adjacent interface, in dotted decimal notation.
Local Address	Specifies the IP address of the local host's interface from which the BGP session is being established.	Type the IP address of the local host's adjacent interface, in dotted decimal notation.

Configuring BGP Sessions with a Configuration Editor

To configure the Services Router as a node in a BGP network, you must perform the following tasks marked *(Required)*.

- Configuring Point-to-Point Peering Sessions (Required) on page 445
- Configuring BGP Within a Network (Required) on page 447
- Configuring a Route Reflector (Optional) on page 448
- Configuring BGP Confederations (Optional) on page 451


For information about using the J-Web and CLI configuration editors, see “Using Services Router Configuration Tools” on page 3.

Configuring Point-to-Point Peering Sessions (Required)

To enable BGP traffic across one or more links, you must configure a BGP peering session with the adjacent host. Generally, such sessions are made at network exit points with neighboring hosts outside the autonomous system. Figure 78 on page 445 shows a network with BGP peering sessions.

In the sample network, a Services Router in AS 17 has BGP peering sessions to a group of peers called **external-peers**. Peers A, B, and C reside in AS 22 and have IP addresses 10.10.10.10, 10.10.10.11, and 10.10.10.12. Peer D resides in AS 79, at IP address 10.21.7.2.

Figure 78: Typical Network with BGP Peering Sessions



To configure the BGP peering sessions shown in Figure 78 on page 445:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 124 on page 446.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To configure IBGP sessions between peers, see “Configuring BGP Within a Network (Required)” on page 447.
 - To configure route reflector clusters, see “Configuring a Route Reflector (Optional)” on page 448.
 - To subdivide autonomous systems (ASs), see “Configuring BGP Confederations (Optional)” on page 451.
 - To check the configuration, see “Verifying a BGP Configuration” on page 452.

Table 124: Configuring BGP Peering Sessions


Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Routing options level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Routing options, click Configure or Edit. 	<p>From the [edit] hierarchy level, enter</p> <p>edit routing-options</p>
Set the network's AS number to 17.	<ol style="list-style-type: none"> 1. In the AS Number box, enter 17. 2. Click OK. 	<p>Set the AS number to 17:</p> <p>set autonomous-system 17</p>
Navigate to the Bgp level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. On the main Configuration page next to Protocols, click Configure or Edit. 2. Next to Bgp, click Configure or Edit. 	<p>From the [edit] hierarchy level, enter</p> <p>edit protocols bgp</p>
Create the BGP group external-peers , and add the external neighbor addresses to the group.	<ol style="list-style-type: none"> 1. In the Group box, click Add new entry. 2. In the Group name box, type the name of the group of external BGP peers—external-peers in this case. 3. In the Neighbor box, click Add new entry. 4. In the Address box, type the IP address of an external BGP peer, in dotted decimal notation, and click OK. 5. Repeat Step 3 and Step 4 for each BGP neighbor within the external group that you are configuring. 	<ol style="list-style-type: none"> 1. Create the group external-peers, and add the address of an external neighbor: <p>set group external-peers neighbor 10.10.10.10</p> 2. Repeat Step 1 for each BGP neighbor within the external peer group that you are configuring.
At the group level, set the AS number for the group external-peers to 22. Because three of the peers in this group (peers A, B, and C) reside in one AS, you can set their AS number as a group.	<ol style="list-style-type: none"> 1. In the Peer as box, type the number of the AS in which most peers in the external-peers group reside. 2. Click OK. 	<p>From the [edit protocols bgp] hierarchy level:</p> <p>set group external-peers peer-as 22</p>
At the individual neighbor level, set the AS number for peer D to 79. Because peer D is a member of the group external-peers , it inherits the peer AS number configured at the group level. You must override this value at the individual neighbor level.	<ol style="list-style-type: none"> 1. Under Neighbor, in the Address column, click the IP address of peer D—10.21.7.2 in this case. 2. In the Peer as box, type the AS number of the peer. 3. Click OK. 	<p>From the [edit protocols bgp group external-peers] hierarchy level:</p> <p>set neighbor 10.21.7.2 peer-as 79</p>
Set the group type to external .	<ol style="list-style-type: none"> 1. From the Type list, select external. 2. Click OK. 	<p>From the [edit protocols bgp group external-peers] hierarchy level:</p> <p>set type external</p>

Configuring BGP Within a Network (Required)

To configure BGP sessions between peering networks, you must configure point-to-point sessions between the external peers of the networks. Additionally, you must configure BGP internally to provide a means by which BGP route advertisements can be forwarded throughout the network. Because of the full mesh requirement of IBGP, you must configure individual peering sessions between all internal nodes of the network—unless you use route reflectors or confederations.

Figure 79 on page 447 shows a typical network with external and internal peer sessions. In the sample network, the Services Router in AS 17 is fully meshed with its internal peers in the group `internal-peers`, which have IP addresses starting at 192.168.6.4.

Figure 79: Typical Network with EBGp External Sessions and IBGP Internal Sessions



To configure IBGP in the network shown in Figure 79 on page 447:

1. Configure all external peering sessions as described in “Configuring Point-to-Point Peering Sessions (Required)” on page 445.
2. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
3. Perform the configuration tasks described in Table 125 on page 448.
4. If you are finished configuring the router, commit the configuration.
5. Go on to one of the following procedures:
 - To configure route reflector clusters, see “Configuring a Route Reflector (Optional)” on page 448.
 - To subdivide autonomous systems (ASs), see “Configuring BGP Confederations (Optional)” on page 451.
 - To check the configuration, see “Verifying a BGP Configuration” on page 452.

Table 125: Configuring IBGP Peering Sessions

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Bgp level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to Bgp, click Edit. 	From the [edit] hierarchy level, enter edit protocols bgp
Create the BGP group internal-peers , and add the internal neighbor addresses to the group. You must configure a full IBGP mesh, which requires that each peer be configured with every other internal peer as a BGP neighbor.	<ol style="list-style-type: none"> 1. In the Group box, click Add new entry. 2. In the Group name box, type the name of the group of internal BGP peers—internal-peers in this case. 3. In the Neighbor box, click Add new entry. 4. In the Address box, type the IP address of an internal BGP peer, in dotted decimal notation. 5. Click OK. 6. Repeat Step 3 and Step 4 for each internal BGP peer within the network. 	<ol style="list-style-type: none"> 1. Create the group internal-peers, and add the address of an internal neighbor: set group internal-peers neighbor 192.168.6.4 2. Repeat Step 1 for each internal BGP neighbor within the network.
Set the group type to internal .	<ol style="list-style-type: none"> 1. From the Type list, select internal. 2. Click OK. 	From the [edit protocols bgp group internal-peers] hierarchy level: set type internal
Configure a routing policy to advertise BGP routes.	See the <i>J-series Services Router Advanced WAN Access Configuration Guide</i> .	

Configuring a Route Reflector (Optional)

Because of the IBGP full-mesh requirement, most networks use route reflectors to simplify configuration. Using a route reflector, you group routers into clusters, which are identified by numeric identifiers unique to the AS. Within the cluster, you must configure a BGP session from a single router (the route reflector) to each internal peer. With this configuration, the IBGP full-mesh requirement is met.




NOTE: You must have an Advanced BGP Feature license installed on each Services Router that uses a route reflector. For license details, see the *J-series Services Router Administration Guide*.

Figure 80 on page 449 shows an IBGP network with a Services Router at IP address 192.168.40.4 acting as a route reflector. In the sample network, each router in Cluster 2.3.4.5 has an internal client relationship to the route reflector. To configure the cluster:

- On the Services Router, create an internal group, configure an internal peer (neighbor) relationship to every other router in the cluster, and assign a cluster identifier.
- On each other router you are assigning to the cluster, create the cluster group and configure a client relationship to the route reflector.

Figure 80: Typical IBGP Network Using a Route Reflector



To configure IBGP in the network using the Services Router as a route reflector:

1. Configure all external peering sessions as described in “Configuring Point-to-Point Peering Sessions (Required)” on page 445.
2. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
3. Perform the configuration tasks described in Table 126 on page 450.
4. If you are finished configuring the router, commit the configuration.
5. Go on to one of the following procedures:
 - To subdivide autonomous systems (ASs), see “Configuring BGP Confederations (Optional)” on page 451.
 - To check the configuration, see “Verifying a BGP Configuration” on page 452.

Table 126: Configuring a Route Reflector

Task	J-Web Configuration Editor	CLI Configuration Editor
On the Services Router that you are using as a route reflector, navigate to the Bgp level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to Bgp, click Edit. 	<p>From the [edit] hierarchy level, enter</p> <p>edit protocols bgp</p>
On the Services Router that you are using as a route reflector, create the BGP group cluster-peers , and add to the group the IP addresses of the internal neighbors that you want in the cluster.	<ol style="list-style-type: none"> 1. In the Group box, click Add new entry. 2. In the Group name box, type the name of the group in which the BGP peer is configured—cluster-peers in this case. 3. In the Neighbor box, click Add new entry. 4. In the Address box, type the IP address of a BGP peer, in dotted decimal notation. 5. Click OK. 6. Repeat Step 3 and Step 4 for each BGP neighbor within the cluster that you are configuring. 	<ol style="list-style-type: none"> 1. Create the group cluster-peers, and add the address of an internal neighbor: <p>set group cluster-peers neighbor 192.168.6.4</p> 2. Repeat Step 1 for each BGP neighbor within the cluster that you are configuring.
On the Services Router that you are using as a route reflector, set the group type to internal .	From the Type list, select internal .	<p>From the [edit protocols bgp group internal-peers] hierarchy level:</p> <p>set type internal</p>
On the Services Router that you are using as a route reflector, configure the cluster identifier for the route reflector.	<ol style="list-style-type: none"> 1. In the Cluster box, enter the unique numeric cluster identifier. 2. Click OK. 	<p>Set the cluster identifier:</p> <p>set cluster 2.3.4.5</p>
<p>On the other routers in the cluster, create the BGP group cluster-peers, and add the internal IP address of the route reflector.</p> <p>You do not need to include the neighbor addresses of the other internal peers, or configure the cluster identifier on these route reflector clients. They need only be configured as internal neighbors.</p> <p>NOTE: If the other routers in the network are Services Routers, follow the steps in this row. Otherwise, consult the router documentation for instructions.</p>	<p>On a client Services Router in the cluster:</p> <ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to Bgp, click Edit. 4. In the Group box, click Add new entry. 5. In the Group name box, type the name of the group in which the BGP peer is configured—cluster-peers in this case. 6. In the Neighbor box, click Add new entry. 7. In the Address box, type the IP address of the route reflector, in dotted decimal notation—in this case, 192.168.40.4. 8. Click OK. 	<p>On a client Services Router in the cluster:</p> <ol style="list-style-type: none"> 1. From the [edit] hierarchy level, enter <p>edit protocols bgp</p> 2. Create the group cluster-peers, and add only the route reflector address to the group: <p>set group cluster-peers neighbor 192.168.40.4</p>


Table 126: Configuring a Route Reflector (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a routing policy to advertise BGP routes.	See the <i>J-series Services Router Advanced WAN Access Configuration Guide</i> .	

Configuring BGP Confederations (Optional)

To help solve BGP scaling problems caused by the IBGP full-mesh requirement, you can divide your AS into sub-ASs called confederations. As Figure 81 on page 451 shows, the connections between the sub-ASs are made through EBGP sessions, and the internal connections are made through standard IBGP sessions.

In the sample network, AS 17 has two separate confederations (sub-AS 64512 and sub-AS 64513), each of which has multiple routers. Within a sub-AS, an IGP (OSPF, for example) is used to establish network connectivity with internal peers. Between sub-ASs, an external BGP peering session is established.

Figure 81: Typical Network Using BGP Confederations

To configure the BGP confederations shown in Figure 81 on page 451:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 127 on page 452.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying a BGP Configuration” on page 452.

Table 127: Configuring BGP Confederations

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Routing options level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the J-Web interface, select Configuration > View and Edit > Edit Configuration. 2. Next to Routing options, click Edit. 	From the [edit] hierarchy level, enter edit routing-options
Set the AS number to the sub-AS number 64512. The sub-AS number is a unique AS number that is usually taken from the pool of private AS numbers—64512 through 65535.	<ol style="list-style-type: none"> 1. In the AS Number box, enter the sub-AS number. 2. Click OK. 	Set the sub-AS number: set autonomous-system 64512
Navigate to the Confederation level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. On the main Configuration page next to Routing options, click Edit. 2. Next to Confederation, click Configure. 	From the [edit] hierarchy level, enter edit routing-options confederation
Set the confederation number to the AS number 17.	In the Confederation as box, enter 17.	Set the confederation AS number: set 17
Add the sub-ASs as members of the confederation. Every sub-AS within the AS must be added as a confederation member.	<ol style="list-style-type: none"> 1. Next to Members, click Add new entry. 2. In the Value box, enter the sub-ASs that are members of this confederation. Separate multiple sub-ASs with a space. 	Add members to the confederation: set 17 members 64512 64513
Using EBGp, configure the peering session between the confederations (from Router A to Router B in this example). When setting the peer AS number for these sessions, use the sub-AS number rather than the AS number.	See “Configuring Point-to-Point Peering Sessions (Required)” on page 445.	
Using IBGP, configure internal sessions within a sub-AS. You can configure an IBGP full mesh, or you can configure a route reflector.	<ul style="list-style-type: none"> ■ To configure an IBGP full mesh, see “Configuring BGP Within a Network (Required)” on page 447. ■ To configure a route reflector, see “Configuring a Route Reflector (Optional)” on page 448. 	

Verifying a BGP Configuration

To verify a BGP configuration, perform these tasks:

- Verifying BGP Neighbors on page 454
- Verifying BGP Groups on page 455

- Verifying BGP Summary Information on page 456
- Verifying Reachability of All Peers in a BGP Network on page 457

Verifying BGP Neighbors

Purpose Verify that BGP is running on configured interfaces and that the BGP session is active for each neighbor address.

Action From the CLI, enter the `show bgp neighbor` command.

```
user@host> show bgp neighbor
Peer: 10.255.245.12+179 AS 35 Local: 10.255.245.13+2884 AS 35
  Type: Internal State: Established (route reflector client)Flags: Sync
  Last State: OpenConfirm Last Event: RecvKeepAlive
  Last Error: None
  Options: Preference LocalAddress HoldTime Cluster AddressFamily Rib-group Refresh

Address families configured: inet-vpn-unicast inet-labeled-unicast
Local Address: 10.255.245.13 Holdtime: 90 Preference: 170
Flags for NLRI inet-vpn-unicast: AggregateLabel
Flags for NLRI inet-labeled-unicast: AggregateLabel
Number of flaps: 0
Peer ID: 10.255.245.12 Local ID: 10.255.245.13 Active Holdtime: 90
Keepalive Interval: 30
NLRI advertised by peer: inet-vpn-unicast inet-labeled-unicast
NLRI for this session: inet-vpn-unicast inet-labeled-unicast
Peer supports Refresh capability (2)
Restart time configured on the peer: 300
Stale routes from peer are kept for: 60
Restart time requested by this peer: 300
NLRI that peer supports restart for: inet-unicast inet6-unicast
NLRI that restart is negotiated for: inet-unicast inet6-unicast
NLRI of received end-of-rib markers: inet-unicast inet6-unicast
NLRI of all end-of-rib markers sent: inet-unicast inet6-unicast
Table inet.0 Bit: 10000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 4
  Received prefixes: 6
  Suppressed due to damping: 0
Table inet6.0 Bit: 20000
  RIB State: restart is complete
  Send state: in sync
  Active prefixes: 0
  Received prefixes: 2
  Suppressed due to damping: 0
Last traffic (seconds): Received 3 Sent 3 Checked 3
Input messages: Total 9 Updates 6 Refreshes 0 Octets 403
Output messages: Total 7 Updates 3 Refreshes 0 Octets 365
Output Queue[0]: 0
Output Queue[1]: 0
Trace options: detail packets
Trace file: /var/log/bgpr size 131072 files 10
```

What It Means The output shows a list of the BGP neighbors with detailed session information. Verify the following information:

- Each configured peering neighbor is listed.
- For State, each BGP session is Established.
- For Type, each peer is configured as the correct type (either internal or external).
- For AS, the AS number of the BGP neighbor is correct.

See Also For a complete description of `show bgp neighbor` output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying BGP Groups

Purpose Verify that the BGP groups are configured correctly.

