



J-series™ Services Router

Configuration Guide

Release 7.4

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, California 94089
USA

408-745-2000

www.juniper.net

Part Number: 530-014089-01, 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, NetScreen Technologies, the NetScreen logo, NetScreen-Global Pro, ScreenOS, and GigaScreen are registered trademarks of Juniper Networks, Inc. in the United States and other countries.

The following are trademarks of Juniper Networks, Inc.: ERX, ESP, E-series, Instant Virtual Extranet, Internet Processor, J2300, J4300, J6300, J-Protect, J-series, J-Web, JUNOS, JUNOScope, JUNOScript, JUNOSe, M5, M7i, M10, M10i, M20, M40, M40e, M160, M320, M-series, MMD, NetScreen-5GT, NetScreen-5XP, NetScreen-5XT, NetScreen-25, NetScreen-50, NetScreen-204, NetScreen-208, NetScreen-500, NetScreen-5200, NetScreen-5400, NetScreen-IDP 10, NetScreen-IDP 100, NetScreen-IDP 500, NetScreen-Remote Security Client, NetScreen-Remote VPN Client, NetScreen-SA 1000 Series, NetScreen-SA 3000 Series, NetScreen-SA 5000 Series, NetScreen-SA Central Manager, NetScreen Secure Access, NetScreen-SM 3000, NetScreen-Security Manager, NMC-RX, SDX, Stateful Signature, T320, T640, T-series, and TX Matrix. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners. All specifications are subject to change without notice.

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.

Copyright © 2005, Juniper Networks, Inc. All rights reserved.

J-series™ Services Router Configuration Guide, Release 7.4
Copyright © 2005, Juniper Networks, Inc.
All rights reserved. Printed in USA.

Writing: Nidhi Bhargava, Michael Bushong, Taffy Everts, Walter Goralski, Joshua Kim, Jerry Isaac, Archana Maheshwari, Laura Phillips, Cheryl Potter, Frank Reade, Swapna Steiger, and Alan Twigg
Editing: Taffy Everts and Stella Hackell
Illustration: Faith Bradford Brown and Nathaniel Woodward
Cover Design: Edmonds Design

Revision History
14 September 2005—Revision 1.

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

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 (including the ERX-310, ERX-705, ERX-710, ERX-1410, ERX-1440, M5, M7i, M10, M10i, M20, M40, M40e, M160, M320, and T320 routers, T640 routing node, and the JUNOS and SDX-300 software) 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.

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.

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 Software solely as embedded in, and for execution on, Juniper equipment originally purchased by Customer from Juniper or an authorized Juniper reseller, unless the applicable Juniper documentation expressly permits installation on non-Juniper equipment.

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.

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, 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. Customer's use of the Software shall be subject to all such limitations and purchase of all applicable licenses.

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 Software on non-Juniper equipment where the Juniper documentation does not expressly permit installation 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; or (k) 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	xxi
-------------------------	-----

Part 1 **Using the Configuration Interfaces**

Chapter 1	Using J-Web Configuration Tools	3
------------------	--	---

Part 2 **Configuring Router Interfaces**

Chapter 2	Interfaces Overview	41
Chapter 3	Configuring Network Interfaces	103
Chapter 4	Configuring Point-to-Point Protocol over Ethernet	151
Chapter 5	Configuring ISDN	169

Part 3 **Configuring Routing Protocols**

Chapter 6	Routing Overview	203
Chapter 7	Configuring Static Routes	237
Chapter 8	Configuring a RIP Network	249
Chapter 9	Configuring an OSPF Network	265
Chapter 10	Configuring the IS-IS Protocol	287
Chapter 11	Configuring BGP Sessions	295

Part 4 Configuring Private Communications over Public Networks with MPLS

Chapter 12	Multiprotocol Label Switching Overview	315
Chapter 13	Configuring Signaling Protocols for Traffic Engineering	331
Chapter 14	Configuring Virtual Private Networks	343
Chapter 15	Configuring CLNS VPNs	367
Chapter 16	Configuring IPSec for Secure Packet Exchange	379

Part 5 Managing Multicast Transmissions

Chapter 17	Multicast Overview	395
Chapter 18	Configuring a Multicast Network	405

Part 6 Configuring DLSw Services

Chapter 19	Configuring Data Link Switching	417
-------------------	--	------------

Part 7 Configuring Routing Policy, Firewall Filters, and Class of Service

Chapter 20	Policy, Firewall Filter, and Class-of-Service Overview	433
Chapter 21	Configuring Routing Policies	459
Chapter 22	Configuring Firewall Filters and NAT	475
Chapter 23	Configuring Class of Service with DiffServ	529

Part 8 Index

Table of Contents

About This Guide	xxi
Objectives	xxi
Audience.....	xxii
How to Use This Guide	xxii
Document Conventions	xxiii
Related Juniper Networks Documentation.....	xxiv
Documentation Feedback.....	xxvi
Requesting Support.....	xxvi

Part 1

Using the Configuration Interfaces

Chapter 1

Using J-Web 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	6
Before You Begin.....	6
Using J-Web Quick Configuration.....	7
Using the J-Web Configuration Editor	8
Editing and Committing the Clickable Configuration	8
Editing the Clickable Configuration.....	8
Discarding Parts of a Candidate Configuration	11
Committing a Clickable Configuration.....	12
Viewing the Configuration Text	12
Editing and Committing the Configuration Text.....	13
Uploading a Configuration File.....	14
Managing Configuration Files with the J-Web Interface	15
Configuration Database and History Overview.....	16
Displaying Users Editing the Configuration	18
Comparing Configuration Files	18
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.....	24
Modifying the Configuration	25

Adding or Modifying a Statement or Identifier	26
Deleting a Statement or Identifier	26
Copying a Statement.....	27
Renaming an Identifier	27
Inserting an Identifier.....	28
Deactivating a Statement or Identifier.....	29
Committing a Configuration with the CLI.....	30
Verifying a Configuration	30
Committing a Configuration and Exiting Configuration Mode	31
Committing a Configuration That Requires Confirmation	31
Scheduling and Canceling a Commit	31
Loading a Previous Configuration File with the CLI	32
Setting or Deleting the Rescue Configuration with the CLI	33
Disabling the CONFIG Button	33
Entering Operational Mode Commands During Configuration.....	34
Managing Configuration Files with the CLI	34
Loading a New Configuration File	34
Saving a Configuration File.....	37

Part 2

Configuring Router Interfaces

Chapter 2

Interfaces Overview	41
Interfaces Terms	42
Network Interfaces	45
Media Types	46
Network Interface Naming	46
J-series Interface Naming Conventions	47
Understanding CLI Output for J-series Interfaces	49
Data Link Layer Overview.....	50
Physical Addressing.....	50
Network Topology.....	50
Error Notification	50
Frame Sequencing	50
Flow Control.....	51
Data Link Sublayers.....	51
MAC Addressing.....	51
Ethernet Interface Overview	52
Ethernet Access Control and Transmission	52
Collisions and Detection.....	53
Collision Detection	53
Backoff Algorithm	53
Collision Domains and LAN Segments	54
Repeaters	54
Bridges and Switches	54
Broadcast Domains	55
Ethernet Frames	55
T1 and E1 Interfaces Overview	56
T1 Overview.....	57

E1 Overview	57
T1 and E1 Signals	57
Encoding	58
AMI Encoding	58
B8ZS and HDB3 Encoding	58
T1 and E1 Framing	59
Superframe (D4) Framing for T1	59
Extended Superframe (ESF) Framing for T1	59
T1 and E1 Loopback Signals	60
T3 and E3 Interfaces Overview	60
Multiplexing DS1 Signals	60
DS2 Bit Stuffing	61
DS3 Framing	61
M13 Asynchronous Framing	62
C-Bit Parity Framing	63
Serial Interface Overview	65
Serial Transmissions	66
Signal Polarity	67
Serial Clocking Modes	67
Serial Interface Transmit Clock Inversion	68
DTE Clock Rate Reduction	68
Serial Line Protocols	68
EIA-530	69
RS-232	69
RS-422/449	70
V.35	70
X.21	71
ADSL Interface Overview	71
ADSL Systems	72
ADSL2 and ADSL2+	73
Asynchronous Transfer Mode	73
SHDSL Interface Overview	74
ISDN Interface Overview	74
ISDN Channels	74
ISDN Interfaces	74
Typical ISDN Network	75
NT Devices and S and T Interfaces	75
U Interface	76
ISDN Call Setup	76
Layer 2 ISDN Connection Initialization	76
Layer 3 ISDN Session Establishment	76
Interface Physical Properties	77
Bit Error Rate Testing	78
Interface Clocking	78
Data Stream Clocking	79
Explicit Clocking Signal Transmission	79
Frame Check Sequences	80
Cyclic Redundancy Checks and Checksums	80
Two-Dimensional Parity	80
Physical Encapsulation on an Interface	81
Frame Relay	81
Virtual Circuits	82
Switched and Permanent Virtual Circuits	82

Data-Link Connection Identifiers	82
Congestion Control and Discard Eligibility	82
Point-to-Point Protocol	83
Link Control Protocol	83
CHAP Authentication	84
Network Control Protocols	84
Magic Numbers	85
CSU/DSU Devices	85
Point-to-Point Protocol over Ethernet	86
PPPoE Discovery	86
PPPoE Sessions	87
High-Level Data Link Control	87
HDLC Stations	87
HDLC Operational Modes	88
Interface Logical Properties	88
Protocol Families	89
Common Protocol Suites	89
Other Protocol Suites	90
IPv4 Addressing	90
IPv4 Classful Addressing	90
IPv4 Dotted Decimal Notation	91
IPv4 Subnetting	91
IPv4 Variable-Length Subnet Masks	92
IPv6 Addressing	93
IPv6 Address Representation	93
IPv6 Address Types	94
IPv6 Address Scope	94
IPv6 Address Structure	94
Virtual LANs	95
Special Interfaces	96
Discard Interface	98
Loopback Interface	99
Management Interface	99
Services Interfaces	100
MLPPP and MLFR	100
MLFR Frame Relay Forum	100
CRTP	100

Chapter 3

Configuring Network Interfaces

103

Before You Begin	103
Configuring Network Interfaces with Quick Configuration	104
Configuring an E1 Interface with Quick Configuration	105
Configuring an E3 Interface with Quick Configuration	108
Configuring a Fast Ethernet Interface with Quick Configuration	113
Configuring a T1 Interface with Quick Configuration	115
Configuring a T3 Interface with Quick Configuration	119
Configuring a Serial Interface with Quick Configuration	122
Configuring Network Interfaces with a Configuration Editor	126
Adding a Network Interface with a Configuration Editor	126