Action From the CLI, enter the `show bgp group` command.

```
user@host> show bgp group
Group Type: Internal  AS: 10045      Local AS: 10045
Name: pe-to-asbr2      Flags: Export Eval
Export: [ match-all ]
Total peers: 1      Established: 1
10.0.0.4+179
bgp.l3vpn.0: 1/1/0
vpn-green.inet.0: 1/1/0

Groups: 1  Peers: 1  External: 0  Internal: 1  Down peers: 0  Flaps: 0
Table      Tot Paths  Act Paths  Suppressed  History Damp State  Pending
bgp.l3vpn.0      1          1          0          0          0          0
```

What It Means The output shows a list of the BGP groups with detailed group information. Verify the following information:

- Each configured group is listed.
- For **AS**, each group's remote AS is configured correctly.
- For **Local AS**, each group's local AS is configured correctly.
- For **Group Type**, each group has the correct type (either internal or external).
- For **Total peers**, the expected number of peers within the group is shown.
- For **Established**, the expected number of peers within the group have BGP sessions in the **Established** state.
- The IP addresses of all the peers within the group are present.

See Also For a complete description of `show bgp group` output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying BGP Summary Information

Purpose Verify that the BGP configuration is correct.

Action From the CLI, enter the `show bgp summary` command.

```
user@host> show bgp summary
Groups: 1 Peers: 3 Down peers: 0
Table Tot Paths Act Paths Suppressed History Damp State Pending
inet.0 6 4 0 0 0 0
Peer AS InPkt OutPkt OutQ Flaps Last Up/Dwn
State|#Active/Received/Damped...
10.0.0.2 65002 88675 88652 0 2 42:38 2/4/0
0/0/0
10.0.0.3 65002 54528 54532 0 1 2w4d22h 0/0/0
0/0/0
10.0.0.4 65002 51597 51584 0 0 2w3d22h 2/2/0
0/0/0
```

What It Means The output shows a summary of BGP session information. Verify the following information:

- For **Groups**, the total number of configured groups is shown.
- For **Peers**, the total number of BGP peers is shown.
- For **Down Peers**, the total number of unestablished peers is 0. If this value is not zero, one or more peering sessions are not yet established.
- Under **Peer**, the IP address for each configured peer is shown.
- Under **AS**, the peer AS for each configured peer is correct.
- Under **Up/Dwn State**, the BGP state reflects the number of paths received from the neighbor, the number of these paths that have been accepted, and the number of routes being damped (such as 0/0/0). If the field is **Active**, it indicates a problem in the establishment of the BGP session.

See Also For a complete description of `show bgp summary` output, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying Reachability of All Peers in a BGP Network

Purpose By using the ping tool on each peer address in the network, verify that all peers in the network are reachable from each Services Router.

Action For each Services Router in the BGP network:

1. In the J-Web interface, select **Diagnose > Ping Host**.
2. In the Remote Host box, type the name of a host for which you want to verify reachability from the Services Router.
3. Click **Start**. Output appears on a separate page.

```
PING 10.10.10.10 : 56 data bytes
64 bytes from 10.10.10.10: icmp_seq=0 ttl=255 time=0.382 ms
64 bytes from 10.10.10.10: icmp_seq=1 ttl=255 time=0.266 ms
```

What It Means If a host is active, it generates an ICMP response. If this response is received, the round-trip time is listed in the **time** field.

See Also For more information about using the J-Web interface to ping a host, see the *J-series Services Router Administration Guide*.
For information about the **ping** command, see the *J-series Services Router Administration Guide* or the *JUNOS System Basics and Services Command Reference*.

Part 4

Index

- Index on page 461

Index

Symbols

#, comments in configuration statements.....xxii
 (), in syntax descriptions.....xxii
 1-port four-wire mode, SHDSL *See* ATM-over-SHDSL interfaces
 2-port two-wire mode, SHDSL *See* ATM-over-SHDSL interfaces
 802.3ad, Gigabit Ethernet.....119
 < >, in syntax descriptions.....xxii
 [], in configuration statements.....xxii
 { }, in configuration statements.....xxii
 | (pipe), in syntax descriptions.....xxii

A

AAL5 multiplex encapsulation.....178
 ATM-over-ADSL for PPP-over-ATM (PPPoA)
 interfaces.....168
 ATM-over-ADSL interfaces.....162
 ATM-over-SHDSL interfaces.....172
 ABM (Asynchronous Balance Mode), HDLC.....90
 ABRs *See* area border routers
 access concentrator
 as a PPPoE server.....192
 naming for PPPoE (configuration editor).....202
 naming for PPPoE (Quick Configuration).....198
 activation priority
 description.....319
 range and default.....330
 active routes, versus passive routes.....385
 Add button.....8
 Add new entry link.....12
 address resolution protocol *See* ARP; static ARP entries
 addresses.....431
 BGP external peer address (configuration editor).....446
 BGP internal peer address (configuration editor).....448
 BGP local address (Quick Configuration).....444
 BGP peer address (Quick Configuration).....444
 IS-IS NETs.....369
 See also NETs

IS-IS NSAP addresses.....431
 physical, in data link layer.....52
 See also IPv4 addressing; IPv6 addressing
 agencies, IS-IS
 hello PDUs.....369
 See also IS-IS
 verifying.....437
 verifying (detail).....438
 ADSL interfaces *See* ATM-over-ADSL interfaces
 ADSL ports *See* ATM-over-ADSL interfaces
 ADSL2+ operating mode.....164, 167
 advertisements *See* LSAs; route advertisements
 aggregated Ethernet, on Gigabit Ethernet ports.....119
 aggregated virtual circuits (AVCs), with MLFR
 FRF.15.....284
 See also MLFR FRF.15; multilink bundles
 aggregation, route.....358
 alternate mark inversion *See* AMI encoding
 always compare, BGP MED option.....377
 AMI (alternate mark inversion) encoding
 E1.....108
 overview.....59
 T1.....123
 analog media module *See* TIM514
 Annex A PIMs
 ATM-over-ADSL interfaces.....164
 See also ATM-over-ADSL interfaces
 ATM-over-SHDSL interfaces.....174
 See also ATM-over-SHDSL interfaces
 ATM-over-SHDSL modes.....169
 G.SHDSL PIMs, setting annex type on.....173, 176
 operating modes (configuration editor).....167
 operating modes (Quick Configuration).....164
 standards supported.....73
 Annex B PIMs
 ATM-over-ADSL interfaces.....164
 See also ATM-over-ADSL interfaces
 ATM-over-SDSL interfaces.....174
 See also ATM-over-SHDSL interfaces
 ATM-over-SHDSL modes.....169
 G.SHDSL PIMs, setting annex type on.....173, 176
 operating modes (configuration editor).....167
 operating modes (Quick Configuration).....164
 standards supported.....73
 ANSI DMT operating mode.....164, 167
 ANSI T1.413 Issue II operating mode.....164, 167

- anycast IPv6 addresses.....95
- Apply button.....9
- area border routers
 - adding interfaces.....420
 - area ID (configuration editor).....420
 - backbone area *See* backbone area
 - backbone area interface.....420
 - description.....365
- areas *See* area border routers; backbone area; IS-IS, areas; NSSAs; stub areas
- ARM (Asynchronous Response Mode), HDLC.....90
- ARP (address resolution protocol), for static ARP entries
 - for Fast Ethernet subnets.....114
See also static ARP entries
 - for Gigabit Ethernet subnets.....117
See also static ARP entries
 - publish (responding to ARP requests), on Fast Ethernet subnets.....114
 - publish (responding to ARP requests), on Gigabit Ethernet subnets.....118
- AS path
 - description.....375
 - forcing by MED.....376
 - role in BGP route selection.....373
- ASs (autonomous systems)
 - area border routers.....365
 - AS number (configuration editor).....446
 - AS number (Quick Configuration).....444
 - breaking into confederations.....381
 - description.....355
 - group AS number (configuration editor).....446
 - individual AS number (configuration editor).....446
 - IS-IS networks.....368
 - sample BGP confederation.....451
 - stub areas *See* stub areas
 - sub-AS number.....452
- asymmetric digital subscriber line (ADSL) *See* ATM-over-ADSL interfaces
- Asynchronous Balance Mode (ABM), HDLC.....90
- asynchronous networks
 - data stream clocking.....80
 - explicit clocking signal transmission.....80
 - overview.....79
- Asynchronous Response Mode (ARM), HDLC.....90
- Asynchronous Transfer Mode (ATM) interfaces *See* ATM-over-ADSL interfaces; ATM-over-SHDSL interfaces
- at-0/0/0 *See* ATM-over-ADSL interfaces; ATM-over-SHDSL interfaces
- ATM interfaces *See* ATM-over-ADSL interfaces; ATM-over-SHDSL interfaces
- ATM NLPID encapsulation
 - ATM-over-ADSL interfaces.....162, 168
 - ATM-over-SHDSL interfaces.....172, 178
- ATM PPP over AAL5 LLC encapsulation
 - ATM-over-ADSL for PPP-over-ATM (PPPoA) interfaces.....162, 168
 - ATM-over-SHDSL interfaces.....172, 178
- ATM PVC encapsulation
 - ATM-over-ADSL interfaces.....163, 167
 - ATM-over-SHDSL interfaces.....173, 176
- ATM SNAP encapsulation
 - ATM-over-ADSL interfaces.....162, 168
 - ATM-over-SHDSL interfaces.....172, 178
- ATM VC multiplex encapsulation
 - ATM-over-ADSL interfaces.....162, 168
 - ATM-over-SHDSL interfaces.....172, 178
- ATM-over-ADSL interfaces.....164
 - adding.....164
 - ADSL overview.....72
 - ADSL systems.....73
 - ADSL topology.....73
 - ADSL2.....74
 - ADSL2 +74
 - ATM interface type.....74
 - CHAP for PPPoA.....179
 - CHAP for PPPoE.....204
 - description.....159
 - encapsulation types, logical (configuration editor).....168
 - encapsulation types, logical (Quick Configuration).....162
 - encapsulation types, physical (configuration editor).....167
 - encapsulation types, physical (Quick Configuration).....163
 - logical properties (configuration editor).....167
 - logical properties (Quick Configuration).....161
 - MTU default and maximum values.....81
 - operating modes (configuration editor).....167
 - operating modes (Quick Configuration).....164
 - physical properties.....165
 - PPPoE configuration.....201
 - PPPoE encapsulation.....200
 - PPPoE session on.....193
 - preparation.....158
 - Quick Configuration.....159
 - statistics.....184
 - VCI162, 169
 - verifying.....181
 - verifying a PPPoA configuration.....185
 - verifying a PPPoE configuration.....206, 207
 - VPI163, 166
 - See also* PPPoE; PPPoE over ATM-over-ADSL; PPPoE over ATM-over-SHDSL
- ATM-over-SHDSL interfaces.....74
 - 1-port four-wire mode.....170
 - 1-port four-wire mode, setting.....173, 175
 - 2-port two-wire mode, overview.....170
 - 2-port two-wire mode, setting.....173, 175

- adding.....174
- annex type, setting.....173, 176
- ATM interface type.....74
- CHAP for PPPoA.....179
- CHAP for PPPoE.....204
- description.....169
- encapsulation types, logical (configuration editor).....178
- encapsulation types, logical (Quick Configuration).....172
- encapsulation types, physical.....173
- encapsulation types, physical (configuration editor).....176
- encapsulation types, physical (Quick Configuration).....173
- line speed.....173
- logical properties (configuration editor).....177
- logical properties (Quick Configuration).....172
- loopback testing.....174
- MTU default and maximum values.....81
- overview.....74
- PPPoE configuration.....201
- PPPoE encapsulation.....200
- PPPoE session on.....193
- preparation.....158
- Quick Configuration.....170
- SNEXT threshold.....174, 177
- SNR margin.....174, 177
- statistics.....189
- status.....188
- VCI173, 179
- verifying.....186
- verifying a PPPoE configuration.....206, 207
- VPI173, 176
- “dying gasp”.....189
- See also* G.SHDSL PIMs
- authentication
 - CHAP, for PPPoE interfaces.....194
 - OSPF, MD5.....425
 - OSPF, plain-text passwords.....425
 - RIPv2, MD5.....405
 - RIPv2, plain-text passwords.....404
- auto operating mode.....164, 167
- autonegotiation, Gigabit Ethernet.....119
- autonomous systems *See* ASs
- Avaya Communication Manager (CM)
 - CAC-BL requirement for WANs.....319
 - description.....318
- Avaya IG550 Integrated Gateway
 - Avaya Communication Manager (CM).....318
 - Avaya manuals, list of.....317
 - description.....314
 - See also* Avaya VoIP modules
 - dynamic CAC *See* dynamic CAC
 - TGM550–JUNOS compatibility.....320
- Avaya Media Gateway Controller (MGC)
 - Avaya Communication Manager (CM).....318
 - Avaya manuals, list of.....317
 - description.....314
 - dynamic CAC *See* dynamic CAC
 - MGC list.....317
 - See also* MGC list
 - supported models.....317
 - verifying MGC list.....341
- Avaya MGC *See* Avaya Media Gateway Controller
- Avaya VoIP
 - Avaya Communication Manager (CM).....318
 - Avaya manuals, list of.....317
 - Avaya Media Gateway Controllers
 - supported.....317
 - bandwidth management *See* dynamic CAC
 - configuration overview.....321
 - Disk-on-Key configuration.....323
 - dynamic CAC *See* dynamic CAC
 - EPW configuration.....323
 - interfaces.....315
 - See also* Avaya VoIP modules
 - IP addressing guidelines.....320
 - modules *See* Avaya VoIP modules
 - network.....315
 - overview.....313
 - prerequisites.....322
 - Quick Configuration.....325
 - TGM550–JUNOS compatibility.....320
 - troubleshooting.....342
 - typical topology.....315
 - verifying available bandwidth.....341
 - verifying configuration.....338
 - version incompatibility, correcting.....342
- Avaya VoIP modules
 - accessing the router from.....337
 - administration.....333
 - Avaya CLI access.....333
 - Avaya Communication Manager (CM).....318
 - Avaya manuals, list of.....317
 - CLI access requirements.....334
 - configuration overview.....321
 - console connection.....334
 - Disk-on-Key configuration.....323
 - dynamic CAC *See* dynamic CAC
 - interface types.....315
 - IP address, modifying (configuration editor).....332
 - JUNOS configurability.....316
 - MGC list, adding.....329
 - MGC list, clearing.....330
 - non-hot-swappability.....316
 - overview.....316
 - prerequisites.....322
 - resetting TGM550.....337
 - saving the configuration.....338
 - SSH connection.....335

Telnet access.....	335
TGM550 IP address, setting (configuration editor).....	327
TGM550-JUNOS compatibility.....	320
AVCs (aggregated virtual circuits), multilink bundles, with MLFR FRF.15.....	284
<i>See also</i> MLFR FRF.15; multilink bundles	

B

B-channel allocation order, on ISDN PRI interfaces.....	149
B-channels	
description.....	75
naming convention.....	214, 215
verifying.....	249
B8ZS encoding.....	59
backbone area	
area ID (configuration editor).....	417
area ID (Quick Configuration).....	414
area type (Quick Configuration).....	414
configuring.....	415
description.....	366
interface.....	420
backoff algorithm, collision detection.....	55
backup connection, ISDN.....	211
backward-explicit congestion notification (BECN)	
bits.....	85
bandwidth for Avaya VoIP, managing <i>See</i> dynamic CAC	
bandwidth on demand, ISDN	
dialer interface (configuration editor).....	234
dialer pool.....	238
ISDN BRI interface (configuration editor).....	238
overview.....	234
BBL (bearer bandwidth limit)	
description.....	319
range and default.....	330
reported (RBBL), description.....	319
verifying available bandwidth.....	341
bc-0/0/0	
ISDN BRI interface.....	214
<i>See also</i> ISDN BRI interfaces; ISDN PRI interfaces	
ISDN PRI interface.....	215
bearer bandwidth limit <i>See</i> BBL	
BECN (backward-explicit congestion notification)	
bits.....	85
BERTs (bit error rate tests)	
on channelized interfaces (configuration editor).....	144
overview.....	79