	Adding an ATM-over-ADSL Network Interface with a Configuration Editor	128
	Configuring CHAP on the ATM-over-ADSL Interface (Optional)	133
	Adding an ATM-over-SHDSL Interface	135
	Configuring Compressed Real-Time Transport Protocol (CRTP)	140
	Deleting a Network Interface with a Configuration Editor	142
	Verifying Interface Configuration	143
	Verifying the Link State of All Interfaces	143
	Verifying Interface Properties	144
	Verifying ADSL Interface Properties	145
	Displaying a PPPoA Configuration for an ATM-over-ADSL Interface	149
Chapter 4	Configuring Point-to-Point Protocol over Ethernet	151
	PPPoE Terms	151
	PPPoE Overview	152
	PPPoE Interfaces	153
	Fast Ethernet Interface	153
	ATM-over-ADSL Interface	153
	PPPoE Stages	154
	PPPoE Discovery Stage	154
	PPPoE Session Stage	154
	Optional CHAP Authentication	155
	Before You Begin	155
	Configuring PPPoE with a Configuration Editor	155
	Setting the Appropriate Encapsulation on the Interface (Required)	155
	Configuring PPPoE Encapsulation on an Ethernet Interface	156
	Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface	157
	Configuring a PPPoE Interface (Required)	158
	Configuring CHAP (Optional)	161
	Verifying a PPPoE Configuration	162
	Displaying a PPPoE Configuration for an Ethernet Interface	162
	Displaying a PPPoE Configuration for an ATM-over-ADSL Interface	163
	Verifying PPPoE Interfaces	164
	Verifying PPPoE Sessions	165
	Verifying the PPPoE Version	166
	Verifying PPPoE Statistics	166
Chapter 5	Configuring ISDN	169
	ISDN Terms	169
	ISDN Overview	170
	ISDN Interfaces	171
	Before You Begin	172
	Configuring ISDN Interfaces with a Configuration Editor	172
	Adding an ISDN Interface (Required)	172
	Configuring a Dialer Interface (Required)	175
	Enabling an ISDN Interface as a Secondary Connection (Optional)	178
	Configuring Dial-on-Demand Connectivity (Optional)	179
	Configuring a Dialer Filter	179
	Applying the Dial-on-Demand Dialer Filter to the Dialer Interface	180
	Configuring Bandwidth-on-Demand (Optional)	181

Configuring a Dialer Interface for Bandwidth-on-Demand	181
Configuring an ISDN Interface for Bandwidth-on-Demand	183
Configuring Dial-on-Demand Routing (Optional)	184
Configuring the Dial-on-Demand Dialer Filter	184
Applying the Dial-on-Demand Dialer Filter to the Dialer Interface...	185
Configuring Dialer Watch (Optional)	186
Adding a Dialer Watch Interface on the Services Router	186
Configuring the ISDN Interface for Dialer Watch	189
Configuring Dial-on-Demand Routing with OSPF Support (Optional)	190
Configuring Dialer Profiles (Optional)	191
Verifying the ISDN Configuration	192
Displaying the ISDN Status	193
Verifying an ISDN Interface	193
Checking B-Channel Statistics	194
Checking D-Channel Interface Statistics	196
Verifying Dialer Interface Configuration	197

Part 3

Configuring Routing Protocols

Chapter 6

Routing Overview 203

Routing Terms	203
Routing Overview	207
Networks and Subnetworks	208
Autonomous Systems	208
Interior and Exterior Gateway Protocols	208
Routing Tables	209
Forwarding Tables	209
Dynamic and Static Routing	210
Route Advertisements	211
Route Aggregation	211
RIP Overview	213
Distance-Vector Routing Protocols	213
Maximizing Hop Count	214
RIP Packets	215
Split Horizon and Poison Reverse Efficiency Techniques	215
Limitations of Unidirectional Connectivity	216
OSPF Overview	217
Link-State Advertisements	218
Role of the Designated Router	218
Path Cost Metrics	219
Areas and Area Border Routers	219
Role of the Backbone Area	220
Stub Areas and Not-So-Stubby Areas	221
IS-IS Overview	222
IS-IS Areas	223
Network Entity Titles and System Identifiers	223
IS-IS Path Selection	224
Protocol Data Units	224

IS-IS Hello PDU	224
Link-State PDU	224
Complete Sequence Number PDU	225
Partial Sequence Number PDU	225
BGP Overview	225
Point-to-Point Connections	226
BGP Messages for Session Establishment	226
BGP Messages for Session Maintenance	227
IBGP and EBGP	227
Route Selection	228
Local Preference	228
AS Path	229
Origin	230
Multiple Exit Discriminator	230
Scaling BGP for Large Networks	231
Route Reflectors—for Added Hierarchy	231
Confederations—for Subdivision	234
Chapter 7	Configuring Static Routes 237
Static Routing Overview	237
Static Route Preferences	238
Qualified Next Hops	238
Control of Static Routes	238
Route Retention	239
Readvertisement Prevention	239
Forced Rejection of Passive Route Traffic	239
Default Properties	239
Before You Begin	240
Configuring Static Routes with Quick Configuration	240
Configuring Static Routes with a Configuration Editor	242
Configuring a Basic Set of Static Routes (Required)	242
Controlling Static Route Selection (Optional)	243
Controlling Static Routes in the Routing and Forwarding Tables (Optional)	245
Defining Default Behavior for All Static Routes (Optional)	246
Verifying the Static Route Configuration	247
Displaying the Routing Table	247
Chapter 8	Configuring a RIP Network 249
RIP Overview	249
RIP Traffic Control with Metrics	249
Authentication	250
Before You Begin	250
Configuring a RIP Network with Quick Configuration	250
Configuring a RIP Network with a Configuration Editor	253
Configuring a Basic RIP Network (Required)	253
Controlling Traffic in a RIP Network (Optional)	256
Controlling Traffic with the Incoming Metric	256
Controlling Traffic with the Outgoing Metric	257
Enabling Authentication for RIP Exchanges (Optional)	259

	Enabling Authentication with Plain-Text Passwords	259
	Enabling Authentication with MD5 Authentication	260
	Verifying the RIP Configuration	261
	Verifying the RIP-Enabled Interfaces	261
	Verifying the Exchange of RIP Messages	261
	Verifying Reachability of All Hosts in the RIP Network	262
Chapter 9	Configuring an OSPF Network	265
	OSPF Overview	265
	Enabling OSPF	265
	OSPF Areas	266
	Path Cost Metrics	266
	OSPF Dial-on-Demand Circuits	266
	Before You Begin	267
	Configuring an OSPF Network with Quick Configuration	267
	Configuring an OSPF Network with a Configuration Editor	269
	Configuring the Router Identifier (Required)	270
	Configuring a Single-Area OSPF Network (Required)	270
	Configuring a Multiarea OSPF Network (Optional)	272
	Creating the Backbone Area	273
	Creating Additional OSPF Areas	273
	Configuring Area Border Routers	274
	Configuring Stub and Not-So-Stubby Areas (Optional)	275
	Tuning an OSPF Network for Efficient Operation	277
	Controlling Route Selection in the Forwarding Table	278
	Controlling the Cost of Individual Network Segments	278
	Enabling Authentication for OSPF Exchanges	279
	Controlling Designated Router Election	280
	Verifying an OSPF Configuration	281
	Verifying OSPF-Enabled Interfaces	281
	Verifying OSPF Neighbors	282
	Verifying the Number of OSPF Routes	283
	Verifying Reachability of All Hosts in an OSPF Network	284
Chapter 10	Configuring the IS-IS Protocol	287
	IS-IS Overview	287
	ISO Network Addresses	287
	System Identifier Mapping	288
	Before You Begin	289
	Configuring IS-IS with a Configuration Editor	289
	Verifying IS-IS on a Services Router	290
	Displaying IS-IS Interface Configuration	291
	Displaying IS-IS Interface Configuration Detail	291
	Displaying IS-IS Adjacencies	292
	Displaying IS-IS Adjacencies in Detail	293
Chapter 11	Configuring BGP Sessions	295

BGP Overview	295
BGP Peering Sessions	295
IBGP Full Mesh Requirement	296
Route Reflectors and Clusters	296
BGP Confederations	296
Before You Begin	297
Configuring BGP Sessions with Quick Configuration	297
Configuring BGP Sessions with a Configuration Editor	299
Configuring a Point-to-Point Peering Session (Required)	299
Configuring BGP Within a Network (Required)	302
Configuring a Route Reflector (Optional)	303
Configuring BGP Confederations (Optional)	306
Verifying a BGP Configuration	308
Verifying BGP Neighbors	308
Verifying BGP Groups	309
Verifying BGP Summary Information	310
Verifying Reachability of All Peers in a BGP Network	311

Part 4

Configuring Private Communications over Public Networks with MPLS

Chapter 12

Multiprotocol Label Switching Overview 315

MPLS and VPN Terms	315
MPLS Overview	317
Label Switching	318
Label-Switched Paths	318
Label-Switching Routers	319
Labels	320
Label Operations	320
Penultimate Hop Popping	321
LSP Establishment	321
Static LSPs	321
Dynamic LSPs	321
Signaling Protocols Overview	322
Label Distribution Protocol	322
LDP Operation	322
LDP Messages	322
Resource Reservation Protocol	322
RSVP Fundamentals	323
Bandwidth Reservation Requirement	323
Explicit Route Objects	323
Constrained Shortest Path First	325
Link Coloring	325
VPN Overview	326
VPN Components	326
VPN Routing Requirements	327
VPN Routing Information	328
VRF Instances	328
Route Distinguishers	328

	Route Targets to Control the VRF Table	329
	Types of VPNs	329
	Layer 2 VPNs	329
	Layer 2 Circuits	329
	Layer 3 VPNs	329
Chapter 13	Configuring Signaling Protocols for Traffic Engineering	331
	Signaling Protocol Overview	331
	LDP Signaling Protocol	332
	RSVP Signaling Protocol	332
	Before You Begin	332
	Configuring LDP and RSVP with a Configuration Editor	333
	Configuring LDP-Signaled LSPs	333
	Configuring RSVP-Signaled LSPs	335
	Verifying an MPLS Configuration	338
	Verifying an LDP-Signaled LSP	338
	Verifying LDP Neighbors	338
	Verifying LDP Sessions	339
	Verifying the Presence of LDP-Signaled LSPs	340
	Verifying Traffic Forwarding over the LDP-Signaled LSP	340
	Verifying an RSVP-Signaled LSP	341
	Verifying RSVP Neighbors	341
	Verifying RSVP Sessions	341
	Verifying the Presence of RSVP-Signaled LSPs	342
Chapter 14	Configuring Virtual Private Networks	343
	VPN Configuration Overview	343
	Sample VPN Topology	344
	Basic Layer 2 VPN Configuration	344
	Basic Layer 2 Circuit Configuration	345
	Basic Layer 3 VPN Configuration	345
	Before You Begin	346
	Configuring VPNs with a Configuration Editor	346
	Configuring Interfaces Participating in a VPN	347
	Configuring Protocols Used by a VPN	349
	Configuring MPLS for VPNs	349
	Configuring a BGP Session	351
	Configuring Routing Options for VPNs	352
	Configuring an IGP and a Signaling Protocol	353
	Configuring LDP for Signaling	353
	Configuring RSVP for Signaling	355
	Configuring a Layer 2 Circuit	356
	Configuring a VPN Routing Instance	357
	Configuring a VPN Routing Policy	359
	Configuring a Routing Policy for Layer 2 VPNs	360
	Configuring a Routing Policy for Layer 3 VPNs	363
	Verifying a VPN Configuration	364
	Pinging a Layer 2 VPN	365

	Pinging a Layer 3 VPN	365
	Pinging a Layer 2 Circuit	365
Chapter 15	Configuring CLNS VPNs	367
	CLNS Terms	367
	CLNS Overview	368
	Before You Begin.....	369
	Configuring CLNS with a Configuration Editor	369
	Configuring a VPN Routing Instance (Required)	370
	Configuring ES-IS	371
	Configuring IS-IS for CLNS	372
	Configuring CLNS Static Routes.....	374
	Configuring BGP for CLNS.....	375
	Verifying CLNS VPN Configuration	376
	Displaying CLNS VPN Configuration	376
Chapter 16	Configuring IPsec for Secure Packet Exchange	379
	IPsec Tunnel Overview	379
	Security Associations	380
	Translating Outgoing Traffic.....	380
	Before You Begin.....	380
	Configuring an IPsec Tunnel with Quick Configuration	380
	Configuring an IPsec Tunnel with a Configuration Editor	382
	Configuring IPsec Services Interfaces	383
	Configuring IPsec Service Sets.....	384
	Configuring an IPsec Stateful Firewall Filter Rule	387
	Configuring a NAT Pool	389
	Verifying the IPsec Tunnel Configuration	391
	Verifying IPsec Tunnel Statistics	391
Part 5	Managing Multicast Transmissions	
Chapter 17	Multicast Overview	395
	Multicast Terms	395
	Multicast Architecture	398
	Upstream and Downstream Interfaces.....	398
	Subnetwork Leaves and Branches	398
	Multicast IP Address Ranges	399
	Notation for Multicast Forwarding States	399
	Dense and Sparse Routing Modes.....	400
	Strategies for Preventing Routing Loops	400
	Reverse-Path Forwarding for Loop Prevention.....	400
	Shortest-Path Tree for Loop Prevention	401
	Administrative Scoping for Loop Prevention.....	401
	Multicast Protocol Building Blocks.....	401

Chapter 18 Configuring a Multicast Network 405

Before You Begin.....	406
Configuring a Multicast Network with a Configuration Editor.....	406
Configuring SAP and SDP (Optional)	406
Configuring IGMP (Required).....	407
Configuring the PIM Static RP (Optional)	408
Configuring a PIM RPF Routing Table (Optional)	410
Verifying a Multicast Configuration.....	411
Verifying SAP and SDP Addresses and Ports.....	412
Verifying the IGMP Version.....	412
Verifying the PIM Mode and Interface Configuration	413
Verifying the PIM RP Configuration	413
Verifying the RPF Routing Table Configuration	414

Part 6 Configuring DLSw Services

Chapter 19 Configuring Data Link Switching 417

DLSw Terms.....	417
Data Link Switching (DLSw) Overview	418
Switch-to-Switch Protocol for DLSw	419
DLSw Operational Stages.....	419
DLSw Capabilities Exchange	419
DLSw Circuits Establishment.....	420
Before You Begin.....	420
Configuring Basic DLSw with a Configuration Editor	420
Configuring LLC2 Properties on an Ethernet Interface (Required)	421
Configuring DLSw on the Local Services Router (Required)	422
Configuring DLSw on the Remote Services Router (Required)	423
Configuring Class-of-Service (CoS) for DLSw (Optional)	424
Verifying DLSw Configuration	426
Displaying LLC2 Properties on a Fast Ethernet Interface	427
Displaying DLSw Capabilities	427
Displaying DLSw Circuit State	427
Displaying Details of a DLSw Circuit State	428
Displaying DLSw Peers	428
Displaying Details of DLSw Peers	428
Displaying DLSw Reachability Information	429

Part 7 Configuring Routing Policy, Firewall Filters, and Class of Service

Chapter 20 Policy, Firewall Filter, and Class-of-Service Overview 433

Policy, Firewall Filter, and CoS Terms	433
--	-----

Routing Policy Overview	435
Routing Policy Components.....	435
Routing Policy Terms.....	436
Routing Policy Match Conditions.....	436
Routing Policy Actions	438
Default and Final Actions.....	440
Applying Routing Policies	440
Firewall Filter Overview	440
Stateful and Stateless Firewall Filters.....	441
Process for Configuring a Stateful Firewall Filter and NAT	442
Summary of Stateful Firewall Filter and NAT Match Conditions and Actions.....	442
Planning a Stateless Firewall Filter	444
Stateless Firewall Filter Match Conditions, Actions, and Action Modifiers.....	445
Class-of-Service Overview.....	449
Benefits of DiffServ CoS	449
DSCPs and Forwarding Service Classes	450
JUNOS CoS Functions.....	451
How Forwarding Classes and Schedulers Work.....	453
Default Forwarding Class Queue Assignments.....	453
Default Scheduler Settings.....	454
Default Behavior Aggregate (BA) Classifiers	455
DSCP Rewrites.....	456
Sample BA Classification	456

Chapter 21**Configuring Routing Policies 459**

Before You Begin.....	460
Configuring a Routing Policy with a Configuration Editor	460
Configuring the Policy Name (Required)	461
Configuring a Policy Term (Required)	461
Rejecting Known Invalid Routes (Optional).....	462
Injecting OSPF Routes into the BGP Routing Table (Optional).....	464
Grouping Source and Destination Prefixes in a Forwarding Class (Optional).....	466
Configuring a Policy to Prepend the AS Path (Optional)	467
Configuring Damping Parameters (Optional)	470

Chapter 22**Configuring Firewall Filters and NAT 475**

Before You Begin.....	476
Configuring a Stateful Firewall Filter with Quick Configuration.....	476
Configuring a Stateful Firewall Filter with a Configuration Editor.....	480
Configuring a Stateless Firewall Filter with Quick Configuration	486
Configuring IPv4 and IPv6 Stateless Firewall Filters	487
Assigning IPv4 and IPv6 Firewall Filters to Interfaces	499
Configuring a Stateless Firewall Filter with a Configuration Editor	501
Stateless Firewall Filter Strategies	502
Strategy for a Typical Stateless Firewall Filter.....	502
Strategy for Handling Packet Fragments	502

Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources	502
Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods	506
Configuring a Routing Engine Firewall Filter to Handle Fragments	511
Applying a Stateless Firewall Filter to an Interface	516
Verifying Firewall Filter Configuration	517
Displaying Firewall Filter Configurations	517
Verifying a Stateful Firewall Filter	522
Displaying Firewall Filter Logs	523
Displaying Firewall Filter Statistics	524
Verifying a Services, Protocols, and Trusted Sources Firewall Filter	525
Verifying a TCP and ICMP Flood Firewall Filter	526
Verifying a Firewall Filter That Handles Fragments	527

Chapter 23

Configuring Class of Service with DiffServ

529

Before You Begin	530
Configuring CoS with DiffServ with a Configuration Editor	530
Configuring a Policer for a Firewall Filter (Required)	531
Configuring and Applying a Firewall Filter for a Multifield Classifier (Required)	532
Assigning Forwarding Classes to Output Queues (Required)	535
Configuring and Applying Rewrite Rules (Required)	537
Configuring and Applying Behavior Aggregate Classifiers (Required)	542
Configuring RED Drop Profiles for Assured Forwarding Congestion Control (Required)	546
Configuring Schedulers (Optional)	548
Configuring and Applying Scheduler Maps (Optional)	552
Configuring and Applying Virtual Channels (Optional)	555
Configuring and Applying Adaptive Shaping (Optional)	559
Verifying a DiffServ Configuration	560
Verifying Multicast Session Announcements	561
Verifying an Adaptive Shaper Configuration	561
Verifying a Virtual Channel Configuration	562
Verifying a Virtual Channel Group Configuration	562

Part 8

Index

Index	565
-------------	-----

About This Guide

This preface provides the following guidelines for using the *J-series™ Services Router Configuration Guide*:

- Objectives on page xxi
- Audience on page xxii
- How to Use This Guide on page xxii
- Document Conventions on page xxiii
- Related Juniper Networks Documentation on page xxiv
- Documentation Feedback on page xxvi
- Requesting Support on page xxvi

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, configure virtual private networks (VPNs), configure and manage multicast networks, and apply routing techniques such as policies, firewall filters, IP Security (IPSec) tunnels, and class-of-service (CoS) classification for safe, efficient routing.



NOTE: This guide documents Release 7.4 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>.

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) to perform the tasks shown in Table 1.

Table 1: Capabilities of J-series Interfaces

J-series Interface	Capabilities
J-Web graphical browser interface	<ul style="list-style-type: none"> ■ Quick (basic) configuration ■ Monitoring, configuration, diagnosis, and management
JUNOS CLI	Monitoring, configuration, diagnosis, and management

J-series Services Router guides provide complete instructions for using the J-Web interface, but they are not a comprehensive resource for using the JUNOS CLI. For CLI information, see the JUNOS software manuals listed in “Related Juniper Networks Documentation” on page xxiv.

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

Because you can configure and manage a Services Router in several ways, most chapters in J-series Services Router guides contain multiple sets of instructions:

- Configuration—For many Services Router features, you can use J-Web Quick Configuration for basic setup. For more extensive configuration of all Services Router features, use the J-Web configuration editor or the JUNOS CLI configuration editor.
- Maintenance—To monitor, diagnose, and manage a Services Router, use the J-Web interface for common tasks, or use CLI operational mode commands.

J-series Services Routers are documented in three guides. Table 2 shows where Services Router instructions are located.

Table 2: Location of Tasks in J-series Guides

Services Router Tasks	Location of Instructions
Installing hardware and establishing basic connectivity	<i>J-series Services Router Getting Started Guide</i>
Configuring interfaces and routing protocols	<i>J-series Services Router Configuration Guide</i>
Managing users and operations, monitoring performance, upgrading software, and diagnosing common problems	<i>J-series Services Router Administration Guide</i>

Document Conventions

Table 3 defines the notice icons used in this guide.

Table 3: 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.

Table 4 defines the text and syntax conventions used in this guide.

Table 4: 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>

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 <i>area-id</i>] 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 three 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 5.