BGP (Border Gateway Protocol)	
AS number (Quick Configuration).....	444
<i>See also</i> ASs (autonomous systems), AS number	
AS path.....	375
<i>See also</i> AS path	
confederations <i>See</i> BGP confederations	
enabling (Quick Configuration).....	444
external.....	372
<i>See also</i> EBGp	
external group type (configuration editor).....	446
external neighbor (peer) address (configuration editor).....	446
full mesh requirement.....	373, 442
internal.....	372
<i>See also</i> IBGP	
internal group type (configuration editor).....	448
internal neighbor (peer) address (configuration editor).....	448
local address (Quick Configuration).....	444
local preference.....	374
MED metric.....	376
<i>See also</i> MED	
origin value.....	376
overview.....	371, 441
peer address (Quick Configuration).....	444
peer AS number (Quick Configuration).....	444
peering sessions <i>See</i> BGP peers; BGP sessions	
point-to-point internal peer session (configuration editor).....	447
point-to-point peer session (configuration editor).....	445
Quick Configuration.....	443
requirements.....	442
route reflectors <i>See</i> BGP route reflectors	
route selection process.....	373
<i>See also</i> route selection	
router ID (Quick Configuration).....	444
routing policy (configuration editor).....	448
<i>See also</i> routing policies	
sample BGP peer network.....	445
sample confederation.....	451
sample full mesh.....	447
sample route reflector.....	449
scaling techniques.....	378
session establishment.....	372
session maintenance.....	372
verifying BGP configuration.....	456
verifying BGP groups.....	455
verifying BGP peers (neighbors).....	454
verifying peer reachability.....	457
BGP confederations	
confederation members.....	452
confederation number.....	452
creating (configuration editor).....	451
description.....	381, 442

- sample network.....451
- sub-AS number.....452
- BGP groups
 - cluster identifier (configuration editor).....450
 - confederations (configuration editor).....451
 - external group type (configuration editor).....446
 - external, creating (configuration editor).....446
 - group AS number (configuration editor).....446
 - internal group type (configuration editor).....448
 - internal, creating (configuration editor).....448
 - internal, creating for a route reflector (configuration editor).....450
 - verifying.....455
- BGP messages
 - to establish sessions.....372
 - update, to maintain sessions.....372
- BGP neighbors *See* BGP peers
- BGP page.....443
- BGP peers
 - directing traffic by local preference.....374
 - external (configuration editor).....445
 - internal (configuration editor).....447
 - internal, sample full mesh.....447
 - internal, sample route reflector.....449
 - monitor probes.....442
 - peer address (Quick Configuration).....444
 - peer address, role in route selection.....374
 - peer AS number (Quick Configuration).....444
 - point-to-point connections.....371
 - routing policy (configuration editor).....448
 - See also* routing policies
 - sample peer network.....445
 - sessions between peers.....441
 - verifying.....454, 456
 - verifying reachability.....457
- BGP route reflectors
 - cluster (configuration editor).....450
 - cluster identifier (configuration editor).....450
 - cluster of clusters.....380
 - clusters, role in route selection.....374
 - creating (configuration editor).....448
 - description.....378, 442
 - group type (configuration editor).....450
 - multiple clusters.....379
 - sample IBGP network.....449
- BGP sessions
 - configured at both ends.....441
 - establishment.....372
 - maintenance.....372
 - point-to-point external (configuration editor).....445
 - point-to-point internal (configuration editor).....447
 - sample peering session.....371
 - types.....442
- bipolar with 8-zero substitution (B8ZS) encoding.....59
- bit error rate tests (BERTs) *See* BERTs
- bit stuffing.....62
- Border Gateway Protocol *See* BGP
- br-0/0/0.....214
 - See also* B-channels; ISDN BRI interfaces
- braces, in configuration statements.....xxii
- brackets
 - angle, in syntax descriptions.....xxii
 - square, in configuration statements.....xxii
- BRI media module *See* TIM521
- bridges, on LAN segments.....56
- buffer size, for Q0 on LFI constituent links.....270
- built-in Ethernet interfaces.....101
 - See also* Fast Ethernet ports; Gigabit Ethernet ports
- buttons.....13
 - Add (Quick Configuration).....8
 - Apply (Quick Configuration).....9
 - Cancel (J-Web configuration editor).....13
 - Cancel (Quick Configuration).....9
 - Commit (J-Web configuration editor).....13
 - CONFIG *See* CONFIG button
 - Delete (Quick Configuration).....8
 - Discard (J-Web configuration editor).....13
 - OK (J-Web configuration editor).....13
 - OK (Quick Configuration).....9
 - Refresh (J-Web configuration editor).....13
 - RESET CONFIG *See* RESET CONFIG button
 - See also* option buttons
- C**
- C-bit parity frame format
 - enable or disable on T3 ports.....127
 - overview.....64
- cables
 - T1 cable length.....124
 - T3 cable length.....127
- CAC *See* dynamic CAC
- CAC-BL requirement for dynamic CAC.....319
- call admission control *See* dynamic CAC
- Call Admission Control: Bandwidth Limitation (CAC-BL), requirement for dynamic CAC.....319
- call setup, ISDN.....77
- callback, ISDN
 - dialer interface (configuration editor).....240
 - encapsulation matching.....239
 - overview.....238
 - rejecting incoming calls (configuration editor).....243
 - screening incoming calls (configuration editor).....242
 - voice not supported.....238
- calling number, ISDN.....218, 225
- Cancel button
 - J-Web configuration editor.....13
 - Quick Configuration.....9
- canceled a commit.....32

carrier sense multiple access with collision detection (CSMA/CD).....	54
ccc protocol family.....	91
Challenge Handshake Authentication Protocol <i>See</i> CHAP	
channel number, in interface name.....	49
channel service unit (CSU) device.....	88
channelized E1 interfaces	
adding.....	142
BERTs (configuration editor).....	144
clear-channel operation (configuration editor).....	143
drop-and-insert (configuration editor).....	145, 146
FAQ.....	153
framing (configuration editor).....	144
ISDN PRI (configuration editor).....	147
MTU default and maximum values.....	81
number of channels supported.....	142
overview.....	61
<i>See also</i> channelized E1 ports	
verifying.....	151
verifying clear-channel interfaces.....	152
channelized E1 ports	
clocking (configuration editor).....	143, 146
clocking for drop-and-insert.....	145
configuring.....	142
drop-and-insert clock combinations	
internal.....	153
FAQ.....	153
ISDN PRI (configuration editor).....	147
link hold time (configuration editor).....	143
overview.....	61
<i>See also</i> channelized E1 interfaces	
per-unit scheduler (configuration editor).....	143
trace options (configuration editor).....	144
channelized interfaces <i>See</i> channelized E1 interfaces;	
channelized T1 interfaces; ISDN PRI interfaces	
channelized T1 interfaces	
adding.....	142
BERTs (configuration editor).....	144
clear-channel operation (configuration editor).....	143
drop-and-insert (configuration editor).....	145, 146
FAQ.....	153
framing (configuration editor).....	144
ISDN PRI (configuration editor).....	147
line encoding (configuration editor).....	144
MTU default and maximum values.....	81
number of channels supported.....	142
overview.....	61
<i>See also</i> channelized T1 ports	
verifying.....	151
verifying clear-channel interfaces.....	152
channelized T1 ports	
clocking (configuration editor).....	143, 146
clocking for drop-and-insert.....	145
configuring.....	142
drop-and-insert clock combinations	
external.....	153
FAQ.....	153
ISDN PRI (configuration editor).....	147
link hold time (configuration editor).....	143
overview.....	61
<i>See also</i> channelized T1 interfaces	
per-unit scheduler (configuration editor).....	143
trace options (configuration editor).....	144
channelized T1/E1/ISDN PRI interfaces,	
overview.....	61, 140
<i>See also</i> channelized E1 interfaces; channelized	
T1 interfaces; ISDN PRI interfaces	
channelized T1/E1/ISDN PRI ports, overview.....	61
<i>See also</i> channelized E1 ports; channelized T1	
ports; ISDN PRI interfaces	
CHAP (Challenge Handshake Authentication Protocol)	
E1 local identity.....	108
E3 local identity.....	110
enabling for PPPoA.....	179
enabling for PPPoE (configuration editor).....	204
enabling for PPPoE (Quick Configuration).....	197
enabling on ATM-over-ADSL interfaces.....	179
enabling on ATM-over-SHDLS interfaces.....	179
enabling on E1.....	108
enabling on E3.....	110
enabling on serial interfaces.....	129
enabling on T1.....	123
enabling on T3.....	127
local identity.....	108, 110
overview.....	86
PPP links.....	86
PPPoE.....	194
serial interface local identity.....	130
T1 local identity.....	123
T3 local identity.....	127
CHAP secret <i>See</i> CHAP, local identity	
checksum	
E1 frame.....	108
E3 frame.....	112
overview.....	81
T1 frame.....	124
T3 frame.....	127
Cisco NLPID encapsulation	
ATM-over-ADSL interfaces.....	162, 168
ATM-over-SHDLS interfaces.....	172, 178
Cisco non-deterministic, BGP MED option.....	378
class of service <i>See</i> CoS components for link services	
classful addressing.....	92
classifiers, defining.....	277
clear system commit command.....	32
clear-channel interface on channelized port	
configuring.....	143
verifying.....	152

- CLI configuration editor
 - activating a configuration.....31
 - ATM-over-ADSL interfaces.....164
 - ATM-over-SHDSL interfaces.....174
 - Avaya VoIP.....327
 - BGP.....444
 - channelized E1 interfaces.....142
 - channelized T1 interfaces.....142
 - CHAP on ATM-over-ADSL interfaces.....179
 - CHAP on ATM-over-SHDSL interfaces.....179
 - command summary.....5
 - committing files.....30
 - confirming a configuration.....31
 - CRTP.....289
 - exiting.....22
 - IS-IS.....433
 - ISDN connections.....223
 - LFI.....273
 - managing files.....35
 - MLPPP bundles.....273
 - modifying a configuration.....25
 - network interfaces.....131
 - network interfaces, adding.....131
 - network interfaces, deleting.....134
 - OSPF.....414
 - PPPoE.....198
 - PPPoE over ATM-over-ADSL.....198
 - PPPoE over ATM-over-SHDSL.....198
 - RIP.....398
 - saving files.....37
 - starting.....22
 - static routes.....387
 - using show commands with.....34
 - verifying a configuration.....31
 - VoIP.....327
- CLI, Avaya VoIP, accessing.....333
- clickable configuration.....10
 - committing.....14
 - discarding changes.....13
 - viewing and editing.....10*See also* J-Web configuration editor
- clock rate, serial interface
 - DTE default reduction.....69
 - values.....131
- clocking
 - channelized ports.....143
 - data stream clocking.....80
 - E1.....107
 - E3.....110
 - explicit clocking signal transmission.....80
 - overview.....79
 - possible combinations for drop-and-insert.....153
 - requirement for drop-and-insert.....145
 - serial interface.....130
 - serial interface, inverting the transmit
 - clock.....69, 130
 - serial interface, modes.....68
 - T1.....122
 - T3.....126
- clusters *See* BGP route reflectors
- CM, Avaya *See* Avaya Communication Manager
- collision detection
 - backoff algorithm.....55
 - overview.....54
- combined stations, HDLC.....90
- comments, in configuration statements.....xxii
- commit and-quit command.....31
- commit at command.....32
- Commit button.....13
- commit check command.....31
- commit command.....30
- commit confirmed command.....32
- committed configuration
 - activating (CLI configuration editor).....31
 - canceling a commit (CLI configuration editor).....32
 - comparing two configurations.....19
 - confirming (CLI configuration editor).....31
 - description.....4
 - methods.....18
 - replacing (CLI configuration editor).....32
 - rescue configuration (CLI configuration
 - editor).....33
 - rescue configuration (J-Web).....21
 - scheduling (CLI configuration editor).....32
 - storage location.....4
 - summaries.....17
 - verifying (CLI configuration editor).....31
 - viewing previous (CLI configuration editor).....33
- Communication Manager (CM), Avaya *See* Avaya Communication Manager
- complete sequence number PDU (CSNP).....370
- Compressed Real-Time Transport Protocol *See* CRTP
- confederations *See* BGP confederations
- CONFIG button
 - default behavior.....21, 33
 - disabling.....34
 - return to factory configuration.....21, 33
- config-button < no-rescue > < no-clear >
 - statement.....34
- configuration
 - activating (CLI configuration editor).....31
 - adding a statement (CLI configuration editor).....25
 - basic.....7
 - changing part of a file (CLI configuration
 - editor).....35
 - CLI commands.....5
 - CLI configuration mode.....22
 - committed.....4
 - committing (CLI configuration editor).....30
 - committing (J-Web).....14
 - committing as a text file, with caution
 - (J-Web).....14

confirming (CLI configuration editor).....	31
copying a statement.....	27
deactivating a statement.....	30
deleting a statement.....	27
deleting with the CONFIG or RESET CONFIG button.....	21, 33
disabling CONFIG or RESET CONFIG button.....	34
discarding changes (J-Web).....	13
downloading (J-Web).....	20
editing (J-Web).....	10
editing as a text file, with caution (J-Web).....	14
finding and replacing values.....	26
history.....	16
<i>See also</i> configuration history	
inserting an identifier.....	28
J-Web options.....	5
loading new (CLI configuration editor).....	35
loading previous (CLI configuration editor).....	32
loading previous (J-Web).....	21
locked, with the configure exclusive command.....	23
managing files (CLI configuration editor).....	35
managing files (J-Web).....	16
merging (CLI configuration editor).....	35
modifying (CLI configuration editor).....	25
modifying a statement (CLI configuration editor).....	25
overriding (CLI configuration editor).....	35
renaming an identifier.....	28
replacing configuration statements (CLI configuration editor).....	36
replacing values.....	26
requirements.....	7
rescuing (CLI configuration editor).....	33
rescuing (J-Web).....	21
rollback (CLI configuration editor).....	32
rollback (J-Web).....	21
saving (CLI configuration editor).....	37
search-and-replace operation.....	26
uploading (J-Web).....	15
users-editors, viewing.....	18
verifying (CLI configuration editor).....	31
viewing as a text file (J-Web).....	9
configuration database, summary.....	17
configuration hierarchy, navigating.....	23
configuration history	
comparing files.....	19
database summary.....	17
displaying.....	16
downloading files.....	20
summary.....	17
users-editors, viewing.....	18
Configuration History page.....	17
configuration mode	
entering and exiting.....	22
using show commands in.....	34
configuration text	
editing and committing, with caution.....	14
viewing.....	9
configuration tools.....	3
<i>See also</i> CLI configuration editor; configuration; configuration history; J-Web configuration editor; Quick Configuration	
configure command.....	23
configure exclusive command.....	23
Configure link.....	12
configure private command.....	23
confirming a configuration.....	31
congestion control, for Frame Relay, with DE bits.....	85
connection process	
ISDN BRI interfaces.....	77
LCP, for PPP.....	85
serial interfaces.....	67
connectivity	
bidirectional (BGP).....	371
bidirectional (OSPF).....	364
unidirectional (RIP).....	363
console port connection to TGM550.....	334
constituent links, queuing <i>See</i> queuing with LFI	
conventions	
for interface names.....	47
how to use this guide.....	xx
notice icons.....	xxi
text and syntax.....	xxi
copy command.....	27
copy running-config startup-config command.....	338
CoS components for link services	
applying on constituent links.....	302
buffer size for Q0.....	270
classifiers (configuration editor).....	277
forwarding classes (configuration editor).....	277
overview.....	269
scheduler maps (configuration editor).....	279
scheduling priority.....	270
shaping rate.....	269
shaping rates (configuration editor).....	283
troubleshooting.....	301
verifying.....	300
verifying configuration.....	294
cost, of a network path <i>See</i> path cost metrics	
CPE device, Services Router as, with PPPoE.....	191
<i>See also</i> PPPoE	
CRC (cyclic redundancy check).....	80
CRTP (Compressed Real-Time Transport Protocol)	
E1 interfaces (configuration editor).....	289
overview.....	102, 266
queuing behavior.....	268
T1 interfaces (configuration editor).....	289
CSMA/CD (carrier sense multiple access with collision detection).....	54
CSNP (complete sequence number PDU).....	370
CSU (channel service unit) device.....	88

curly braces, in configuration statements.....xxii
 customer premises equipment (CPE) device, Services
 Router as, with PPPoE.....191
 See also PPPoE
 customer support.....xxv
 contacting JTAC.....xxv
 cyclic redundancy check (CRC).....80