Table 5: J-series Guides and Related JUNOS Software Publications

Chapter in a J-series Guide	Corresponding JUNOS Software Manual
J-series Services Router Getting Started Guide	
“J-series User Interface Overview”	<i>JUNOS System Basics Configuration Guide</i>
“Establishing Basic Connectivity”	
“Configuring Autoinstallation”	
J-series Services Router Configuration Guide	
“Using J-Web Configuration Tools”	<i>JUNOS System Basics Configuration Guide</i>
“Interfaces Overview”	■ <i>JUNOS Network Interfaces Configuration Guide</i>
“Configuring Network Interfaces”	■ <i>JUNOS Interfaces Command Reference</i>
“Configuring Point-to-Point Protocol over Ethernet”	
“Configuring ISDN”	
“Routing Overview”	■ <i>JUNOS Routing Protocols Configuration Guide</i>
“Configuring Static Routes”	■ <i>JUNOS Routing Protocols and Policies Command Reference</i>
“Configuring a RIP Network”	
“Configuring an OSPF Network”	
“Configuring the IS-IS Protocol”	
“Configuring BGP Sessions”	
“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, Firewall Filter, and Class-of-Service Overview”	<i>JUNOS Policy Framework Configuration Guide</i>
“Configuring Routing Policies”	<i>JUNOS Routing Protocols and Policies Command Reference</i>

Chapter in a J-series Guide	Corresponding JUNOS Software Manual
“Configuring Firewall Filters and NAT”	<ul style="list-style-type: none"> ■ <i>JUNOS Network Interfaces Configuration Guide</i> ■ <i>JUNOS Policy Framework Configuration Guide</i> ■ <i>JUNOS Services Interfaces Configuration Guide</i> ■ <i>JUNOS System Basics and Services Command Reference</i> ■ <i>JUNOS Routing Protocols and Policies Command Reference</i>
“Configuring Class of Service with DiffServ”	<ul style="list-style-type: none"> ■ <i>JUNOS Class of Service Configuration Guide</i> ■ <i>JUNOS System Basics and Services Command Reference</i>
J-series Services Router Administration Guide	
“Managing Users and Operations”	<i>JUNOS System Basics Configuration Guide</i>
“Configuring SNMP for Network Management”	<i>JUNOS Network Management Configuration Guide</i>
“Configuring the DHCP Server”	<i>JUNOS System Basics Configuration Guide</i>
“Configuring and Monitoring Alarms”	<i>JUNOS System Basics Configuration Guide</i>
“Monitoring a Services Router”	<ul style="list-style-type: none"> ■ <i>JUNOS System Basics and Services Command Reference</i>
“Using Services Router Diagnostic Tools”	<ul style="list-style-type: none"> ■ <i>JUNOS Interfaces Command Reference</i> ■ <i>JUNOS Routing Protocols and Policies Command Reference</i>
“Monitoring Real-Time Performance”	<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 J-Web Configuration Tools on page 3

Chapter 1

Using J-Web Configuration Tools

Use J-Web 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 6
- Using J-Web Quick Configuration on page 7
- Using the J-Web Configuration Editor on page 8
- Managing Configuration Files with the J-Web Interface on page 15
- Using the CLI Configuration Editor on page 22
- Managing Configuration Files with the CLI on page 34

Configuration Tools Terms

Before using the Services Router configuration tools, become familiar with the terms defined in Table 6.

Table 6: 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 6: 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 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. For a comparison of configuration interfaces, see the *J-series Services Router Getting Started Guide*.

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.

J-Web Configuration Options

You access the J-Web interface configuration tools by selecting **Configuration** in the task bar. Table 7 describes the J-Web configuration options.

Table 7: J-Web Configuration Options

Option	Purpose	Description
Quick Configuration	Basic configuration	Displays options for quick Services Router configuration— Set Up , SSL , Interfaces , Users , SNMP , Routing , Firewall/NAT , and IPSec Tunnels . 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 , Edit Configuration , Edit Configuration Text , and Upload Configuration File . For more information, see “Using the J-Web Configuration Editor” on page 8.
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 15.
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 8 provides a summary of the top-level CLI configuration commands.

Table 8: Top-Level CLI Configuration Commands

Command	Function
Managing the Configuration and Configuration Files	
<code>commit</code>	Commit the set of configuration changes in the candidate configuration to take operational effect.

Table 8: Top-Level CLI Configuration Commands (continued)

Command	Function
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.
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. For more information about configuring

access privilege levels, see the *J-series Services Router Administration Guide* and the *JUNOS System Basics Configuration Guide*.

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 either the side pane or the main pane (see Figure 1). To configure the Services Router using Quick Configuration, see the configuration sections in this manual.

Figure 1: J-Web Quick Configuration Options

Juniper NETWORKS

ROUTER - J4300

Logged in as: **regress**

[Help](#) [About](#) [Logout](#)

Monitor **Configuration** **Diagnose** **Manage**

[Configuration > Quick Configuration > Summary](#)

Quick Configuration

Summary

Router Configuration

The following pages help you to configure your router quickly and easily. They provide access to the most commonly configured parameters and are useful in generating the initial configuration of the router.

- ▶ **Set Up**
Define network identification, default gateway, name and time servers, root user authentication, and basic local network access to the system.
- ▶ **SSL**
Configure certificates and SSL access methods.
- ▶ **Interfaces**
List all interfaces installed on system and configure logical interfaces and common interface parameters.
- ▶ **Users**
Define users allowed to access the router and configure authentication servers. Pick authorization level for each user.

▶ **Quick Configuration**

- Set Up
- SSL
- Interfaces
- Users
- SNMP
- Routing
- Firewall/NAT
- IPSec Tunnels
- Realtime Performance Monitoring

▶ **View and Edit**

▶ **History**

▶ **Rescue**

Table 9 describes the functions of the buttons that appear in the J-Web Quick Configuration pages.

Table 9: J-Web Quick Configuration Buttons

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

- Editing and Committing the Clickable Configuration on page 8
- Viewing the Configuration Text on page 12
- Editing and Committing the Configuration Text on page 13
- Uploading a Configuration File on page 14

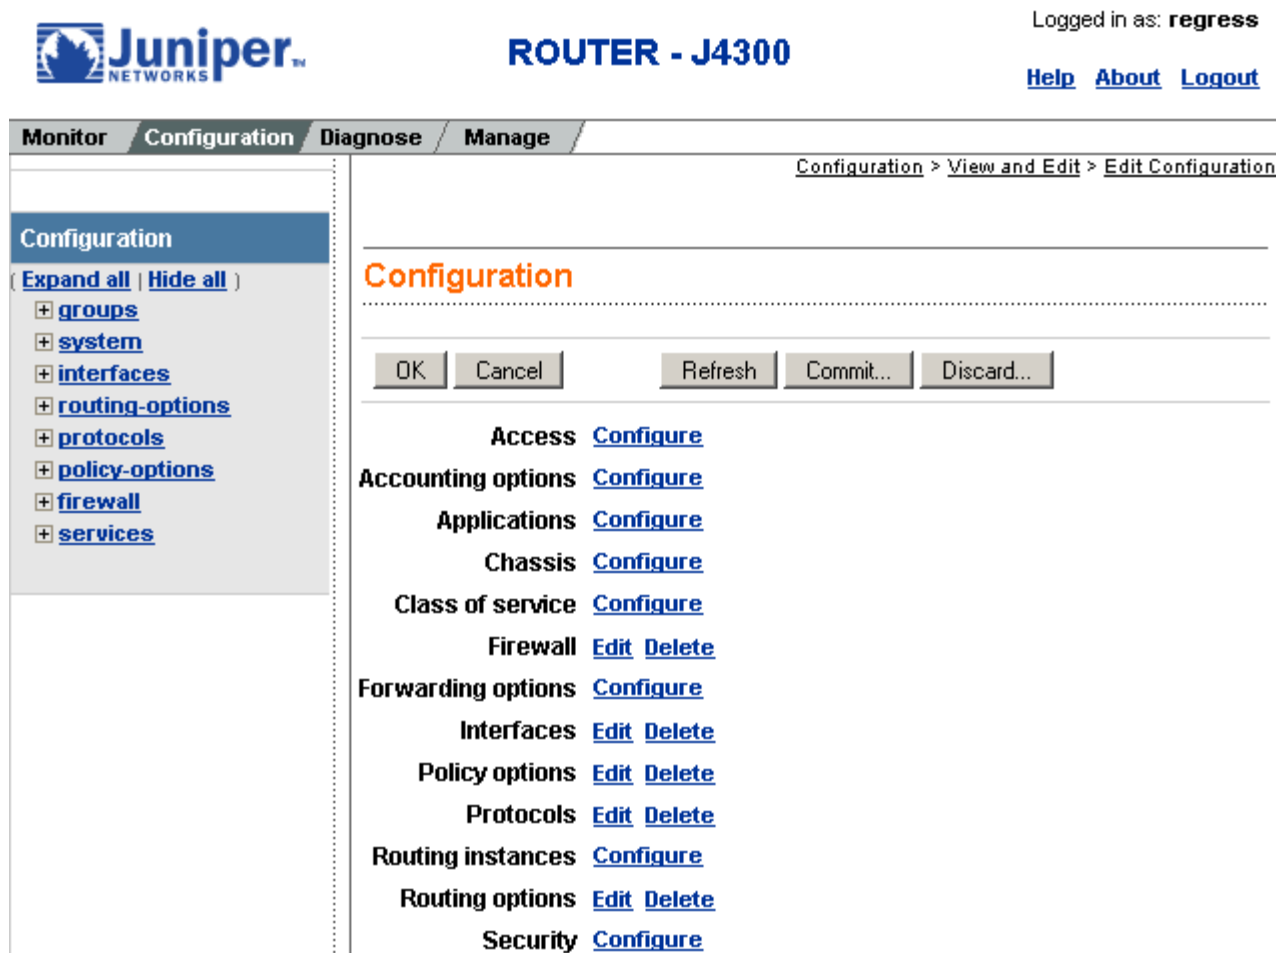
Editing and Committing the Clickable Configuration

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 8
- Discarding Parts of a Candidate Configuration on page 11
- Committing a Clickable Configuration on page 12

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 2).

Figure 2: 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 10 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 10: 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 11 describes the meaning of these icons.

Table 11: 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 29 and the *JUNOS System Basics Configuration Guide*.

After typing or selecting your configuration edits, click a button in the main pane (described in Table 12) 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 12: 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 12.
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 11.

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 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.

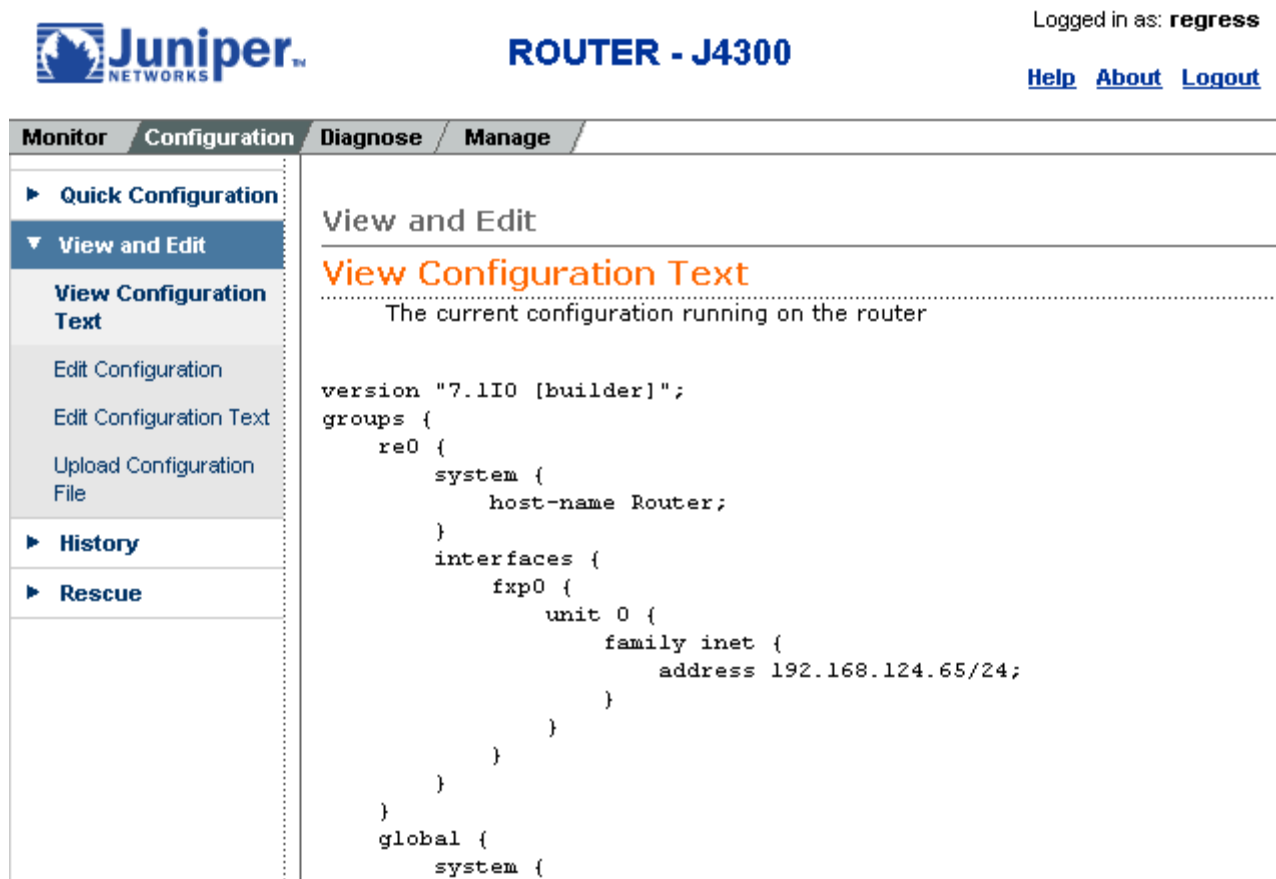
3. To display all the edits applied to the running configuration, click **Refresh**.

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 3).

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 3: View Configuration Text Page


Juniper NETWORKS **ROUTER - J4300** Logged in as: **regress**
[Help](#) [About](#) [Logout](#)

Monitor Configuration Diagnose Manage

► Quick Configuration
 ▼ View and Edit
 View Configuration Text
 Edit Configuration
 Edit Configuration Text
 Upload Configuration File
 ► History
 ► Rescue

View and Edit

View Configuration Text

The current configuration running on the router

```
version "7.1I0 [builder]";
groups {
  re0 {
    system {
      host-name Router;
    }
    interfaces {
      fxp0 {
        unit 0 {
          family inet {
            address 192.168.124.65/24;
          }
        }
      }
    }
  }
}
global {
  system {
```

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).

For more information about the format of an ASCII configuration file, see “Viewing the Configuration Text” on page 12.

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

The screenshot shows the Juniper J-Web Configuration Editor interface for a J4300 router. The top navigation bar includes the Juniper logo, the router model 'ROUTER - J4300', and a user login status 'Logged in as: regress'. Below this is a secondary navigation bar with tabs: Monitor, Configuration (selected), Diagnose, and Manage. On the left, a sidebar menu lists options: Quick Configuration, View and Edit (selected), View Configuration Text, Edit Configuration, Edit Configuration Text (highlighted), Upload Configuration File, History, and Rescue. The main content area is titled 'View and Edit' and 'Edit Configuration Text'. It contains a description: 'Edit the configuration. When you click "Commit", the edited configuration replaces the current configuration. If any errors occur when the configuration is loading or committed, they are displayed and the configuration is restored.' Below this, a 'Configuration' section shows a text editor with the following code:

```
version "7.1I0 [builder]";
groups {
  re0 {
    system {
      host-name Router;
    }
    interfaces {
      fxp0 {
        unit 0 {
          family inet {
            address 192.168.1.1;
          }
        }
      }
    }
  }
}
```

Uploading a Configuration File

To upload a configuration file from your local system:

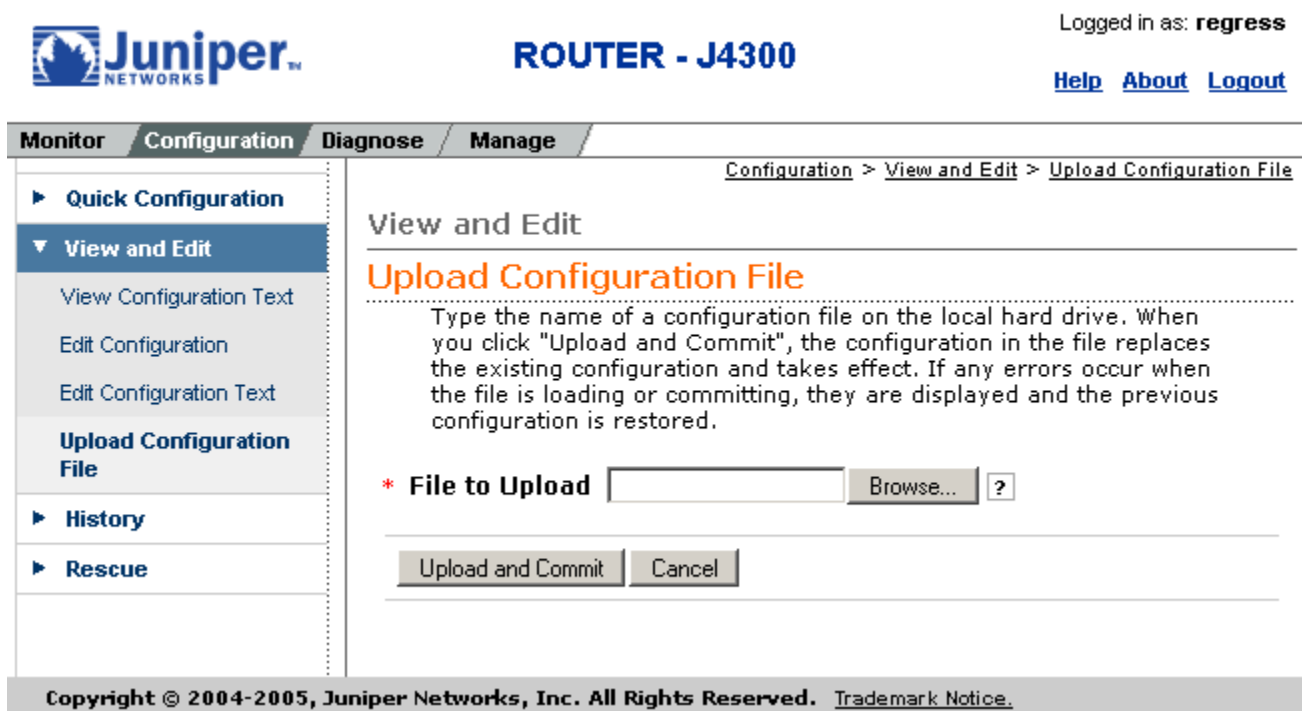
1. Select **Configuration > View and Edit > Upload Configuration File**.

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

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



Juniper NETWORKS

ROUTER - J4300

Logged in as: **regress**

[Help](#) [About](#) [Logout](#)

Monitor **Configuration** **Diagnose** **Manage**

[Configuration](#) > [View and Edit](#) > [Upload Configuration File](#)

View and Edit

Upload Configuration File

Type the name of a configuration file on the local hard drive. When you click "Upload and Commit", the configuration in the file replaces the existing configuration and takes effect. If any errors occur when the file is loading or committing, they are displayed and the previous configuration is restored.

* **File to Upload**

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. [Trademark Notice](#).

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 18
- 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).

Table 13 and Table 14 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 13: 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 14: 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.
Client	Method by which the configuration was committed: <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 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.

Table 14: J-Web Configuration History Summary (continued)

Field	Description
Log Message	<p>Method used to edit the configuration:</p> <ul style="list-style-type: none"> ■ Imported via <i>paste</i>—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 13. ■ Imported upload [<i>filename</i>]—Configuration was uploaded with the Configuration > View and Edit > Upload Configuration File option. For more information, see “Uploading a Configuration File” on page 14. ■ Modified via <i>quick-configuration</i>—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 <i>user-interface</i>—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). Table 13 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). Table 14 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):

- 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

[edit system]	[edit system]
	autoinstallation; radius-server { 10.10.10.10; }
[edit system tacplus-server]	[edit system tacplus-server]
	192.17.8.2;
[edit system tacplus-server]	[edit system tacplus-server]
10.7.7.9 secret "\$9\$l.le87-ds4JDbSz6A0hcbs2goG"; ## SECRET-DATA	
[edit]	[edit]
	chassis { alarm { ethernet { link-down yellow; } } }
[edit interfaces fe-0/0/0 unit 0 family inet]	[edit interfaces fe-0/0/0 unit 0 family inet]
service { input { service-set jweb-wan-sfw-service-set; } output { service-set jweb-wan-sfw-service-set; } }	
[edit interfaces fe-0/0/0 unit 0 family inet]	[edit interfaces fe-0/0/0 unit 0 family inet]
	address 192.168.124.75/24;

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). Table 14 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). Table 14 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** 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** 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** button. For more information, see “Disabling the CONFIG Button” on page 33.

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 24
- Modifying the Configuration on page 25
- Committing a Configuration with the CLI on page 30
- Disabling the CONFIG Button on page 33
- 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 6.

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 15) 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 System Basics Configuration Guide*.

Table 15: 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 <code>request system logout user</code> 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 fe-0/0/0
unit 0 {
  family inet {
    address 192.168.4.1/30;
  }
}
```

```
[edit]
user@host# edit interfaces fe-0/0/0
[edit interfaces fe-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 26
- Deleting a Statement or Identifier on page 26
- Copying a Statement on page 27
- Renaming an Identifier on page 27
- Inserting an Identifier on page 28
- Deactivating a Statement or Identifier on page 29

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
```

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 fe-0/0/0] hierarchy level:

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

[edit interfaces fe-0/0/0]
user@host# copy unit 0 to unit 1

[edit interfaces fe-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 27.

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 fe-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 `fe-0/0/0` at the `[edit interfaces]` hierarchy level:

```
[edit interfaces]
user@host# deactivate fe-0/0/0
```

```
[edit interfaces]
user@host# show
inactive: fe-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 30
- 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 31
- 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.
- *yyyy-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** button on the Services Router. For more information, see the *J-series Services Router Getting Started Guide*.)

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** 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** 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 Button

You can change the default behavior of the **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 button from loading the rescue configuration. The `no-clear` option prevents the CONFIG button from deleting all configurations on the router.

To return the function of the 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 34
- 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. (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 System Basics Configuration 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.
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 through Figure 10 show the results of override, replace, and merge operations.