D

D-channel
 description.....75
 naming convention.....214, 215
 verifying.....251
 D4 framing.....59
 data communications equipment *See* DCE
 data inversion
 E1.....108
 T1.....123
 data link layer
 error notification.....52
 flow control.....52
 frame sequencing.....52
 MAC addresses.....52
 network topology.....52
 physical addressing.....52
 purpose.....52
 sublayers.....52
 data packets
 integrating with voice, with drop-and-insert.....145
 LFI handling.....265
 load-balancing and queuing behavior.....269
 data service unit (DSU) device.....88
 data stream clocking.....80
 data terminal equipment *See* DTE
 data-link connection identifiers *See* DLCIs
 Database Information page.....17
 dc-0/0/0
 ISDN BRI interface.....214
 See also D-channel; ISDN BRI interfaces; ISDN
 PRI interfaces
 ISDN PRI interface.....215
 DCE (data communications equipment)
 serial connection process.....67
 serial device.....66
 DCE clocking mode.....68
 DDR *See* dial-on-demand routing backup, ISDN
 DE (discard eligibility) bits
 BECN bits.....85
 FECN bits.....85
 deactivate command.....30
 deactivating configuration statements or
 identifiers.....30
 default gateway, static routing.....387
 delay-sensitive packets, LFI handling.....265
 See also LFI

Delete button.....8
 delete command.....27
 Delete Configuration Below This Point option
 button.....13
 Delete link.....12
 deleting
 current rescue configuration (CLI configuration
 editor).....34
 current rescue configuration (J-Web).....21
 network interfaces.....134
 designated router, OSPF
 controlling election.....425
 description.....365
 destination prefix lengths.....93
 Deutsche Telekom UR-2 operating mode.....164, 167
 diagnosis
 BERT.....79
 channelized T1/E1 interfaces.....153
 displaying IS-IS-enabled interfaces.....435
 displaying IS-IS-enabled interfaces (detail).....436
 displaying static routes in the routing table.....393
 IS-IS adjacencies.....437
 IS-IS adjacencies (detail).....438
 IS-IS neighbors.....437
 IS-IS neighbors (detail).....438
 load balancing on the link services interface.....307
 packet encapsulation on link services
 interfaces.....306
 PPP magic numbers.....87
 verifying B-channels.....249
 verifying BGP configuration.....456
 verifying BGP groups.....455
 verifying BGP peer reachability.....457
 verifying BGP peers (neighbors).....454
 verifying D-channels.....251
 verifying dialer interfaces.....255
 verifying ISDN BRI interfaces.....247
 verifying ISDN call status.....253
 verifying ISDN PRI interfaces.....249
 verifying ISDN status.....246
 verifying link services CoS.....300
 verifying link services interface status.....296
 verifying OSPF host reachability.....430
 verifying OSPF neighbors.....428
 verifying OSPF routes.....429
 verifying OSPF-enabled interfaces.....427
 verifying PPPoA for ATM-over-ADSL
 configuration.....185
 verifying PPPoE interfaces.....208
 verifying PPPoE over ATM-over-ADSL
 configuration.....206, 207
 verifying PPPoE over ATM-over-SHDSL
 configuration.....206, 207
 verifying PPPoE sessions.....209
 verifying PPPoE statistics.....210
 verifying PPPoE version information.....209

verifying RIP host reachability	409
verifying RIP message exchange.....	408
verifying RIP-enabled interfaces.....	407
VoIP interface.....	342
dial backup	
configuring (configuration editor).....	228
configuring (Quick Configuration—ISDN BRI).....	222
interfaces to back up (configuration editor).....	229
interfaces to back up (Quick Configuration).....	222
selecting (Quick Configuration—ISDN BRI).....	220
dial-in, ISDN	
dialer interface (configuration editor).....	240
encapsulation matching.....	239
overview.....	238
rejecting incoming calls (configuration editor).....	243
screening incoming calls (configuration editor).....	242
voice not supported.....	238
dial-on-demand filter <i>See</i> dialer filter, ISDN	
dial-on-demand routing backup, ISDN	
dialer filter.....	229
<i>See also</i> dialer filter, ISDN	
dialer watch.....	231
<i>See also</i> dialer watch	
OSPF support.....	233
<i>See also</i> dialer watch	
dialer filter, ISDN	
applying to the dialer interface.....	231
configuring.....	230
overview.....	229
dialer interface, ISDN	
adding.....	226
bandwidth on demand (configuration editor).....	234
callback (configuration editor).....	240
dial-in (configuration editor).....	240
dialer filter.....	229
<i>See also</i> dialer filter, ISDN	
dialer watch <i>See</i> dialer watch	
disabling dial-out (configuration editor).....	243
encapsulation matching for dial-in or callback.....	239
limitations.....	215
multiple, ensuring different IPv4 subnBRI et addresses on.....	222
naming convention.....	215
rejecting incoming calls (configuration editor).....	243
restrictions.....	215
screening incoming calls (configuration editor).....	242
secondary (backup) connection.....	228
verifying.....	255
dialer interface, ISDN BRI (Quick Configuration).....	219
dialer options, ISDN	
for ISDN BRI service.....	219
for ISDN PRI service.....	150
dialer pools, ISDN	
for bandwidth on demand (configuration editor).....	238
for dialer watch (configuration editor).....	233
ISDN BRI physical interface (configuration editor).....	224
Quick Configuration.....	218
dialer watch	
adding a dialer watch interface (configuration editor).....	232
configuring (Quick Configuration—ISDN BRI).....	222
dialer pool (configuration editor).....	233
ISDN interface for (configuration editor).....	232
overview.....	231
selecting (Quick Configuration—ISDN BRI).....	220
watch list (configuration editor).....	232
watch list (Quick Configuration).....	222
digital subscriber line (DSL) <i>See</i> ATM-over-ADSL interfaces; ATM-over-SHDSL interfaces; DSLAM connection	
Discard All Changes option button.....	13
Discard button.....	13
Discard Changes Below This Point option button.....	13
discard eligibility bits <i>See</i> DE bits	
discard interface.....	100
discarding configuration changes.....	13
discovery packets, PPPoE.....	88, 194
Disk-on-Key configuration	
description.....	323
procedure.....	324
requirements.....	323
RESET CONFIG button caution.....	324
distance-vector routing protocols.....	360
<i>See also</i> RIP	
dl0.....	215
<i>See also</i> dialer interface, ISDN	
DLCIs (data-link connection identifiers)	
in MLFR FRF.16 bundles (configuration editor).....	287
overview.....	84
documentation set	
comments on.....	xxv
domains	
broadcast domains.....	56
collision domains.....	55
dotted decimal notation.....	92
downloading, configuration files (J-Web).....	20
drop-and-insert of time slots, on channelized ports	
clock source requirement.....	145
configuring.....	146
overview.....	141, 145
possible clock combinations.....	153

sample configuration.....	154
signaling channel requirement.....	145
DS0 time slots	
channelization.....	61
<i>See also</i> channelized E1 interfaces;	
channelized T1 interfaces	
drop-and-insert, on channelized T1/E1	
interfaces.....	145
DS1 interfaces <i>See</i> E1 interfaces; T1 interfaces	
DS1 ports <i>See</i> E1 ports; T1 ports	
DS1 signals	
E1 and T1.....	58
<i>See also</i> E1 interfaces; T1 interfaces	
multiplexing into DS2 signal.....	62
DS2 signals	
bit stuffing.....	62
frame format.....	62
DS3 interfaces <i>See</i> E3 interfaces; T3 interfaces	
DS3 ports <i>See</i> E3 ports; T3 ports	
DS3 signals	
DS3 C-bit parity frame format.....	64
M13 frame format.....	63
dsc interface.....	100
DSL <i>See</i> ATM-over-ADSL interfaces; ATM-over-SHDSL	
interfaces; DSLAM connection	
DSL access multiplexer <i>See</i> DSLAM connection	
DSLAM connection	
ATM-over-ADSL interface for.....	164
ATM-over-SHDSL interface for.....	174
PPPoE over ATM-over-ADSL topology.....	193
DSU (data service unit) device.....	88
DTE (data terminal equipment)	
default clock rate reduction.....	69
serial connection process.....	67
serial device	66
DTE clocking mode <i>See</i> internal clocking mode	
dying gasp message, SHDSL.....	189
dynamic CAC	
activation priority, description.....	319
BBL, description.....	319
CAC-BL requirement for WANs.....	319
configuring on WAN interfaces (configuration	
editor).....	330
overview.....	318
supported interfaces.....	318
verifying available bandwidth.....	341
dynamic call admission control <i>See</i> dynamic CAC	
dynamic routing.....	357

E

E1 interfaces	
AMI encoding.....	59
CRTP (configuration editor).....	289
data stream.....	58
encoding.....	58
framing.....	59
HDB3 encoding.....	59
loopback.....	60
multilink bundles (Quick Configuration).....	271
overview.....	58
<i>See also</i> E1 ports; channelized E1 interfaces	
Quick Configuration.....	105
signals.....	58
E1 ports	
CHAP.....	108
clocking.....	107
data inversion.....	108
encapsulation type.....	107
fractional, channel number.....	49
frame checksum.....	108
framing.....	108
logical interfaces.....	107
MTU.....	107
MTU default and maximum values.....	81
overview.....	58
<i>See also</i> E1 interfaces; channelized E1 ports	
Quick Configuration.....	105
time slots.....	108
E1/T1 media module <i>See</i> TIM510	
E3 interfaces	
bit stuffing.....	62
data stream.....	62
DS3 framing.....	63
multilink bundles (Quick Configuration).....	271
multiplexing on.....	63
overview.....	62
<i>See also</i> E3 ports	
Quick Configuration.....	109
E3 ports	
CHAP.....	110
clocking.....	110
encapsulation type.....	110
frame checksum.....	112
logical interfaces.....	110
MTU.....	110
MTU default and maximum values.....	81
overview.....	62
<i>See also</i> E3 interfaces	
Quick Configuration.....	109
EBGP (external BGP)	
description.....	372
sample network.....	447
edit command.....	23
Edit Configuration page.....	11
Edit Configuration Text page.....	15
Edit link.....	12
EGPs (exterior gateway protocols).....	355
EIA-232.....	70
EIA-422.....	71
EIA-449.....	71
EIA-530.....	70

Electronic Preinstallation Worksheet (EPW), for Avaya	
VoIP configuration.....	323
<i>See also</i> EPW configuration	
encapsulation overhead, PPP and MLPPP.....	306, 307
encapsulation type	
ATM-over-ADSL logical interfaces.....	162, 168
ATM-over-ADSL physical interfaces.....	163, 167
ATM-over-SHDSL logical interfaces	172, 178
ATM-over-SHDSL physical interfaces.....	173, 176
E1	107
E3.....	110
Frame Relay.....	83
HDLC.....	89
ISDN dial-in and callback, monitoring.....	239
overview.....	83
PPP.....	85
PPPoE.....	191
PPPoE for Ethernet.....	199
PPPoE, over ATM-over-ADSL.....	200
PPPoE, over ATM-over-SHDSL.....	200
PPPoE, overview.....	88
serial interfaces.....	129
T1.....	122
T3.....	126
verifying for LFI and load balancing.....	306
encoding	
AMI.....	59
B8ZS.....	59
channelized T1 (configuration editor).....	144
HDB3.....	59
EPW (Electronic Preinstallation Worksheet)	
configuration	
description.....	323
procedure.....	324
requirements.....	323
RESET CONFIG button caution.....	324
error notification, in the data link layer.....	52
ESF (extended superframe) framing.....	60
Ethernet interfaces.....	53, 112, 115
access control.....	54
broadcast domains.....	56
collision detection.....	54
collision domains.....	55
CSMA/CD.....	54
frame format.....	56
IS-IS, NET address.....	434
overview.....	53
Quick Configuration.....	112, 115
<i>See also</i> Fast Ethernet ports	
<i>See also</i> Fast Ethernet ports; Gigabit Ethernet ports	
<i>See also</i> Gigabit Ethernet ports	
Ethernet over ATM encapsulation.....	167
ATM-over-ADSL interfaces.....	162, 163
ATM-over-SHDSL interfaces.....	173, 176
for PPPoE.....	201
Ethernet over ATM LLC encapsulation	
ATM-over-ADSL interfaces.....	168
ATM-over-SHDSL interfaces.....	172, 178
Ethernet ports <i>See</i> Ethernet interfaces; Fast Ethernet ports; Gigabit Ethernet ports	
Ethernet switches	
configuring uPIMs as.....	343
ETSI TS 101 388 V1.3.1 operating mode.....	164, 167
EU-64 addresses.....	53
exit command	
to leave configuration mode.....	22
to navigate the configuration hierarchy.....	24
exit configuration-mode command.....	22
explicit clocking signal transmission.....	80
extended superframe (ESF) framing.....	60
exterior gateway protocols.....	355
external BGP <i>See</i> EBGp	
external paths, role in BGP route selection.....	373
F	
failover connection, ISDN.....	211
FAQ (frequently asked questions)	
Are LFI and load balancing working	
correctly?.....	304
What causes jitter and latency on multilink	
bundles?.....	303
What clock combinations are possible for	
channelized T1/E1 drop-and-insert?.....	153
Which CoS components apply on link services	
interface?.....	302
Why Are Packets Dropped on a PVC Between a	
J-series Router and Another Vendor?.....	310
Why is the VoIP interface unavailable?.....	342
Fast Ethernet ports	
ARP address.....	113
CHAP for PPPoA.....	179
CHAP for PPPoE.....	204
logical interfaces.....	113
MAC address.....	113
MTU.....	115
MTU default and maximum values.....	81
overview.....	53
PPPoE configuration.....	201
PPPoE encapsulation.....	199
PPPoE session on.....	193
Quick Configuration.....	112
static ARP entries (configuration editor).....	133
FCS (frame check sequence)	
checksums.....	81
CRCs.....	80
overview.....	80
two-dimensional parity.....	81
fe-0/0/0, management interface.....	101
<i>See also</i> Fast Ethernet ports	
FEAC C-bit condition indicators.....	65

FECN (forward-explicit congestion notification)
 bits.....85
 file management
 configuration files (CLI configuration editor).....35
 configuration files (J-Web).....16
 flow control
 data link layer.....52
 Gigabit Ethernet.....118
 font conventions.....xxi
 forward-explicit congestion notification (FECN)
 bits.....85
 forwarding classes, defining.....277
 forwarding table
 controlling OSPF routes in.....422
 controlling static routes in.....384, 391
 description.....356
 MED to determine routes in.....376
 four-wire mode (1 port), SHDSL *See* ATM-over-SHDSL
 interfaces
 FPC (PIM slot on a Services Router) *See* PIMs
 fragmentation, verifying on the link services
 interface.....305
 frame check sequence *See* FCS
 Frame Relay encapsulation
 congestion control.....85
 DLCIs.....84
 overview.....83
 PVCs.....84
 SVCs.....84
 virtual circuits.....84
 Frame Relay network, typical.....84
 frames
 DS2 M-frame format.....62
 DS3 C-bit parity frame format.....64
 DS3 M13 frame format.....63
 Ethernet frame format.....56
 sequencing, data link layer.....52
 framing
 channelized E1 (configuration editor).....144
 channelized T1 (configuration editor).....144
 E1.....108
 T1.....123
 T3.....127
 frequently asked questions *See* FAQ
 FRF.15 and FRF.16 *See* MLFR FRF.15; MLFR FRF.16
 full mesh requirement
 description.....373
 fulfilling with confederations.....381
 fulfilling with route reflectors.....378
 sample network.....447
 fxp interfaces (not supported).....97