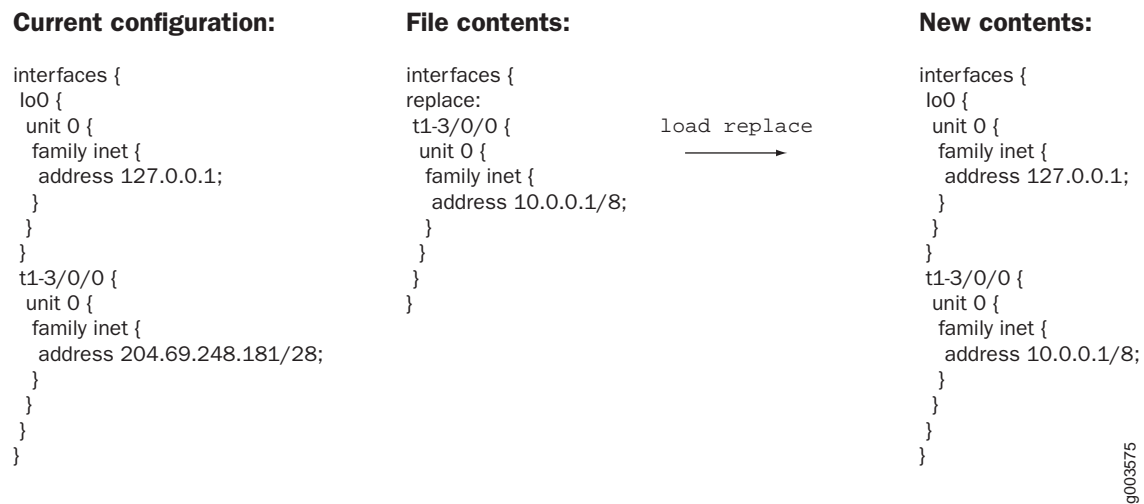
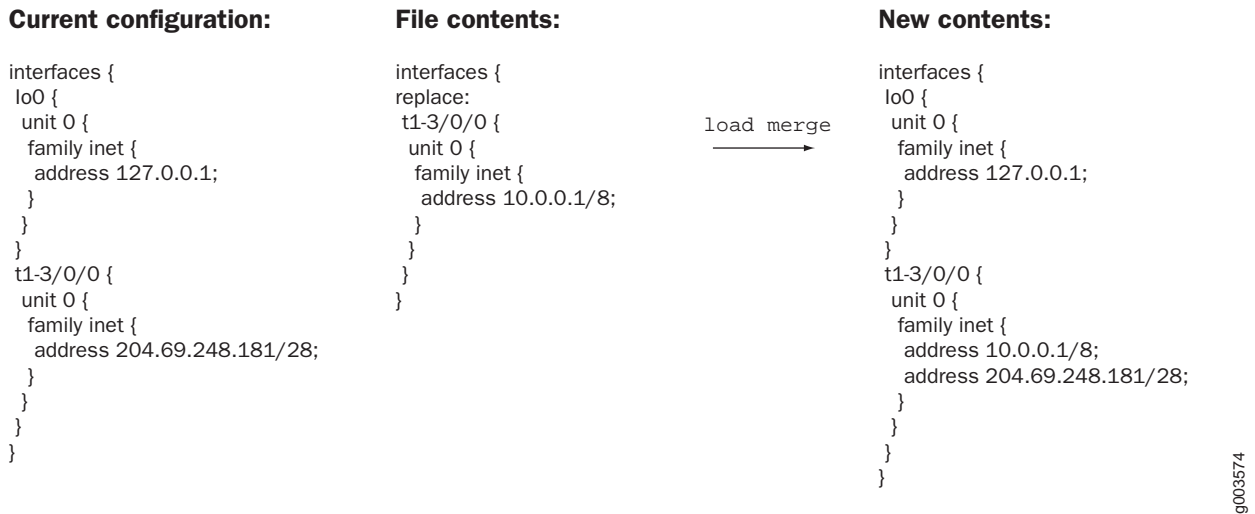
Figure 8: Loading a Configuration with the Override Operation**Figure 9: Loading a Configuration with the Replace Operation**

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 System Basics Configuration Guide*.

Part 2

Configuring Router Interfaces

- Interfaces Overview on page 41
- Configuring Network Interfaces on page 103
- Configuring Point-to-Point Protocol over Ethernet on page 151
- Configuring ISDN on page 169

Chapter 2

Interfaces Overview

J-series Services Routers support network interfaces for E1, E3, T1, T3, Fast 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 45
- Data Link Layer Overview on page 50
- Ethernet Interface Overview on page 52
- T1 and E1 Interfaces Overview on page 56
- T3 and E3 Interfaces Overview on page 60
- Serial Interface Overview on page 65
- ADSL Interface Overview on page 71
- SHDSL Interface Overview on page 74
- ISDN Interface Overview on page 74
- Interface Physical Properties on page 77
- Physical Encapsulation on an Interface on page 81
- Interface Logical Properties on page 88
- Special Interfaces on page 96

Interfaces Terms

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

Table 17: Network Interfaces Terms

Term	Definition
alternate mark inversion (AMI)	Original method of formatting T1 and E1 data streams.
asymmetrical 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	See <i>frame checksum sequence</i> .
channel service unit (CSU)	Unit that connects a digital telephone line to a multiplexer or other signal service.
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.

Table 17: Network Interfaces Terms (continued)

Term	Definition
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.
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.
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.
fractional E1	Service also called channelized E1, in which a 2.048-Mbps E1 link is subdivided into 32 DS0 time slots (channels) in which time slot 0 is reserved. The individual channels or groups of channels connect to different destinations, and customers pay for only the channels used and not for the entire line.
fractional T1	Service also called channelized T1, in which a 1.544-Mbps T1 link is subdivided into 24 DS0 time slots (channels) in which time slot 0 is reserved. The individual channels or groups of channels connect to different destinations, and customers pay for only the channels used and not for the entire line.
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.

Table 17: Network Interfaces Terms (continued)

Term	Definition
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. MLFR is often used in conjunction with Multilink Point-to-Point Protocol (MLPPP).
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.
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"> ■ Two Fast Ethernet LAN interfaces ■ Two T1 or two E1 WAN interfaces ■ Single E3 or T3 (DS3) WAN interface (J6300 model 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 (J4300 and J6300 models) ■ Single ISDN S/T or U interface (J2300 model) or four ISDN S/T or U interfaces (J4300 and J6300 models) ■ Two serial interfaces ■ Symmetric high-speed digital subscriber line (SHDSL) WAN interface—Annex A or Annex B to support SHDSL on J4300 and J6300 models.
Point-to-Point Protocol (PPP)	Link-layer protocol that provides multiprotocol encapsulation. PPP is used for link-layer and network-layer configuration.

Table 17: Network Interfaces Terms (continued)

Term	Definition
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 68. <p>For cable details, see the <i>J-series Services Router Getting Started Guide</i>.</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 asymmetrical 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.
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 46

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 asymmetrical digital subscriber line (ATM-over-ADSL) interface (J4300 and J6300 models only)



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 (J4300 and J6300 models only)



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

- E1 WAN interface
- E3 WAN interface (J6300 models only)
- Fast Ethernet LAN interface
- Integrated Services Digital Network (ISDN) BRI WAN interface
- Serial interface (EIA-530, RS-449/422, RS-232, V.35, and X.21 line protocols)
- T1 WAN interface
- T3 WAN interface (also called DS3) (J6300 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 the larger 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.

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>Network interface identifiers:</p> <ul style="list-style-type: none"> ■ at—ATM-over-ADSL or ATM-over-SHDSL WAN interface ■ bc—Bearer channel on an ISDN BRI ■ br—Basic Rate Interface for establishing ISDN connections ■ dc—Delta channel on an ISDN BRI ■ dl—Dialer interface for initiating ISDN connections ■ e1—E1 WAN interface ■ e3—E3 WAN interface ■ fe—Fast Ethernet LAN interface ■ se—Serial interface (either RS-232, RS-422/499, RS-530, V.35, or X.21) ■ t1—T1 (also called DS1) WAN interface ■ t3—T3 (also called DS3) WAN interface <p>Special interface identifiers: (See “Special Interfaces” on page 96.)</p> <ul style="list-style-type: none"> ■ dsc ■ gr, gre ■ ip, ipip ■ lo ■ ls ■ lsi ■ mtun ■ pd, pimd ■ pe, pime ■ sp ■ tap
<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. <p>The PIM number 0 is reserved for the out-of-band management ports. (See “Management Interface” on page 99.)</p>
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 T1 or E1 interface.	<ul style="list-style-type: none"> ■ On an E1 interface, a value from 0 through 32. The 0 and 1 time slots are reserved. ■ On a T1 interface, a value from 0 through 24. The 0 time slot is reserved.
<i>unit</i>	Number of the logical interface created on a physical interface.	<p>A value from 0 through 16384.</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 88.</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

```

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

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 52
- Collisions and Detection on page 53
- Collision Domains and LAN Segments on page 54
- Broadcast Domains on page 55
- Ethernet Frames on page 55

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 19 shows collision rounds up to round 10.

Table 19: 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}

Table 19: Collision Backoff Algorithm Rounds (continued)

Round	Size of Set	Elements in the Set
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 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 57
- T1 and E1 Signals on page 57
- 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][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.

T3 and E3 Interfaces Overview

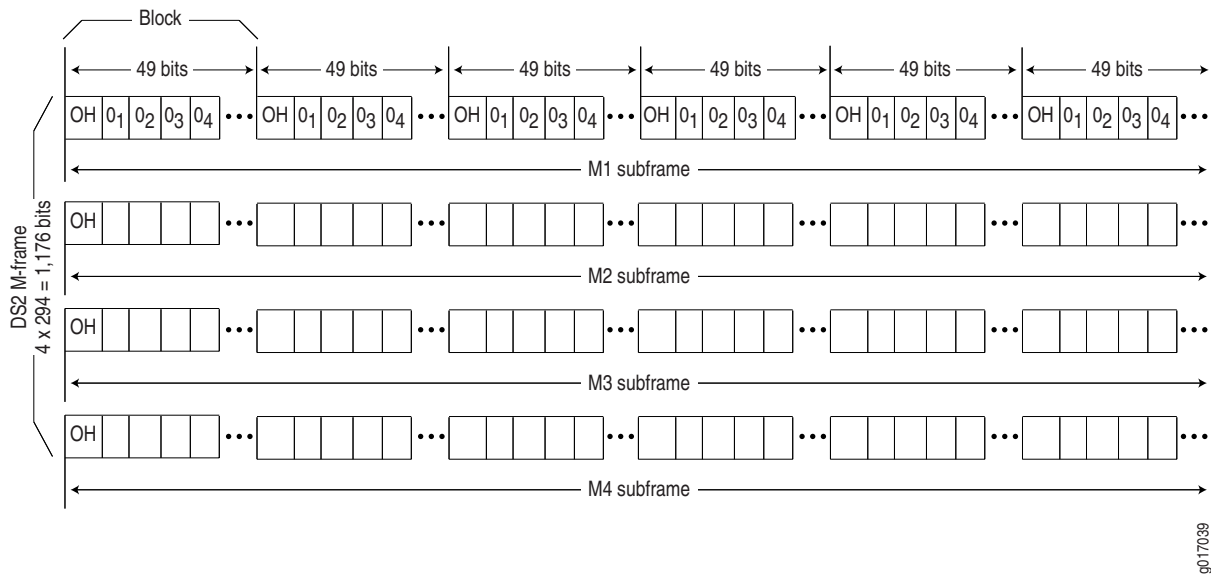
T3 is a high-speed data-transmission medium formed by multiplexing 28 DS1 signals into seven separate DS2 signals, and combining the DS2 signals into a single DS3 signal. T3 links operate at 43.736 Mbps. T3 is also called DS3.

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 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 and Figure 14.

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.

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, C₁₁, C₁₂, and C₁₃ 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.

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 1111111, in which x can be either 1 or 0. Bits are transmitted from right to left.

Table 20 lists some C-bit code words and the alarm or status condition indicated.

Table 20: FEAC C-Bit Condition Indicators

Alarm or Status Condition	C-Bit Code Word
DS3 equipment failure requires immediate attention.	00110010 11111111
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 21 lists and defines serial signals and their sources.

Table 21: 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.
- DTE clocking mode—Uses an internally generated clock. The speed of this clock is configured locally. DTE 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 shows the clock sources for loop, DCE, and DTE 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 69

- RS-232 on page 69
- RS-422/449 on page 70
- V.35 on page 70
- X.21 on page 71

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).

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.

J4300 and J6300 Services Routers 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.

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

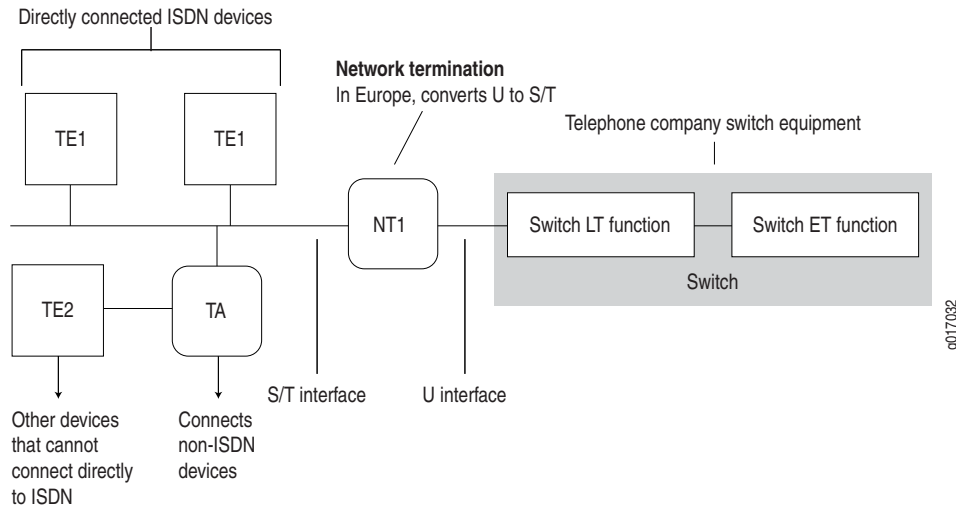
ISDN provides two basic types of service, Basic Rate Interface (BRI) and Primary Rate Interface (PRI). Services Routers support ISDN BRI only.

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.

Typical ISDN Network

Figure 17 shows a typical ISDN network.

Figure 17: ISDN Network



In Figure 17, 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. 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 22 summarizes some key physical properties of J-series Services Router interfaces.

Table 22: 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 78.
bert-period	Bit error rate test (BERT) time period over which bit errors are sampled. See “Bit Error Rate Testing” on page 78.
chap	Challenge Handshake Authentication Protocol (CHAP). Specifying chap enables CHAP authentication on the interface. See “CHAP Authentication” on page 84.
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 78.
description	A user-defined text description of the interface, often used to describe the interface’s purpose.
disable	Administratively disables the interface.

Table 22: Interface Physical Properties (continued)

Physical Property	Description
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 81.
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 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.
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.
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.

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 81
- Point-to-Point Protocol on page 83
- Point-to-Point Protocol over Ethernet on page 86
- High-Level Data Link Control on page 87

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 shows a typical Frame Relay network.

Figure 18: Frame Relay Network




Figure 18 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.

CHAP 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.

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.

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 89
- IPv4 Addressing on page 90
- IPv6 Addressing on page 93
- Virtual LANs on page 95

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.
- `vpls`—Virtual private LAN service (VPLS).

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 shows two subnets in a network.


Figure 19: Subnets in a Network

Figure 19 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 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 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 23 lists each special interface and briefly describes its use.

Table 23: Special Interfaces on a Services Router

Interface Name	Description
dsc	Discard interface. See “Discard Interface” on page 98.
fxp0	This interface is not supported on a J-series Services Router. (On an M-series or T-series router, fxp0 is used for out-of-band management.) For more information about the J-series Services Router management port interface, see “Management Interface” on page 99.
gr-0/0/0	Configurable generic routing encapsulation (GRE) interface. GRE allows the encapsulation of one routing protocol over another routing protocol. 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.
gre	Internally generated GRE interface. This interface is generated by the JUNOS software to handle GRE. It is not a configurable interface.

Table 23: 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 99.
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 100.</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>Configurable logical tunnel interface. The tunnel interface is used to provide services such as Layer 3 MPLS VPNs over GRE, IPsec over GRE, GRE over IPsec, PIM sparse mode multicast, multicast over Layer 3 VPNs, virtual private LAN service (VPLS), VPLS or Layer 2 VPNs terminated into Layer 3 VPNs, IPv6-over-IPv4 encapsulation, and logical routers.</p> <p>Within a Services Router, packets are routed to this internal interface for tunnel services. The logical 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 tunnel services.</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>
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.

Table 23: Special Interfaces on a Services Router (continued)

Interface Name	Description
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 151.</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.

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 interface (also called the out-of-band management interface) on a J-series Services Router can either be `fe-0/0/0` or `fe-0/0/1`. The management interface is a Fast Ethernet interface with a permanent port on the front of the router chassis.

The management interface is the primary interface for accessing the router remotely. Typically, the management interface is not connected to the in-band network, but is connected instead to the router's internal network. Through the 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 the 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

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.

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 Network Interfaces

Each Services Router can support types of interfaces that perform different functions. The router uses 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.

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 PPPoE interfaces, see “Configuring Point-to-Point Protocol over Ethernet” on page 151. To configure ISDN interfaces, see “Configuring ISDN” on page 169.

- Before You Begin on page 103
- Configuring Network Interfaces with Quick Configuration on page 104
- Configuring Network Interfaces with a Configuration Editor on page 126
- Verifying Interface Configuration on page 143

Before You Begin

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

- Install Services Router hardware. For more information, see the *J-series Services Router Getting Started Guide*.
- Establish basic connectivity. For more information, see the *J-series 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.

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.

Configuring Network Interfaces with Quick Configuration

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

Figure 22: Quick Configuration Interfaces Page

Logged in as: **regress**
[Help](#) [About](#) [Logout](#)

Router - J4300

Monitor Configuration Diagnose Manage

Quick Configuration

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

Quick Configuration

Interfaces

Interface Name	Link State	Configured	Description
fe-0/0/0	Down	Yes	Fast Ethernet Interface 'fe-0/0/0'
fe-0/0/0.0	Down	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'
fxp0	Up	Yes	Management Interface 'fxp0'
fxp0.0	Up	Yes	Logical Unit 0 on Management Interface 'fxp0'
lo0	Up	Yes	Loopback Interface 'lo0'
lo0.0	Up	Yes	Logical Unit 0 on Loopback Interface 'lo0'

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. [Trademark Notice.](#)

To configure a network interface with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Quick Configuration > Interfaces**. You can select **Interfaces** in the list under Router Configuration or from the left pane.

A list of the network interfaces present on the Services Router is displayed, as shown in Figure 22. The third column indicates whether the interface has been configured.

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 108
- Configuring a Fast Ethernet Interface with Quick Configuration on page 113
- Configuring a T1 Interface with Quick Configuration on page 115
- Configuring a T3 Interface with Quick Configuration on page 119
- Configuring a Serial Interface with Quick Configuration on page 122

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, select the E1 interface you want to configure.

The properties you can configure on an E1 interface are displayed, as shown in Figure 23.

Figure 23: E1 Interfaces Quick Configuration Page

The screenshot shows the Juniper J6300 Router configuration interface. The top navigation bar includes the Juniper logo, the router model "ROUTER - J6300", and the user "regress" is logged in. There are links for "Help", "About", and "Logout". Below the navigation bar, the "Configuration" tab is selected, and the breadcrumb trail is "Configuration > Quick Configuration > Interfaces".

The left sidebar contains a "Quick Configuration" menu with options: "Set Up", "SSL", "Interfaces" (selected), "Users", "SNMP", "Routing", "Firewall/NAT", "IPSec Tunnels", "Realtime Performance Monitoring", "View and Edit", "History", and "Rescue".

The main content area is titled "Quick Configuration" and "Interfaces". It shows the "Physical Interface: 'e1-1/0/0'". Under "Logical Interfaces", it states "No logical interfaces configured." with an "Add..." button. The "Physical Interface Description" field is empty. The "Encapsulation" section has a dropdown menu and an "Enable CHAP" checkbox. The "CHAP Local Identity" section includes a "Use System Host Name" checkbox (checked), a "Local Name" field, and fields for "CHAP Peer Identity" and "CHAP Secret", each preceded by a red asterisk.

2. Enter information into the Quick Configuration page, as described in Table 24.
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 143.

Table 24: 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.	<p>From the list, select one of the following:</p> <ul style="list-style-type: none"> ■ internal—Services Router's own system clock (the default) ■ external—Clock received from the E1 interface
Encapsulation		
Encapsulation	Specifies the encapsulation type for traffic on the interface.	<p>From the list, select the encapsulation for this E1 interface:</p> <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.

Table 24: E1 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 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, select the interface you want to configure.

The properties you can configure on an E3 interface are displayed, as shown in Figure 24.

Figure 24: E3 Interfaces Quick Configuration Page

Juniper NETWORKS ROUTER - J6300

Logged in as: regress [Help](#) [About](#) [Logout](#)

Monitor Configuration Diagnose Manage Alarms

Quick Configuration

Set Up
Secure Access
Interfaces
Users
SNMP
Routing
Firewall/NAT
DHCP
IPSec Tunnels
Realtime Performance Monitoring

View and Edit
History
Rescue

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

Quick Configuration

Interfaces Physical Interface: 'e3-3/0/0'

Logical Interfaces

	Logical Interface Name	Link State	Configured	Description
<input type="checkbox"/>	e3-3/0/0.0	Up	Yes	Logical Unit 0 on E3 Interface 'e3-3/0/0'

[Add...](#) [Delete](#)

Physical Interface Description

MTU (bytes) ?

Clocking (internal) ?

Encapsulation

Encapsulation

Enable CHAP ☐

- Enter information into the Quick Configuration page, as described in Table 25.
- 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**.
- To verify that the E3 interface is configured correctly, see “Verifying Interface Configuration” on page 143.