G

G.992.1 Deutsche Telekom UR-2 operating
 mode.....164, 167

G.992.1 Non-UR-2 operating mode.....164, 167
 G.SHDSL PIMs.....74
 Annex A or Annex B modes.....169
 configuring.....169
 default mode.....175
 standard supported.....74
 See also ATM-over-SHDSL interfaces
 Gateway Module *See* TGM550
 ge-0/0/0, management interface.....101
 See also Gigabit Ethernet ports
 Gigabit Ethernet ports
 (copper) manual speed and link mode
 configuration.....120
 802.3ad.....119
 aggregated Ethernet.....119
 ARP address.....117
 autonegotiation.....119
 CHAP for PPPoA.....179
 CHAP for PPPoE.....204
 dynamic CAC for voice packets (configuration
 editor).....331
 See also Avaya VoIP
 flow control.....118
 logical interfaces.....117
 MAC address.....117
 MAC address learning.....120
 MTU.....118
 MTU default and maximum values.....81
 overview.....53
 PPPoE configuration.....201
 PPPoE encapsulation.....199
 PPPoE session on.....193
 Quick Configuration.....115
 source filtering.....118
 source filtering, for MAC addresses.....119
 static ARP entries (configuration editor).....133
 TPIDs.....119
 Gigabit Ethernet uPIMs
 as switches.....343
 global unicast IPv6 addresses.....95
 glossary
 Avaya VoIP.....311
 channelized T1/E1/ISDN PRI.....139
 configuration.....3
 DSL.....157
 interfaces.....42
 ISDN.....211
 link services.....261
 ports.....42
 PPPoE.....191
 routing protocols.....349
 gr-0/0/0 interface.....97
 gre interface.....97
 grounding cable required for TGM550.....322

groups		incoming metric (RIP)	
BGP <i>See</i> BGP groups		description.....	396
OSPF areas.....	417	modifying.....	402
RIP routers.....	398	inet protocol family.....	91
		MTU value for PPPoE.....	203
		inet6 protocol family.....	91
		MTU value for PPPoE.....	203
		insert command.....	28
		inserting configuration identifiers.....	28
		Integrated Services Digital Network <i>See</i> ISDN	
		interface naming conventions.....	47
		interfaces	
		ATM-over-ADSL interfaces.....	72
		ATM-over-SHDSL interfaces.....	74
		Avaya VoIP.....	315
		<i>See also</i> Avaya VoIP modules	
		channelized T1/E1/ISDN PRI interfaces.....	61
		clocking.....	79
		data link layer.....	52
		E1 interfaces.....	57
		E3 interfaces.....	62
		Ethernet interfaces.....	53
		FCS.....	80
		G.SHDSL interfaces.....	74
		IPv4 addressing.....	91
		IPv6 addressing.....	94
		ISDN interfaces.....	75
		logical properties.....	90
		MTU values.....	81
		overview.....	41
		<i>See also</i> ATM-over-ADSL interfaces;	
		ATM-over-SHDSL interfaces; channelized	
		interfaces; ISDN interfaces; link services	
		interface; loopback interface; management	
		interfaces; network interfaces; ports;	
		special interfaces; VoIP interface	
		physical encapsulation.....	83
		<i>See also</i> encapsulation type	
		physical properties.....	78
		protocol families.....	91
		Quick Configuration.....	104
		serial interfaces.....	66
		special interfaces.....	97
		supported for dynamic CAC, for Avaya VoIP.....	318
		<i>See also</i> dynamic CAC	
		T1 interfaces.....	57
		T3 interfaces.....	61
		VLANs.....	96
		VoIP.....	315
		<i>See also</i> Avaya VoIP modules	
		Interfaces page.....	104
		for E1.....	106
		for E3.....	109
		for Fast Ethernet.....	113
		for Gigabit Ethernet.....	116
		for serial interfaces.....	128
H			
HDB3 encoding.....	59		
HDLC (High-Level Data Link Control)			
encapsulation.....	89		
HDLC operational modes.....	90		
HDLC stations.....	89		
hello PDUs.....	369		
hierarchy, configuration.....	23		
high-density bipolar 3 code (HDB3) encoding.....	59		
High-Level Data Link Control <i>See</i> HDLC			
history <i>See</i> configuration history			
hold time, to maintain a session.....	372		
hop count, maximizing.....	361		
<i>See also</i> RIP			
host reachability			
verifying for an OSPF network.....	430		
verifying for RIP network hosts.....	409		
hostname			
for PPPoA CHAP.....	180		
for PPPoE CHAP (configuration editor).....	205		
for PPPoE CHAP (Quick Configuration).....	197		
IS-IS identifier-to-hostname mapping.....	432		
hovering the J-Web mouse pointer.....	7		
how to use this guide.....	xx		
I			
IBGP (internal BGP)			
description.....	372		
full mesh (configuration editor).....	447		
full mesh requirement.....	442		
sample network.....	447		
sample route reflector.....	449		
identifier link.....	12		
identifiers, configuration			
adding or modifying.....	25		
deactivating.....	30		
deleting.....	27		
inserting.....	28		
renaming.....	28		
IG550 Integrated Gateway <i>See</i> Avaya IG550 Integrated			
Gateway; Avaya VoIP modules			
IGP plus MED, BGP option.....	378		
IGP route metric, role in BGP route selection.....	374		
IGPs (interior gateway protocols).....	355		
incoming calls			
rejecting.....	243		
screening.....	242		

- for T1.....121
 - for T3 (DS3).....125
- interior gateway protocols.....355
- internal BGP *See* IBGP
- internal clocking mode.....68
- Internet routing, with BGP.....441
- invalid configuration, replacing
 - with J-Web.....21
 - with the CLI.....33
- inverting the transmit clock.....130
- IP addresses.....91
 - as IS-IS system identifiers.....432
 - Avaya VoIP module, guidelines for.....320
 - TGM550, guidelines for.....320
 - See also* addresses; IPv4 addressing; IPv6 addressing
- ip telnet port command.....336
- ip telnet-client command.....336
- ip-0/0/0 interface.....98
- ip-ip interface.....98
- IPv4 addressing
 - assigning for PPPoE (configuration editor).....203
 - assigning for PPPoE (Quick Configuration).....197
 - classful addressing.....92
 - dotted decimal notation.....92
 - MAC-48 address format.....53
 - overview.....91
 - subnets.....93
 - VLSMs.....93
- IPv4 MTU value, PPPoE.....203
- IPv6 addressing
 - address format.....94
 - address scope.....95
 - address structure.....95
 - address types.....95
 - assigning for PPPoE.....203
 - overview.....94
- IPv6 MTU value, PPPoE.....203
- IPv6 support.....349
- IS-IS (Intermediate System-to-Intermediate System)
 - adjacency establishment with hello PDUs.....369
 - areas.....368
 - ASs.....368
 - CSNPs.....370
 - enabling on router interfaces.....433
 - hello PDUs.....369
 - LSPs.....370
 - NETs.....369
 - See also* NETs
 - NSAP addresses.....431
 - overview.....368, 431
 - path selection.....369
 - preparation.....432
 - PSNPs.....370
 - system identifiers.....369
 - See also* system identifiers
- verifying adjacencies.....437
 - verifying adjacencies (detail).....438
 - verifying interface configuration.....435
 - verifying interface configuration (detail).....436
 - verifying neighbors.....437
 - verifying neighbors (detail).....438
- ISDN BRI Dialer Logical Interface page.....221
- ISDN BRI interfaces
 - adding an interface.....224
 - B-channel interface.....214
 - bandwidth on demand (configuration editor).....238
 - call setup.....77
 - callback *See* callback
 - calling number.....218, 225
 - connection initialization.....77
 - D-channel interface.....214
 - dial backup.....220, 228 *See* dial backup
 - dial-in *See* dial-in
 - dial-on-demand routing backup, with OSPF.....233
 - dialer filter.....229
 - dialer interface *See* dialer interface, ISDN
 - dialer watch *See* dialer watch
 - dialer watch (configuration editor).....232
 - disabling dial-out (configuration editor).....243
 - disabling ISDN signaling (configuration editor).....244
 - ISDN channels.....75
 - MTU default and maximum values.....81
 - naming conventions.....214
 - NT1 devices.....76
 - overview.....75
 - See also* ISDN connections
 - PIMs supported.....214
 - Q.931 timer.....219, 225
 - Quick Configuration.....216
 - requirements.....215
 - S/T interfaces.....76, 214
 - screening incoming calls.....242
 - session establishment.....77
 - SPID.....218, 225
 - static TEI.....219, 225
 - switch types.....218, 225
 - TEI option.....219, 226
 - typical network.....75
 - U interface.....76, 214
 - verifying B-channels.....249
 - verifying call status.....253
 - verifying D-channels.....251
 - verifying ISDN interfaces.....247
 - verifying ISDN status.....246
- ISDN BRI Physical Interface page.....216
- ISDN connections
 - adding an ISDN BRI interface.....224
 - adding an ISDN PRI interface.....147
 - bandwidth on demand.....234

callback <i>See</i> callback	
calling number.....	218, 225
configuring.....	211
dial backup <i>See</i> dial backup	
dial-in <i>See</i> dial-in	
dial-on-demand routing backup, with OSPF.....	233
dialer filter <i>See</i> dialer filter	
dialer interface <i>See</i> dialer interface, ISDN	
dialer watch <i>See</i> dialer watch	
disabling dial-out (configuration editor).....	243
disabling ISDN signaling (configuration editor).....	244
interface naming conventions.....	214
ISDN interface types.....	214
overview.....	214
<i>See also</i> dialer interfaces; ISDN BRI interfaces; ISDN PRI interfaces	
Q.931 timer.....	219, 225
Quick Configuration (ISDN BRI).....	216
requirements.....	215
SPID.....	218, 225
static TEI.....	219, 225
switch types.....	141, 218, 225
TEI option.....	219, 226
verifying B-channels.....	249
verifying call status.....	253
verifying D-channels.....	251
verifying dialer interfaces.....	255
verifying ISDN BRI interfaces.....	247
verifying ISDN PRI interfaces.....	249
verifying ISDN status.....	246
ISDN PRI interfaces	
adding.....	147
B-channel allocation order.....	149
B-channel interface.....	215
bandwidth on demand.....	238
callback <i>See</i> callback	
channelized interface.....	214
D-channel interface.....	215
dial backup.....	228
dial-in <i>See</i> dial-in	
dialer filter.....	229
dialer interface <i>See</i> dialer interface, ISDN	
dialer options.....	150
dialer watch.....	232
disabling dial-out.....	243
disabling ISDN signaling.....	244
overview.....	140
PIM supported.....	214
Q.931 timers.....	150
screening incoming calls.....	242
supported switch types.....	141
transmission.....	141
verifying B-channels.....	249
verifying call status.....	253
verifying configuration.....	153
verifying D-channels.....	251
verifying ISDN interfaces.....	249
verifying ISDN status.....	246
ISO network addresses, for IS-IS routers.....	431
ISO protocol family.....	91
ITU Annex B non-UR-2 operating mode.....	164
ITU Annex B UR-2 operating mode.....	164
ITU Annex B non-UR-2 operating mode.....	167
ITU Annex B UR-2 operating mode.....	167
ITU DMT bis operating mode.....	164, 167
ITU DMT operating mode.....	164, 167
ITU G.992.1 operating mode.....	164, 167
ITU G.992.3 operating mode.....	164
ITU G.992.5 operating mode.....	164, 167
J	
J-series	
Avaya VoIP connectivity.....	311
BGP routing.....	441
channelized T1/E1/ISDN PRI interfaces.....	139
configuration tools.....	3
DSL.....	157
interfaces overview.....	41
IS-IS protocol.....	431
ISDN connections.....	139, 211
link services interface.....	261
link services interface, implementation	
exceptions.....	264
network interfaces.....	103
OSPF routing.....	411
PPPoE.....	191
release notes, URL.....	xix
RIP routing.....	395
routing protocols overview.....	349
static routing.....	383
J-Web configuration editor	
ATM-over-ADSL interfaces.....	164
ATM-over-SHDSDL interfaces.....	174
Avaya VoIP.....	327
BGP.....	444
channelized E1 interfaces.....	142
channelized T1 interfaces.....	142
CHAP on ATM-over-ADSL interfaces.....	179
CHAP on ATM-over-SHDSDL interfaces.....	179
clickable configuration, committing.....	14
clickable configuration, discarding changes.....	13
clickable configuration, editing.....	10
committing a text file, with caution.....	14
configuration text, viewing.....	9
CRTP.....	289
editing a text file, with caution.....	14
IS-IS.....	433
ISDN connections.....	223
LFI.....	273
managing files.....	16

- MLPPP bundles.....273
 - network interfaces.....131
 - network interfaces, adding.....131
 - network interfaces, deleting.....134
 - OSPF.....414
 - PPPoE.....198
 - PPPoE over ATM-over-ADSL.....198
 - PPPoE over ATM-over-SHDLSL.....198
 - RIP.....398
 - static routes.....387
 - uploading a file.....15
 - VoIP.....327
 - J-Web interface.....5
 - comparing configuration differences.....19
 - configuration history.....16
 - See also* configuration history
 - configuration options.....5
 - mouse-over.....7
 - See also* J-Web configuration editor
 - J2300 routers
 - ADSL support.....46, 73
 - built-in Ethernet interfaces.....101
 - built-in ISDN BRI S/T interface.....214
 - built-in ISDN BRI U interface.....214
 - DSL support.....45
 - ISDN support.....45
 - MTU values.....81
 - PIM number (0).....48
 - J4300 routers
 - built-in Ethernet interfaces.....101
 - MTU values.....81
 - PIM number.....48
 - J4350 routers
 - Avaya VoIP connectivity.....311
 - built-in Ethernet interfaces.....101
 - Gigabit Ethernet support.....45
 - manual copper Gigabit Ethernet speed and link
 - mode configuration.....120
 - MTU values.....82
 - PIM number.....48
 - T3 (DS3) and E3 support.....45
 - J6300 routers
 - built-in Ethernet interfaces.....101
 - MTU values.....81
 - PIM number.....48
 - T3 (DS3) and E3 support.....45
 - J6350 routers
 - Avaya VoIP connectivity.....311
 - built-in Ethernet interfaces.....101
 - Gigabit Ethernet support.....45
 - manual copper Gigabit Ethernet speed and link
 - mode configuration.....120
 - MTU values.....82
 - PIM number.....48
 - T3 (DS3) and E3 support.....45
 - jitter, removing on multilink bundles.....303
 - JTAC (Juniper Networks Technical Assistance Center)
 - See* technical support
 - Juniper Networks Technical Assistance Center *See*
 - technical support
 - JUNOS Internet software
 - Avaya VoIP configurability with.....316
 - Avaya VoIP connectivity.....311
 - configuration.....3
 - See also* configuration
 - ISDN connections.....211
 - release notes, URL.....xix
 - TGM550 firmware compatibility with.....320
- K**
- keepalive messages, for session hold time.....372
- L**
- LANs
 - bridges on LAN segments.....56
 - collision domains55
 - repeaters on LAN segments.....55
 - topology.....96
 - latency, reducing on multilink bundles.....303
 - lazy quantifiers, and replace command.....26
 - LCP (Link Control Protocol), connection process.....85
 - Level 1 areas, IS-IS.....368
 - Level 2 areas, IS-IS.....368
 - LFI (link fragmentation and interleaving)
 - enabling (configuration editor).....276
 - load-balancing behavior.....268
 - overview.....265
 - See also* link services interface
 - queuing behavior for data vs. voice packets.....268
 - queuing on constituent links.....267
 - See also* queuing with LFI
 - with CoS components.....269
 - line buildout
 - T1.....124
 - T3.....127
 - line speed
 - ATM-over-SHDLSL interfaces.....173
 - serial interfaces.....131
 - line timing.....68
 - link fragmentation and interleaving *See* LFI
 - link hold time, channelized ports.....143
 - link services.....101
 - See also* link services interface; ls-0/0/0
 - link services interface
 - applying CoS components on constituent
 - links.....302
 - channels, with MLFR FRF.16 (configuration
 - editor).....287
 - classifiers and forwarding classes (configuration
 - editor).....277

configuring.....	261	rollback (J-Web)	21
CoS components.....	269	rollback command.....	32
<i>See also</i> CoS components for link services		uploading (J-Web).....	15
CRTP (configuration editor).....	289	without specifying full hierarchy.....	36
displaying CoS configurations.....	294	local preference	
FAQ.....	301	description.....	374
fragmentation, troubleshooting.....	305	high value preferred.....	375
J-series implementation exceptions.....	264	role in BGP route selection.....	373
LFI <i>See</i> LFI		locked configuration.....	23
load balancing, troubleshooting.....	307	logical interfaces	
MLFR bundles (Quick Configuration).....	271	adding (configuration editor).....	133
MLFR FRF.15 bundles (configuration editor).....	284	ATM-over-ADSL (configuration editor).....	167
MLFR FRF.16 bundles (configuration editor).....	287	ATM-over-ADSL (Quick Configuration).....	161
MLPPP bundles (Quick Configuration).....	271	ATM-over-SHDSL.....	177
MLPPP header overhead.....	306	ATM-over-SHDSL (Quick Configuration).....	172
multilink bundles <i>See</i> multilink bundles		E1.....	107
overview.....	262	E3.....	110
<i>See also</i> ls-0/0/0		Fast Ethernet.....	113
packet encapsulation, troubleshooting.....	306	Gigabit Ethernet.....	117
PPP header overhead.....	306	serial	129
preventing dropped packets on PVCs.....	310	T1.....	122
Quick Configuration.....	271	T3.....	126
reducing jitter and latency on multilink		logical units	
bundles.....	303	adding (configuration editor)	133
requirements.....	270	ATM-over-ADSL interface (Quick	
sample CoS configuration.....	294	Configuration).....	161
scheduler maps (configuration editor).....	279	ATM-over-SHDSL interface (Quick	
services on.....	263	Configuration).....	172
shaping rates, applying (configuration		E1 interface.....	107
editor).....	283	E3 interface.....	110
troubleshooting.....	301	Fast Ethernet interface.....	113
troubleshooting LFI and load balancing.....	304	Gigabit Ethernet interface.....	117
verifying.....	291	number in interface name.....	49
verifying CoS configuration.....	300	pp0 interface.....	201
verifying status.....	296	PPPoE encapsulation.....	199
link states, verifying.....	136	PPPoE over ATM-over-ADSL encapsulation.....	200
link-local unicast IPv6 addresses.....	95	PPPoE over ATM-over-SHDSL encapsulation.....	200
link-state advertisements <i>See</i> LSAs		serial interface.....	129
link-state PDUs <i>See</i> LSPs		T1 interface.....	122
lo0 interface functions.....	100	T3 interface.....	126
<i>See also</i> loopback interface		long buildout <i>See</i> line buildout	
lo0.16385, internal loopback address.....	98	loop clocking mode.....	68
load balancing on link services interfaces		loopback address, internal, lo0.16385.....	98
description.....	268	loopback interface	
FAQ.....	304	functions.....	100
troubleshooting.....	304	NET on for IS-IS.....	433
verifying.....	307	loopback signals, E1 and T1	60
load command.....	35	loopback testing, SHDSL.....	174
load merge command.....	35	ls-0/0/0	
load override command.....	35	configuring.....	261
load patch command.....	35	<i>See also</i> link services interface	
load replace command.....	36	interface description.....	98
loading a configuration file		LSAs (link-state advertisements)	
CLI configuration editor.....	35	description.....	364
downloading (J-Web).....	20	three-way handshake.....	364
examples.....	36	lsi interface.....	98

LSPs (link-state PDUs)	
CSNPs.....	370
overview.....	370
PSNPs.....	370
lt-0/0/0 interface.....	98

M

M13 frame format.....	63
MAC (media access control) addresses	
as IS-IS system identifiers.....	432
associating with IP addresses on Ethernet subnets.....	133
EUI-64 addresses.....	53
for static ARP on Fast Ethernet subnets.....	114
<i>See also</i> static ARP entries	
for static ARP on Gigabit Ethernet subnets.....	117
<i>See also</i> static ARP entries	
in static ARP entries (configuration editor).....	133
learning, on Gigabit Ethernet ports.....	120
MAC-48 address format.....	53
overview.....	53
physical addressing.....	52
source filtering on Gigabit Ethernet ports.....	119
MAC-48 addresses.....	53
magic numbers, PPP.....	87
management interfaces, overview.....	101
managing files <i>See</i> file management	
manuals	
Avaya VoIP.....	317
comments on.....	xxv
maximum hop count, RIP.....	361
maximum transmission unit <i>See</i> MTU	
MED (multiple exit discriminator)	
always compare option.....	377
Cisco non-deterministic option.....	378
default use.....	376
description.....	376
path selection options.....	377
plus IGP option.....	378
role in BGP route selection.....	373
media access control <i>See</i> MAC addresses	
Media Gateway Controller <i>See</i> Avaya Media Gateway Controller; MGC list	
media types supported.....	46
memory stick, USB, for Avaya VoIP	
configuration.....	323
merging a configuration file.....	35
example.....	37
metrics <i>See</i> MED; path cost metrics	
MGC <i>See</i> Avaya Media Gateway Controller; MGC list	
MGC list	
clearing.....	330
configuring.....	329
overview.....	317
<i>See also</i> Avaya VoIP	

Quick Configuration.....	325, 326
verifying.....	341
MLFR (Multilink Frame Relay)	
multilink bundles (Quick Configuration).....	271
overview.....	101
<i>See also</i> link services interface; multilink bundles	
MLFR bundles <i>See</i> MLFR; multilink bundles	
MLFR FRF.15	
multilink bundles (configuration editor).....	284
overview.....	102
MLFR FRF.16	
multilink bundles (configuration editor).....	287
overview.....	102
mlfr-end-to-end protocol family.....	91
mlfr-uni-nni protocol family.....	91
MLPPP (Multilink Point-to-Point Protocol)	
multilink bundles (configuration editor).....	274
multilink bundles (Quick Configuration).....	271
overview.....	101
<i>See also</i> link services interface; multilink bundles	
queuing behavior, with CRTP.....	269
queuing behavior, with LFI.....	269
sample topology.....	274
MLPPP bundles <i>See</i> MLPPP; multilink bundles	
MLPPP encapsulation, on the link services interface.....	306
mlppp protocol family.....	91
mouse-over, for J-Web selection.....	7
MPLS protocol family.....	91
MTU value for PPPoE.....	203
mt-0/0/0 interface.....	98
MTU (maximum transmission unit)	
default values for all interfaces.....	81
E1.....	107
E3.....	110
Fast Ethernet.....	115
Gigabit Ethernet.....	118
maximum values for all interfaces.....	81
serial.....	129
T1.....	122
T3.....	126
mtun interface.....	99
multiarea network, OSPF.....	417
multicast IPv6 addresses.....	95
multilink bundles	
buffer size for Q0.....	270
classifiers and forwarding classes (configuration editor).....	277
displaying configurations.....	292
LFI (configuration editor).....	276
MLFR FRF.15 (configuration editor).....	284
MLFR FRF.16 (configuration editor).....	287
overview.....	264
preventing dropped packets.....	310

queuing, on Q0 of constituent links.....	268
queuing, on Q2 of constituent links.....	268
Quick Configuration options.....	272
reducing latency.....	303
removing jitter.....	303
sample configuration.....	292
sample topology.....	274
scheduler maps (configuration editor).....	279
scheduling priority.....	270
shaping rate.....	269
shaping rates (configuration editor).....	283
Multilink Frame Relay <i>See</i> MLFR	
Multilink Frame Relay end-to-end <i>See</i> MLFR FRF.15	
Multilink Frame Relay Forum <i>See</i> MLFR FRF.15; MLFR FRF.16	
Multilink Point-to-Point Protocol <i>See</i> MLPPP	
multilink services	
configuring.....	261
overview.....	101
<i>See also</i> CRTP; link services interface; MLFR; MLPPP	
multiple exit discriminator <i>See</i> MED	
N	
n-selectors, in IS-IS NET addresses.....	432
names, of network interfaces.....	48
NCPs (Network Control Protocols).....	87
neighbors <i>See</i> adjacencies, IS-IS; BGP peers; OSPF neighbors; RIP neighbors	
NETs (network entity titles)	
n-selectors.....	432
on an Ethernet interface.....	434
on the loopback interface.....	433
parts.....	369
system identifier.....	369
Network Control Protocols (NCPs).....	87
network entity titles <i>See</i> NETs	
network interfaces	
adding.....	131
ATM-over-ADSL configuration.....	164
ATM-over-ADSL interfaces.....	72
ATM-over-SHDSL configuration.....	174
ATM-over-SHDSL interfaces.....	74
channelized E1 configuration.....	142
channelized T1 configuration.....	142
channelized T1/E1/ISDN PRI interfaces.....	61
clocking.....	79
deleting.....	134
DS3 configuration.....	124
E1 configuration.....	105
E1 interfaces.....	57
E3 configuration.....	109
E3 interfaces.....	62
enabling RIP on.....	398
Ethernet interfaces.....	53
Fast Ethernet configuration.....	112
FCS.....	80
G.SHDSL interfaces.....	74
Gigabit Ethernet configuration.....	115
IPv4 addressing.....	91
IPv6 addressing.....	94
ISDN interfaces.....	75
link services interface.....	261
logical properties.....	90
media types.....	46
MTU values.....	81
names.....	48
naming conventions.....	47
output, understanding.....	49
physical encapsulation.....	83
<i>See also</i> encapsulation type	
physical properties.....	78
preparation.....	103, 142
protocol families.....	91
Quick Configuration.....	104
sample name.....	49
serial configuration.....	127
serial interfaces.....	66
supported.....	46
T1 configuration.....	120
T1 interfaces.....	57
T3 configuration.....	124
T3 interfaces.....	61
verifying ATM-over-ADSL properties.....	181
verifying ATM-over-SHDSL configuration.....	186
verifying channelized interfaces.....	151
verifying clear-channel interfaces.....	152
verifying ISDN PRI configuration.....	153
verifying link states.....	136
verifying properties.....	137
verifying properties of uPIM switch ports.....	346
verifying RIP message exchange.....	408
verifying RIP on.....	407
VLANs.....	96
network service access point (NSAP) addresses for IS-IS routers.....	431
networks	
Avaya VoIP.....	315
description.....	354
designated router <i>See</i> designated router, OSPF	
IPv4 subnets.....	93
path cost metrics <i>See</i> path cost metrics	
PPPoE session on an ATM-over-ADSL loop.....	194
PPPoE session on an Ethernet loop.....	193
sample BGP AS path.....	376
sample BGP confederation.....	451
sample BGP confederations.....	381
sample BGP external and internal links.....	447
sample BGP local preference use.....	375
sample BGP MED use.....	377
sample BGP peer network.....	445

sample BGP peer session.....371

sample BGP route reflector (one cluster).....379, 449

sample BGP route reflectors (cluster of clusters).....380

sample BGP route reflectors (multiple clusters).....380

sample distance-vector routing.....360

sample LFI and multilink bundle topology.....274

sample multiarea OSPF routing.....366

sample multilink bundle and LFI topology.....274

sample OSPF backbone area.....367

sample OSPF multiarea network.....417

sample OSPF network with stubs and NSSAs.....367

sample OSPF single-area network.....416

sample OSPF stub areas and NSSAs.....421

sample OSPF topology.....429

sample poison reverse routing.....363

sample RIP network with incoming metric.....401

sample RIP network with outgoing metric.....403

sample RIP topology.....399

sample route advertisement.....358

sample route aggregation.....359

sample routing topology.....356

sample split horizon routing.....362

sample static route, preferred path.....389

sample stub network for static routes.....388

sample unidirectional routing.....363

static routing.....357

VoIP.....315

next hop

 address for static routes.....387

 defining for static routes.....389

 qualified, defining for static routes.....390

 qualified, for static routes.....384

 role in BGP route selection.....374

no ip telnet command.....337

no ip telnet-client command.....337

non-LFI packets *See* data packets

non-UR-2 operating mode.....164, 167

Normal Response Mode, HDLC.....90

not-so-stubby areas *See* NSSAs

notice icons.....xxi

NRM, HDLC.....90

NSAP (network service access point) addresses for IS-IS routers.....431

NSSAs (not-so-stubby areas)

 area ID (configuration editor).....419

 area ID (Quick Configuration).....414

 area type (Quick Configuration).....414

 creating (configuration editor).....420

 description.....367

 example.....368

 sample topology.....421

NT1 devices.....76

O

OK button

 J-Web configuration editor.....13

 Quick Configuration.....9

Open Shortest Path First protocol *See* OSPF

operational mode, entering during configuration.....34

option buttons

 Delete Configuration Below This Point.....13

 Discard All Changes.....13

 Discard Changes Below This Point.....13

origin, of BGP route.....376

OSPF (Open Shortest Path First)

 area border routers *See* area border routers

 area type (Quick Configuration).....414

 areas.....365, 412

See also area border routers; backbone area; NSSAs; stub areas

 authenticating exchanges (OSPFv2 only).....424

 backbone area *See* backbone area

 controlling designated router election.....425

 controlling route cost.....423

 designated router *See* designated router, OSPF

 designating OSPF interfaces (configuration editor).....417, 419

 designating OSPF interfaces (Quick Configuration).....414

 dial-on-demand routing backup support,

 ISDN.....233

 enabling (Quick Configuration).....414

 enabling, description.....411

 ensuring efficient operation.....422

 ISDN dial-on-demand routing backup support.....233

 LSAs.....364

 multiarea network (configuration editor).....417

 NSSAs *See* NSSAs

 overview.....364, 411

 path cost metrics *See* path cost metrics

 Quick Configuration.....413

 requirements.....412

 route preferences.....422

 router ID (configuration editor).....415

 router ID (Quick Configuration).....414

 sample multiarea network.....417

 sample network topology.....429

 sample NSSAs.....421

 sample single-area network.....416

 sample stub areas.....421

 single-area network (configuration editor).....415

 stub areas *See* stub areas

 supported versions.....364

 three-way handshake.....364

 tuning an OSPF network.....422

 verifying host reachability.....430

 verifying neighbors.....428

verifying RIP-enabled interfaces.....	427
verifying routes.....	429
OSPF interfaces	
enabling.....	414
enabling (configuration editor).....	417, 419
enabling, for area border routers.....	420
verifying.....	427
OSPF neighbors, verifying.....	428
OSPF page.....	413
field summary.....	414
out-of-band management interfaces.....	101
outgoing metric (RIP)	
description.....	396
modifying.....	404
overriding a configuration file.....	35
example.....	36
P	
packet encapsulation	
overview.....	83
<i>See also</i> encapsulation type	
troubleshooting on the link services	
interface.....	304
verifying on the link services interface.....	306
packet fragmentation	
troubleshooting on the link services	
interface.....	304
verifying on the link services interface.....	305
packets	
PPPoE discovery.....	88, 194
RIP, description.....	361
PADI packets.....	88
PADO packets.....	88
PADR packets.....	89
PADS packets.....	89
PADT packets.....	89
parentheses, in syntax descriptions.....	xxii
partial sequence number PDU (PSNP).....	370
passive routes, rejection, in static routing.....	385
password	
for OSPFv2 authentication.....	425
for RIPv2 authentication.....	404
for TGM550 access.....	334
patching a configuration file.....	35
path cost metrics	
for BGP, description.....	376
<i>See also</i> MED	
for OSPF routes, description.....	365, 412
for OSPF routes, modifying.....	423
for RIP routes, description.....	395
for RIP routes, modifying.....	401
path selection, IS-IS.....	369
path-vector protocol <i>See</i> BGP	
pd-0/0/0 interface.....	99
PDU (protocol data units)	
CSNPs.....	370
hello PDUs.....	369
LSPs.....	370
overview.....	369
PSNPs.....	370
pe-0/0/0 interface.....	99
peering sessions <i>See</i> BGP peers; BGP sessions	
per-unit scheduler, channelized ports.....	143
permanent routes, adding.....	383
permanent virtual circuits <i>See</i> PVCs	
Physical Interface Module <i>See</i> PIMs	
physical interface properties	
BERT.....	79
encapsulation.....	83
FCS.....	80
interface clocking.....	79
key properties.....	78
MTU values.....	81
PIC (PIM on a Services Router) <i>See</i> PIMs	
pimd interface.....	99
pime interface.....	99
PIMs (Physical Interface Modules)	
abbreviations.....	50
G.SHDSL.....	169
<i>See also</i> G.SHDSL PIMs	
initial configuration of interfaces.....	131
names.....	50
output about, understanding.....	49
PIM number, always 0.....	48
PIM slot number.....	48
Ping Host page, output for BGP.....	457
ping, verifying link states.....	136
plesiochronous networks.....	79
Point-to-Point Protocol <i>See</i> PPP	
Point-to-Point Protocol over ATM <i>See</i> PPPoA	
Point-to-Point Protocol over Ethernet <i>See</i> PPPoE	
poison reverse technique.....	362
polarity, signal.....	68
policy <i>See</i> routing policies	
ports	
Avaya VoIP.....	315
<i>See also</i> Avaya VoIP modules	
DS1 <i>See</i> E1 ports; T1 ports	
DS3 <i>See</i> E3 ports; T3 ports	
E1 <i>See</i> E1 ports	
E3 <i>See</i> E3 ports	
Fast Ethernet <i>See</i> Fast Ethernet ports	
Gigabit Ethernet <i>See</i> Gigabit Ethernet ports	
interfaces overview.....	41
<i>See also</i> ATM-over-ADSL interfaces;	
ATM-over-SHDSL interfaces; ISDN	
interfaces; link services interface; loopback	
interface; management interfaces; network	
interfaces; special interfaces; VoIP interface	
number in interface name.....	49

- serial *See* serial ports
- T1 *See* T1 ports
- T3 *See* T3 ports
- verifying status of uPIM ports in switching mode.....346
- VoIP.....315
 - See also* Avaya VoIP modules
- pp0
 - creating.....201
 - enabling CHAP.....204
 - information about.....208
 - interface description.....99
 - logical Ethernet interface on (configuration editor).....202
 - logical Ethernet interface on (Quick Configuration).....197
- PPP (Point-to-Point Protocol) *See* MLPPP; PPP encapsulation; PPPoA; PPPoE
- PPP encapsulation
 - CHAP authentication.....86
 - CSU/DSU devices.....88
 - LCP connection process.....85
 - magic numbers.....87
 - NCPs.....87
 - on the link services interface.....306
 - overview.....85
- PPP over ATM *See* PPPoA
- PPP over ATM-over-ADSL *See* PPPoA
- PPP over ATM-over-SHDSL *See* PPPoA
- PPP over Ethernet *See* PPPoE
- PPPoA (Point-to-Point Protocol over ATM)
 - CHAP.....179
 - logical encapsulation.....168
 - logical encapsulation (ATM-over-ADSL).....168
 - logical encapsulation (ATM-over-SHDSL).....178
 - physical encapsulation (ATM-over-ADSL).....167
 - physical encapsulation (ATM-over-SHDSL).....173, 176
 - verifying ATM-over-ADSL configuration.....185
- PPPoE (Point-to-Point Protocol over Ethernet)
 - address assignment (configuration editor).....203
 - address assignment (Quick Configuration).....197
 - CHAP (configuration editor).....204
 - CHAP (Quick Configuration).....197
 - CHAP local identity (Quick Configuration).....197
 - CHAP, overview.....194
 - client and server.....192
 - creating the pp0 interface (configuration editor).....201
 - discovery packets.....88, 194
 - encapsulation on an Ethernet interface.....88, 199
 - interfaces (Quick Configuration).....195
 - interfaces, overview.....193
 - logical interfaces (Quick Configuration).....197
 - MTU values.....203
 - overview.....192
 - See also* PPPoE over ATM-over-ADSL; PPPoE over ATM-over-SHDSL
 - preparation.....195
 - sample topology.....193
 - service type (configuration editor).....203
 - service type (Quick Configuration).....198
 - session limit (Quick Configuration).....198
 - session overview.....89, 194
 - session reconnection time (configuration editor).....202
 - session reconnection time (Quick Configuration).....198
 - underlying interface (Quick Configuration).....198
 - verifying interfaces.....208
 - verifying sessions.....209
 - verifying statistics.....210
 - verifying version information.....209
- PPPoE Active Discovery Initiation (PADI) packets.....88
- PPPoE Active Discovery Offer (PADO) packets.....88
- PPPoE Active Discovery Request (PADR) packets.....89
- PPPoE Active Discovery Session-Confirmation (PADS) packets.....89
- PPPoE Active Discovery Termination (PADT) packets.....89
- PPPoE encapsulation *See* PPPoE
- PPPoE interfaces *See* PPPoE
- PPPoE Interfaces Quick Configuration page.....196
- PPPoE over ATM LLC encapsulation
 - ATM-over-ADSL interfaces.....162, 168
 - ATM-over-SHDSL interfaces.....172, 178
- PPPoE over ATM-over-ADSL
 - CHAP.....204
 - creating the pp0 interface.....201
 - encapsulation.....200
 - overview.....193
 - See also* PPPoE
 - preparation.....195
 - sample topology.....193
 - verifying configuration.....206, 207
- PPPoE over ATM-over-SHDSL
 - CHAP.....204
 - creating the pp0 interface.....201
 - encapsulation.....200
 - overview.....193
 - See also* PPPoE
 - preparation.....195
 - verifying configuration.....206, 207
- PPPoEoA *See* PPPoE over ATM-over-ADSL; PPPoE over ATM-over-SHDSL
- preferences
 - for OSPF routes.....422
 - for static routes.....384
 - setting for static routes.....390
- primary stations, HDLC.....89

properties, verifying	
for ATM-over-ADSL network interfaces.....	181
for ATM-over-SHDSL network interfaces.....	186
for network interfaces.....	137
protocol data units <i>See</i> PDUs	
protocol families	
ccc.....	91
common protocol suites.....	91
inet.....	91
inet6.....	91
ISO.....	91
mlfr-end-to-end.....	91
mlfr-uni-nni.....	91
mlppp.....	91
MPLS.....	91
overview.....	91
tcc.....	91
tnp.....	91
protocols	
ARP.....	133
BGP <i>See</i> BGP	
CRTP.....	266, 289
distance vector <i>See</i> RIP	
EGPs.....	355
EIA-530.....	70
IGPs.....	355
IS-IS <i>See</i> IS-IS	
OSPF <i>See</i> OSPF	
overview.....	349
path vector <i>See</i> BGP	
PPPoE <i>See</i> PPPoE	
RIP <i>See</i> RIP	
RS-232.....	70
RS-422/449.....	71
serial.....	69
V.35.....	71
X.21.....	72
PSNP (partial sequence number PDU).....	370
publishing responses to ARP requests	
on Fast Ethernet subnets (Quick Configuration).....	114
on Gigabit Ethernet subnets (Quick Configuration).....	118
static ARP entries (configuration editor).....	133
PVCs (permanent virtual circuits)	
in multilink bundles, with MLFR FRF.15.....	284
<i>See also</i> MLFR FRF.15; multilink bundles	
in multilink bundles, with MLFR FRF.16.....	287
<i>See also</i> MLFR FRF.16; multilink bundles	
overview.....	84
preventing dropped packets on.....	310

Q

Q.931 timer, ISDN.....	150, 219, 225
------------------------	---------------

queuing with LFI	
data packets.....	269
on Q0 of constituent links.....	268
on Q2 of constituent links.....	268
overview.....	267
voice packets.....	269
Quick Configuration	
ATM-over-ADSL Interfaces page.....	159
ATM-over-SHDSL Interfaces page.....	170
Avaya VoIP.....	325
BGP page.....	443
buttons.....	8
E1 Interfaces page.....	106
E3 Interfaces page.....	109
Fast Ethernet Interfaces page.....	113
Gigabit Ethernet Interfaces page.....	116
Interfaces page.....	104
ISDN BRI Dialer Logical Interface page.....	221
ISDN BRI Physical Interface page.....	216
network interfaces.....	104
OSPF page.....	413
overview.....	7
PPPoE Interfaces page.....	196
RIP page.....	397
serial Interfaces page.....	128
Static Routes page.....	386
Summary page.....	8
T1 Interfaces page.....	121
T3 (DS3) Interfaces page.....	125
TGM550.....	325
VoIP.....	325

R

radio buttons <i>See</i> option buttons	
RADIUS authentication, of PPP sessions.....	195
RBBL <i>See</i> BBL	
reachability	
verifying for a RIP network.....	409
verifying for BGP peers.....	457
verifying for OSPF network hosts.....	430
reactivate command.....	30
real-time performance monitoring (RPM), for BGP	
peers.....	442
Refresh button.....	13
rejecting incoming calls, ISDN.....	243
relative option.....	36
release notes, URL.....	xix
Remote Authentication Dial-In User Service (RADIUS)	
authentication, of PPP sessions.....	195
rename command.....	28
renaming configuration identifiers.....	28
repeaters, on LAN segments.....	55
replace command.....	26
upto option.....	26

- replacing a configuration file.....36
 - example.....37
- reported bearer bandwidth limit *See* BBL
- request chassis fpc slot slot-number restart
 - command.....338
- request system configuration rescue delete
 - command.....34
- request system configuration rescue save
 - command.....34
- rescue configuration
 - deleting (CLI configuration editor).....33
 - deleting (J-Web).....21
 - disabling CONFIG or RESET CONFIG button
 - for.....34
 - loading with the CONFIG or RESET CONFIG
 - button.....21, 33
 - setting (CLI configuration editor).....33
 - setting (J-Web).....21
 - viewing (CLI configuration editor).....33
 - viewing (J-Web).....21
- reset button, for return to factory configuration *See*
 CONFIG button *See* CONFIG button; RESET CONFIG
 button
- RESET CONFIG button
 - default behavior.....21, 33
 - disabling.....34
 - return to factory configuration.....21, 33
- RIP (Routing Information Protocol)
 - authentication (RIPv2 only).....396
 - authentication (RIPv2 only), configuring.....404
 - basic network (configuration editor).....398
 - designating RIP interfaces.....398
 - distance vector protocol.....360
 - efficiency techniques.....362
 - enabling (Quick Configuration).....397
 - maximum hop count.....361
 - overview.....360, 395
 - packets.....361
 - path cost metrics *See* path cost metrics
 - poison reverse technique.....362
 - Quick Configuration.....396
 - requirements.....396
 - routing policy (configuration editor).....398
 - sample network with incoming metric.....401
 - sample network with outgoing metric.....403
 - sample topology.....399
 - split horizon technique.....362
 - supported versions.....360
 - traffic control with metrics *See* path cost metrics
 - traffic control with metrics, configuring.....401
 - unidirectional limitations.....363
 - verifying host reachability409
 - verifying RIP message exchange408
 - verifying RIP-enabled interfaces407
- RIP neighbors, verifying.....407
- RIP page.....397
 - field summary.....397
- rollback ? command.....33
- rollback command.....32
- rollback rescue command.....33
- rolling back a configuration file
 - during configuration (CLI configuration
 - editor).....32
 - during configuration (J-Web).....21
- route advertisements
 - AS path in.....375
 - BGP, update messages.....372
 - description.....357
 - external, EBGP.....372
 - internal, IBGP.....373
 - LSAs.....364
 - stub areas and NSSAs, to control.....367
- route aggregation.....358
- route origin, role in BGP route selection.....373
- route reflectors *See* BGP route reflectors
- route selection
 - BGP process for.....373
 - BGP, determining by AS path.....375
 - BGP, determining by local preference.....374
 - BGP, determining by MED metric.....376
 - BGP, lowest origin value preferred.....376
 - static routes, defining.....389
- router *See* Services Router
- router ID, role in BGP route selection.....374
- routing.....349
 - advertisements.....357
 - aggregation.....358
 - BGP *See* BGP
 - configuring PPPoE.....191
 - dynamic.....357
 - forwarding tables.....356
 - in multiple ASs with BGP.....441
 - in one AS with OSPF.....411
 - in one AS with RIP.....395
 - IS-IS *See* IS-IS
 - neighbors *See* BGP peers; OSPF neighbors; RIP
 - neighbors
 - OSPF *See* OSPF
 - protocol overview.....349
 - RIP *See* RIP
 - RIP statistics.....408
 - routing tables.....355
 - static *See* static routing
 - See also* protocols; routing policies; routing
 - solutions
- Routing Information Protocol *See* RIP
- routing mode, multi-port uPIMs.....343
- routing policies
 - BGP routing policy (configuration editor).....448
 - RIP routing policy (configuration editor).....398
- routing protocols *See* protocols

routing solutions	
applying CoS components on link services	
interface.....	302
BGP confederations, for scaling problems.....	451
BGP route reflectors, for scaling problems.....	448
BGP scaling techniques.....	378
controlling designated router election.....	425
controlling OSPF route cost.....	423
controlling OSPF route selection.....	422
controlling RIP traffic with the incoming	
metric.....	401
controlling RIP traffic with the outgoing	
metric.....	403
designated router, to reduce flooding.....	365
directing BGP traffic by local preference.....	374
drop-and-insert clock combinations.....	153
load balancing on link services interfaces.....	304
managing VoIP bandwidth <i>See</i> dynamic CAC	
NSSAs, to control route advertisement.....	367
path cost metrics, for packet flow control <i>See</i> path	
cost metrics	
point-to-point sessions over Ethernet.....	191
poison reverse, for traffic reduction.....	362
preventing dropped packets on PVCs.....	310
reducing jitter and latency on multilink	
bundles.....	303
securing OSPF routing (OSPFv2 only).....	424
split horizon, for traffic reduction.....	362
static route control techniques.....	384
stub areas, to control route advertisement.....	367
routing table	
controlling static routes in.....	384, 391
description.....	355
displaying static routes in.....	393
sample distance-vector routing.....	360
updates, limitations in RIP.....	363
verifying OSPF routes.....	429
RPM, for BGP peers.....	442
RS-232.....	70
RS-422/449.....	71
RS-530.....	70
RST (reset) button, TGM550.....	337
run command.....	34
S	
S/T interface	
overview.....	76
PIMs.....	214
samples	
drop-and-insert clock combinations.....	154
link services CoS.....	294
multilink bundle.....	292
PPPoA for ATM-over-ADSL configuration.....	185
PPPoE over ATM-over-ADSL configuration.....	207
PPPoE over ATM-over-SHDSL configuration.....	207
PPPoE over Ethernet configuration.....	206
saving configuration files.....	37
scaling BGP <i>See</i> BGP confederations; BGP route	
reflectors	
scheduler maps	
defining and applying.....	279
scheduling priority, overview.....	270
scheduling a commit.....	32
scope, IPv6 addresses	
global unicast.....	95
link-local unicast.....	95
multicast types.....	95
site-local unicast.....	95
screening incoming calls, ISDN.....	242
search-and-replace in configuration files.....	26
secondary stations, HDLC.....	89
secret, CHAP <i>See</i> CHAP, local identity	
security	
MD5 authentication for OSPF.....	425
MD5 authentication for RIPv2.....	405
password authentication for OSPFv2.....	425
password authentication for RIPv2.....	404
self-near-end crosstalk <i>See</i> SNEXT	
serial interfaces	
clocking modes.....	68
connection process.....	67
DTE default clock rate reduction.....	69
EIA-530.....	70
inverting the transmit clock.....	69
line protocols.....	69
MLPPP bundles and LFI (configuration	
editor).....	273
multilink bundles (Quick Configuration).....	271
overview.....	66
<i>See also</i> serial ports	
Quick Configuration.....	127
RS-232.....	70
RS-422/449.....	71
signal polarity.....	68
transmission signals.....	67
V.35.....	71
X.21.....	72
serial numbers, in MAC addresses.....	53
serial ports	
CHAP.....	129
clock rate.....	131
clocking.....	130
clocking, inverting the transmit clock.....	130
encapsulation type.....	129
line speed.....	131
logical interfaces.....	129
MTU.....	129
MTU default and maximum values.....	81

- overview.....66
 - See also* serial interfaces
- Quick Configuration.....127
- service provider ID *See* SPID
- service types, naming for PPPoE.....203
- services interfaces, overview.....101
 - See also* link services interface; multilink services
- Services Router
 - as a PPPoE client.....192
 - Avaya VoIP connectivity.....311
 - BGP routing.....441
 - channelized T1/E1/ISDN PRI interfaces.....139
 - configuration tools.....3
 - CPE, with PPPoE.....191
 - See also* PPPoE
 - DSL.....157
 - interfaces overview.....41
 - IS-IS protocol.....431
 - ISDN connections.....139, 211
 - link services interface.....261
 - link services interface, implementation
 - exceptions.....264
 - network interfaces.....103
 - OSPF routing.....411
 - PPPoE.....191
 - RIP routing.....395
 - routing protocols overview.....349
 - static routing.....383
- sessions
 - BGP session establishment.....372
 - BGP session maintenance.....372
 - ISDN session establishment.....77
 - limit on PPPoE sessions.....194
 - PPPoE.....89, 194
 - PPPoE, reconnection time (configuration editor).....202
 - PPPoE, reconnection time (Quick Configuration).....198
- set tgm fpc slot media-gateway-controller
 - command.....329
- shaping rate
 - applying.....283
 - overview.....269
 - requirement.....283
- SHDSL interfaces *See* ATM-over-SHDSL interfaces
- SHDSL page.....171
- SHDSL ports *See* ATM-over-SHDSL interfaces
- shortest path first algorithm.....364
- show access command.....185
- show bgp group command.....455
 - explanation.....455
- show bgp neighbor command.....454
 - explanation.....454
- show bgp summary command.....456
 - explanation.....456
- show chassis hardware command.....49
- show class-of-service classifier name command.....300
- show class-of-service command.....294
- show class-of-service interface command.....300
- show class-of-service scheduler-map command.....300
- show cli history command.....34
- show command.....24
- show interfaces at-3/0/0 command.....185
- show interfaces bc-0/0/4:1 extensive command.....249
- show interfaces br-6/0/0 extensive command.....247
- show interfaces command
 - for channelized interfaces.....151
 - for clear-channel channelized interfaces.....152
 - for multilink bundles.....292
 - for PPPoE over ATM-over-ADSL.....207
 - for PPPoE over Ethernet.....206
 - for the link services interface.....292
- show interfaces ct1-3/0/1 command.....151
- show interfaces dc-0/0/4 extensive command.....251
- show interfaces detail command.....137
- show interfaces dl0 extensive command.....255
- show interfaces e1-3/0/1 command.....152
- show interfaces extensive command.....181
 - explanation, for ATM-over-ADSL interfaces.....184
 - explanation, for ATM-over-SHDSL interfaces.....186, 188
 - explanation, for ISDN interfaces.....247, 250, 251
 - explanation, for VoIP interfaces.....340
- show interfaces ls-0/0/0 statistics detail
 - command.....296
 - explanation.....298
- show interfaces ppo command.....208
- show interfaces switch-port command.....346
- show interfaces vp-3/0/0 extensive command.....339
- show isdn calls command.....253
- show isdn status command.....246
- show isis adjacency brief command.....437
- show isis adjacency extensive command.....438
 - explanation.....438
- show ospf interface command.....427
 - explanation.....427
- show ospf neighbor command.....428
 - explanation.....428
- show ospf route command.....429
 - explanation.....430
- show pppoe interfaces command.....209
- show pppoe statistics command.....210
- show pppoe version command.....209
- show rip neighbor command.....407
 - explanation.....407
- show rip statistics command.....408
- show route terse command.....393
 - explanation.....393
- show system reboot command.....34
- show tgm dynamic-call-admission-control
 - command.....341

show tgm fpc slot-number media-gateway-controller command.....	341
show isis interface brief command.....	435
show isis interface detail command.....	436
explanation.....	436
signal-to-noise ratio <i>See</i> SNR	
signals	
DS1.....	58
E1 loopback (control).....	60
explicit clocking signal transmission.....	80
ISDN, disabling.....	244
multiplexing DS1 into DS2 signal.....	62
serial polarity.....	68
serial transmission.....	67
T1 loopback (control).....	60
V.35.....	71
X.21.....	72
single-area network, OSPF.....	415
site-local unicast IPv6 addresses.....	95
SNEXT (self-near-end crosstalk) threshold, SHDSL.....	174, 177
SNR (signal-to-noise ratio) margin, SHDSL.....	174, 177
source filtering, Gigabit Ethernet.....	118
for MAC addresses.....	119
sp-0/0/0 interface.....	99
special interfaces	
CRTP.....	102, 266, 289
dsc interface.....	100
IPv4 addressing.....	91
IPv6 addressing.....	94
logical properties.....	90
loopback interface.....	100
management interface.....	101
MLFR.....	101
<i>See also</i> link services interface; MLFR	
MLFR FRF.15 and FRF.16.....	102
<i>See also</i> link services interface	
MLPPP.....	101
<i>See also</i> link services interface; MLPPP	
names.....	48
naming conventions.....	47
output, understanding.....	49
overview.....	97
physical properties.....	78
protocol families.....	91
services interfaces.....	101
<i>See also</i> link services interface; multilink services	
summary.....	97
SPF (shortest path first) algorithm.....	364
SPID (service provider ID), ISDN.....	218, 225
split horizon technique.....	362
SSH connection to TGM550.....	335
statements	
adding or modifying.....	25
copying.....	27
deactivating.....	30
deleting.....	27
replacing.....	36
static ARP entries	
Fast Ethernet interface.....	113
Gigabit Ethernet interface.....	117
overview.....	133
static routes	
configuring basic routes (configuration editor).....	388
controlling.....	384
controlling in routing and forwarding tables.....	391
default properties.....	385
default properties, setting.....	392
defining route selection.....	389
preferences.....	384
preventing readvertisement.....	385
qualified next hops.....	384
Quick Configuration.....	386
rejecting passive traffic.....	385
requirements.....	385
route retention.....	385
sample preferred path.....	389
sample stub network.....	388
verifying.....	393
Static Routes page.....	386
field summary.....	387
static routing	
default gateway.....	387
description.....	357
overview.....	383
<i>See also</i> static routes	
static TEI (terminal endpoint identifier), ISDN.....	219, 225
statistics	
ATM-over-ADSL interfaces.....	184
ATM-over-SHDSL interfaces.....	189
ISDN B-channel interfaces.....	249
ISDN D-channel interfaces.....	251
link services interface.....	296
PPPoE.....	210
RIP.....	408
VoIP interface.....	339
status	
ATM-over-SHDSL interfaces, verifying.....	188
ISDN calls, verifying.....	253
ISDN interfaces, verifying.....	246
link services interface, verifying.....	296
link states, verifying.....	136
VoIP interface.....	339
status command.....	22
stub areas	
area ID (configuration editor).....	419
area ID (Quick Configuration).....	414
area type (Quick Configuration).....	414
controlling OSPF route cost.....	424

creating (configuration editor).....	420
description.....	367
example.....	368
sample topology.....	421
sub-ASs, BGP.....	381
subautonomous systems, BGP.....	381
subnet masks.....	93
subnets <i>See</i> subnetworks	
subnetworks	
description.....	354
IPv4 subnet addresses for multiple ISDN dialer	
interfaces.....	222
IPv4 subnets.....	93
route aggregation.....	359
Summary Quick Configuration page.....	8
superframe framing.....	59
support, technical <i>See</i> technical support	
SVCs (switched virtual circuits).....	84
switch types, supported, ISDN	
for ISDN BRI service.....	218, 225
for ISDN PRI service.....	141
switched virtual circuits (SVCs).....	84
switches	
configuring uPIMs as.....	343
on LAN segments.....	56
switching mode, multi-port uPIMs.....	343
symmetric high-speed digital subscriber line (SHDSL)	
<i>See</i> ATM-over-SHDSL interfaces	
synchronous networks.....	79
syntax conventions.....	xxi
system clock <i>See</i> clocking	
system identifier, IS-IS	
all zeros not supported.....	432
formats, MAC or IP address.....	432
identifier-to-hostname mapping.....	432
overview.....	369

T

T1 interfaces	
AMI encoding.....	59
B8ZS encoding.....	59
CRTP (configuration editor).....	289
D4 framing.....	59
data stream.....	57
dynamic CAC for voice packets (configuration	
editor).....	331
<i>See also</i> Avaya VoIP	
encoding.....	58
ESF framing.....	60
framing.....	59
loopback.....	60
multilink bundles (Quick Configuration).....	271
overview.....	57
<i>See also</i> T1 ports; channelized T1 interfaces	
Quick Configuration.....	120

signals.....	58
superframe framing.....	59
T1 ports	
cable length.....	124
CHAP.....	123
clocking.....	122
data inversion.....	123
encapsulation type.....	122
fractional, channel number.....	49
frame checksum.....	124
framing.....	123
logical interfaces.....	122
MTU.....	122
MTU default and maximum values.....	81
overview.....	57
<i>See also</i> T1 interfaces; channelized T1 ports	
Quick Configuration.....	120
time slots.....	123
T1/E1 media module <i>See</i> TIM510	
T3 interfaces	
bit stuffing.....	62
data stream.....	61
DS3 framing.....	63
multilink bundles (Quick Configuration).....	271
multiplexing on.....	63
overview.....	61
<i>See also</i> T3 ports	
Quick Configuration.....	124
T3 ports	
C-bit parity.....	127
cable length.....	127
CHAP.....	127
clocking.....	126
encapsulation type.....	126
frame checksum.....	127
framing.....	127
logical interfaces.....	126
MTU.....	126
MTU default and maximum values.....	81
overview.....	61
<i>See also</i> T3 interfaces	
Quick Configuration.....	124
tag protocol IDs (TPIDs), Gigabit Ethernet.....	119
tap interface.....	99
tcc protocol family.....	91
technical support	
contacting JTAC.....	xxv
TEI option, ISDN.....	219, 226
telephone calls	
rejecting incoming ISDN.....	243
screening incoming ISDN.....	242
verifying status.....	253
Telephony Gateway Module <i>See</i> TGM550	
Telephony Interface Module <i>See</i> TIM510; TIM514;	
TIM521	

Telnet access to TGM550	
connecting to TGM550.....	336
disabling Telnet service.....	336
enabling Telnet service.....	336
overview.....	335
security caution.....	335
telnet command.....	336
terminal endpoint identifier <i>See</i> static TEI; TEI option	
terminology	
Avaya VoIP.....	311
channelized T1/E1/ISDN PRI.....	139
configuration.....	3
DSL.....	157
interfaces.....	42
ISDN.....	211
link services.....	261
ports.....	42
PPPoE.....	191
routing protocols.....	349
TGM550	
accessing the router from.....	337
administration.....	333
Avaya CLI access.....	333
Avaya Media Gateway Controllers	
supported.....	317
CLI access requirements.....	334
console connection.....	334
grounding cable requirement.....	322
interfaces.....	316
IP address change caution.....	332
IP address, modifying (configuration editor).....	332
IP address, setting (configuration editor).....	327
IP addressing guidelines.....	320
JUNOS compatibility.....	320
MGC list, adding.....	329
MGC list, clearing.....	330
MGCs supported.....	317
Quick Configuration.....	325
reset on address change.....	332
resetting.....	337
RST (reset) button.....	337
saving the configuration.....	338
SSH connection.....	335
Telnet access.....	335
Telnet connection to router.....	337
verifying MGC list.....	341
verifying VoIP interface.....	339
three-way handshake.....	364
TIM510 interfaces.....	316
TIM514 interfaces.....	316
TIM521 interfaces.....	316
time slots	
dropping and inserting, on channelized T1/E1	
interfaces.....	145
E1.....	108
number in interface name.....	49
T1.....	123
tnp protocol family.....	91
top command.....	24
topology	
Avaya VoIP.....	315
data link layer.....	52
IPv4 subnets.....	93
PPPoE session on an ATM-over-ADSL loop.....	194
PPPoE session on an Ethernet loop.....	193
sample ATM-over-ADSL.....	73
sample BGP AS path.....	376
sample BGP confederation.....	451
sample BGP confederations.....	381
sample BGP external and internal links.....	447
sample BGP local preference use.....	375
sample BGP MED use.....	377
sample BGP peer network.....	445
sample BGP peer session.....	371
sample BGP route reflector (one	
cluster).....	379, 449
sample BGP route reflectors (cluster of	
clusters).....	380
sample BGP route reflectors (multiple	
clusters).....	380
sample distance-vector routing.....	360
sample Frame Relay network.....	84
sample ISDN network.....	75
sample LAN.....	96
sample LFI and multilink bundle network.....	274
sample multiarea OSPF routing.....	366
sample multilink bundle and LFI network.....	274
sample OSPF backbone area.....	367
sample OSPF multiarea network.....	417
sample OSPF network.....	429
sample OSPF network with stubs and	
NSSAs.....	367
sample OSPF single-area network.....	416
sample OSPF stub areas and NSSAs.....	421
sample poison reverse routing.....	363
sample RIP network.....	399
sample RIP network with incoming metric.....	401
sample RIP network with outgoing metric.....	403
sample route advertisement.....	358
sample route aggregation.....	359
sample router network.....	356
sample split horizon routing.....	362
sample static route.....	357
sample static route, preferred path.....	389
sample stub network for static routes.....	388
sample unidirectional routing.....	363
sample VLAN.....	97
VoIP.....	315
topology database, OSPF.....	411
TPIDs, Gigabit Ethernet.....	119
trace options, channelized ports.....	144

Traceroute page	
results for OSPF.....	430
results for RIP.....	409
traffic	
controlling with incoming RIP metric.....	401
controlling with outgoing RIP metric.....	403
transmit clock source <i>See</i> clocking	
troubleshooting	
applying CoS components on link services	
interface.....	302
Avaya VoIP.....	342
channelized T1/E1 interfaces.....	153
dialer interfaces, packet loss due to duplicate IP	
subnet addresses.....	222
dropped packets on PVCs.....	310
jitter and latency on multilink bundles.....	303
LFI and load balancing on multilink bundles.....	304
link services interface.....	301
two-dimensional parity.....	81
two-wire mode (2 ports), SHDSL <i>See</i> ATM-over-SHDSL	
interfaces	
types of interfaces.....	48

U

U interface	
overview.....	76
PIMs.....	214
umd0 interface.....	100
unicast IPv6 addresses.....	95
up command.....	24
uPIMs	
verifying port status.....	346
uploading a configuration file.....	15
upto option, replace command.....	26
UR-2 operating mode.....	164, 167
URLs	
Avaya VoIP support.....	317
release notes.....	xix
USB memory stick, for Avaya VoIP configuration.....	323
USB modem interface.....	100

V

V.35.....	71
variable-length subnet masks (VLSMs).....	93
VCI (virtual channel identifier)	
ATM-over-ADSL interfaces.....	162, 169
ATM-over-SHDSL interfaces.....	173, 179
PPPoE over ATM-over-ADSL interfaces.....	201
PPPoE over ATM-over-SHDSL interfaces.....	201
verification	
ATM-over-ADSL interface properties.....	181
ATM-over-SHDSL interface configuration.....	186
Avaya VoIP.....	338
B-channels.....	249

BGP configuration.....	456
BGP groups.....	455
BGP peer reachability.....	457
BGP peers (neighbors).....	454
channelized interfaces.....	151
channelized T1/E1/ISDN PRI interfaces.....	150
clear-channel interfaces.....	152
configuration syntax.....	31
D-channels.....	251
dialer interfaces.....	255
interface properties.....	137
interface properties for uPIM switches.....	346
IS-IS adjacencies.....	437
IS-IS adjacencies (detail).....	438
IS-IS interface configuration.....	435
IS-IS interface configuration (detail).....	436
IS-IS neighbors.....	437
IS-IS neighbors (detail).....	438
ISDN BRI interfaces.....	247
ISDN call status.....	253
ISDN PRI interface configuration.....	153
ISDN PRI interface operation.....	249
ISDN status.....	246
link services CoS.....	300
link services interface CoS configuration.....	294
link services interface status.....	296
link states.....	136
load balancing on the link services interface.....	307
multilink bundle configuration.....	292
network interfaces.....	135
OSPF host reachability.....	430
OSPF neighbors.....	428
OSPF routes.....	429
OSPF-enabled interfaces.....	427
packet encapsulation on link services	
interface.....	306
PPPoA for ATM-over-ADSL configuration.....	185
PPPoE interfaces.....	208
PPPoE over ATM-over-ADSL	
configuration.....	206, 207
PPPoE over ATM-over-SHDSL	
configuration.....	206, 207
PPPoE sessions.....	209
PPPoE statistics.....	210
PPPoE version.....	209
RIP host reachability.....	409
RIP message exchange.....	408
RIP-enabled interfaces.....	407
static routes in the routing table.....	393
VoIP.....	338
version	
OSPF, supported.....	364
PPPoE, verifying.....	209
RIP, supported.....	360
View Configuration Text page.....	10
virtual channel identifier <i>See</i> VCI	

virtual circuits	
DLCIs.....	84
overview.....	84
PVCs.....	84
SVCs.....	84
virtual LANs <i>See</i> VLANs	
virtual link, through the backbone area.....	366
virtual path identifier <i>See</i> VPI	
virtual router interface (VRI).....	343
VLAN tagging <i>See</i> Gigabit Ethernet ports, TPIDs	
VLANs (virtual LANs)	
LAN comparison.....	96
overview.....	96
topology.....	97
VLSMs (variable-length subnet masks).....	93
voice calls, not supported in dial-in or callback.....	238
voice over IP <i>See</i> Avaya VoIP	
voice packets	
integrating with data, with drop-and-insert.....	145
LFI handling.....	265
load-balancing and queuing behavior.....	269
speeding transmission with CRTP.....	266, 289
voice traffic latency, controlling with shaping	
rates.....	283
<i>See also</i> multilink bundling	
VoIP (voice over IP) <i>See</i> Avaya VoIP; VoIP interface	
VoIP interface	
addressing guidelines.....	320
correcting version incompatibility problem.....	342
IP address, modifying (configuration editor).....	332
IP address, setting (configuration editor).....	327
naming convention.....	316
Quick Configuration.....	325
unavailability, correcting.....	342
verifying.....	339
vp-0/0/0.....	316
<i>See also</i> VoIP interface	
VPI (virtual path identifier)	
ATM-over-ADSL interfaces.....	163, 166
ATM-over-SHDSL interfaces.....	173, 176
PPPoE over ATM-over-ADSL interfaces	
(configuration editor).....	200
PPPoE over ATM-over-SHDSL interfaces	
(configuration editor).....	200
VRI.....	343

W

WAN interfaces, configuring dynamic CAC on for Avaya	
VoIP.....	330
watch list, for ISDN backup.....	222, 232

X

X.21.....	72
-----------	----