Table 25: 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.	<p>From the list, select one of the following:</p> <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.	<p>From the list, select the encapsulation for this E3 interface:</p> <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 25: 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.

Table 25: E3 Quick Configuration Summary (continued)

Field	Function	Your Action
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>
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.	<p>From the Idle Cycle Flag list, select one of the following:</p> <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>	<p>From the Loopback list, select one of the following:</p> <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.

Table 25: E3 Quick Configuration Summary (continued)

Field	Function	Your Action
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.	<p>From the Start End Flag list, select one of the following:</p> <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, select the interface you want to configure.

The properties you can configure on a Fast Ethernet interface are displayed, as shown in Figure 25.

Figure 25: Fast Ethernet Interfaces Quick Configuration Page

The screenshot shows the Juniper J4300 Router configuration interface. The top navigation bar includes 'Monitor', 'Configuration', 'Diagnose', and 'Manage'. The 'Configuration' tab is active, and the left sidebar shows 'Quick Configuration' expanded with options like 'Set Up', 'SSL', 'Interfaces', 'Users', 'SNMP', 'Routing', 'Firewall/NAT', 'IPSec Tunnels', and 'Realtime Performance Monitoring'. The 'Interfaces' section is selected. The main content area displays 'Quick Configuration' for 'Physical Interface: fe-0/0/0/0'. It includes a table for 'Logical Interfaces' with columns for 'Logical Interface Name', 'Link State', 'Configured', and 'Description'. The table shows one entry for 'fe-0/0/0.0' with a 'Down' link state and 'Yes' configured. Below the table are 'Add...' and 'Delete' buttons. A 'Physical Interface Description' field is present with 'OK', 'Cancel', and 'Apply' buttons at the bottom.

Juniper NETWORKS **ROUTER - J4300** Logged in as: **regress** [Help](#) [About](#) [Logout](#)

Monitor Configuration Diagnose Manage

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

Quick Configuration

Interfaces Physical Interface: 'fe-0/0/0/0'

Logical Interfaces

	Logical Interface Name	Link State	Configured	Description
<input type="checkbox"/>	fe-0/0/0.0	Down	Yes	Logical Unit 0 on Fast Ethernet Interface 'fe-0/0/0/0'

[Add...](#) [Delete](#)

Physical Interface Description

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

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. [Trademark Notice.](#)

2. Enter information into the Quick Configuration page, as described in Table 26.
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 143.

Table 26: Fast Ethernet Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
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.
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.	Type a value between 256 and 9192 bytes. The default MTU for Fast Ethernet interfaces is 1504 .

Configuring a T1 Interface with Quick Configuration

To configure properties on a T1 interface:

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

The properties you can configure on a T1 interface are displayed, as shown in Figure 26.

Figure 26: T1 Interfaces Quick Configuration Page

Juniper NETWORKS **ROUTER - J6300** Logged in as: **regress**
[Help](#) [About](#) [Logout](#)

Monitor **Configuration** **Diagnose** **Manage**

Quick Configuration

- Set Up
- SSL
- Interfaces**
- Users
- SNMP
- Routing
- Firewall/NAT
- IPSec Tunnels
- Realtime Performance Monitoring

► **View and Edit**

► **History**

► **Rescue**

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

Quick Configuration

Interfaces **Physical Interface: 't1-6/0/1'**

Logical Interfaces

No logical interfaces configured.

Physical Interface Description

Encapsulation

Encapsulation

Enable CHAP ☐

CHAP Local Identity

Use System Host Name ☒

Local Name

*** CHAP Peer Identity**

*** CHAP Secret**

- Enter information into the Quick Configuration page, as described in Table 27.
- 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**.
- To verify that the T1 interface is configured correctly, see “Verifying Interface Configuration” on page 143.

Table 27: 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.	<p>From the list, select one of the following:</p> <ul style="list-style-type: none"> ■ internal—Services Router's own system clock (the default) ■ external—Clock received from the T1 interface
Encapsulation		
Encapsulation	Specifies the encapsulation type for traffic on the interface.	<p>From the list, select the encapsulation for this T1 interface:</p> <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.

Table 27: T1 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 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 27: 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 a T3 Interface with Quick Configuration

To configure properties on a T3 (DS3) interface:

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

The properties you can configure on a T3 interface are displayed, as shown in Figure 27.

Figure 27: T3 Interfaces Quick Configuration Page

Juniper NETWORKS **ROUTER - J6300** Logged in as: **regress**
[Help](#) [About](#) [Logout](#)

Monitor **Configuration** **Diagnose** **Manage**

Quick Configuration

- Set Up
- SSL
- Interfaces**
- Users
- SNMP
- Routing
- Firewall/NAT
- IPSec Tunnels
- Realtime Performance Monitoring

► **View and Edit**

► **History**

► **Rescue**

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

Quick Configuration

Interfaces **Physical Interface: 't3-4/0/0'**

Logical Interfaces

No logical interfaces configured.

Physical Interface Description

Encapsulation

Encapsulation

Enable CHAP ☐

CHAP Local Identity

Use System Host Name ☒

Local Name

*** CHAP Peer Identity**

*** CHAP Secret**

- Enter information into the Quick Configuration page, as described in Table 28.
- 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**.
- To verify that the T3 interface is configured correctly, see “Verifying Interface Configuration” on page 143.

Table 28: 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.	<p>From the list, select one of the following:</p> <ul style="list-style-type: none"> ■ internal—Services Router's own system clock (the default) ■ external—Clock received from the T3 interface
Encapsulation		
Encapsulation	Specifies the encapsulation type for traffic on the interface.	<p>From the list, select the encapsulation for this T3 interface:</p> <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.

Table 28: T3 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 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).	<div>■ To enable long buildout, select the check box.</div> <div>■ To disable long buildout, clear the check box.</div>
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.	<div>■ To disable, select the check box.</div> <div>■ To enable, clear the check box.</div>

Configuring a Serial Interface 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, select the interface you want to configure.

The properties you can configure on a serial interface are displayed, as shown in Figure 28.

Figure 28: Serial Interfaces Quick Configuration Page

The screenshot shows the Juniper J6300 Router configuration interface. The top navigation bar includes 'Monitor', 'Configuration', 'Diagnose', and 'Manage'. The 'Configuration' tab is active, and the breadcrumb trail is 'Configuration > Quick Configuration > Interfaces'. The left sidebar shows a tree view with 'Quick Configuration' expanded, containing 'Set Up', 'SSL', 'Interfaces' (selected), 'Users', 'SNMP', 'Routing', 'Firewall/NAT', 'IPSec Tunnels', and 'Realtime Performance Monitoring'. Below this are 'View and Edit', 'History', and 'Rescue' buttons. The main content area is titled 'Quick Configuration' and 'Interfaces'. It shows 'Physical Interface: 'se-5/0/0'' and 'Logical Interfaces' with a message 'No logical interfaces configured.' and an 'Add...' button. The 'Physical Interface Description' field is empty. The 'Encapsulation' section has a dropdown menu and an 'Enable CHAP' checkbox. The 'CHAP Local Identity' section includes a checked 'Use System Host Name' checkbox, and fields for 'Local Name', '* CHAP Peer Identity', and '* CHAP Secret'.

2. Enter information into the Quick Configuration page, as described in Table 29.
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 143.

Table 29: 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"> Type one or more IPv4 addresses and prefixes. For example: 10.10.10.10/24 Click Add. 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.
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)		
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		

Table 29: Serial Quick Configuration Summary (continued)

Field	Function	Your Action
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 [edit interfaces se-pim / 0 / port serial-options] 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>
Clock Rate	<p>Specifies the line speed in kilohertz or megahertz for serial interfaces that use the DTE clocking mode.</p>	<p>From the list, select one of the following clock rates:</p> <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 Network 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 126
- Adding an ATM-over-ADSL Network Interface with a Configuration Editor on page 128
- Configuring CHAP on the ATM-over-ADSL Interface (Optional) on page 133
- Adding an ATM-over-SHDSL Interface on page 135
- Configuring Compressed Real-Time Transport Protocol (CRTP) on page 140
- Deleting a Network Interface with a Configuration Editor on page 142

For information about using the J-Web and CLI configuration editors, see “Using J-Web Configuration Tools” on page 3.

Adding a Network Interface with a Configuration Editor

To configure network interfaces 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 30.
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 143.

Table 30: 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 configuration editor hierarchy, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Configure or Edit. 	<p>From the top of the configuration hierarchy, create and name the interface:</p> <pre>edit interfaces interface-name</pre>
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 46. 3. Click OK. 	
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. 	<p>Enter values for physical interface properties as needed. Examples include changes to the default values for physical encapsulation or MTU. For example:</p> <pre>set interface-name encapsulation ppp</pre>

Table 30: Adding an Interface (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
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"> Under Nested configuration, click Configure for the appropriate interface type. In the interface-specific page that appears, enter the values you need to supply or change the default values. 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"> From the [edit interfaces <i>interface-name</i>] hierarchy level, type edit interface-options Enter the statement for each interface-specific property for which you need to change the default value.
Add logical interfaces.	<ol style="list-style-type: none"> In the main Interface page for this interface, next to Unit, click Add new entry. On the Unit page for logical interfaces that appears, type a number from 0 through 16384 in the Interface unit number box. Enter values in other fields as required for your network. To configure protocol family values if needed, under Family, click Configure next to the appropriate protocol. To access additional subordinate hierarchies under Nested configuration, click Configure next to any parameter you want to configure. When you are finished, click OK. 	<ol style="list-style-type: none"> From the [edit interfaces <i>interface-name</i>] hierarchy level, type set unit logical-unit-number Replace <i>logical-unit-number</i> with a value from 0 through 16384. Enter additional values for properties you need to configure on the logical interface, such as logical encapsulation or protocol family.

Adding an ATM-over-ADSL Network Interface with a Configuration Editor

J4300 and J6300 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). ATM-over-ADSL interfaces are not currently supported on J2300 Services Routers.



NOTE: You can configure J4300 and J6300 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`. Multiple encapsulation types are supported on both the physical and logical ATM-over-ADSL interface.

To configure ATM-over-ADSL network interfaces 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 31.
3. Go on to one of the following procedures:
 - To enable authentication on the interface, see “Configuring CHAP on the ATM-over-ADSL Interface (Optional)” on page 133.
 - 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 151.

Table 31: 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 configuration editor hierarchy, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	<p>From the top of the configuration hierarchy, create and name the interface:</p> <p>edit interfaces at-2/0/0</p>
Create the new interface—for example, <code>at-2/0/0</code> .	<ol style="list-style-type: none"> 1. Next to Interface, click Add new entry. 2. In the Interface name box, type <code>at-2/0/0</code>. 3. Click OK. 	
Configuring Physical Properties		

Table 31: 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.	1. Next to Atm options, click Configure .	1. To configure the VPI value, enter
<ul style="list-style-type: none"> ■ ATM VPI—A number between 0 and 255—for example, 25. 	2. Next to Vpi , click Add new entry .	<code>set atm-options vpi 25</code>
<ul style="list-style-type: none"> ■ 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. 	3. In the Vpi number box, type 25.	2. To configure OAM liveness values on a VPI, enter
	4. In the Actions box, click Edit .	<code>set atm-options vpi 25 oam-liveness up-count 200 down-count 200</code>
	5. Next to Oam liveness , click Configure .	3. To configure the OAM period, enter
	6. In the Down count box, type 200.	<code>set atm-options vpi 25 oam-period 100</code>
	7. In the Up count box, type 200.	
<ul style="list-style-type: none"> ■ 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. 	8. Click OK .	
	9. Next to Oam period box, click Configure .	
	10. From the Oam period choices list, select Oam period .	
	11. In the Oam period box, type 100.	
	12. Click OK .	

Table 31: 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
<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. 	3. Click OK .	
Annex A supports the following operating modes:		
<ul style="list-style-type: none"> ■ 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. 		
Annex B supports the following operating modes:		
<ul style="list-style-type: none"> ■ 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 .	1. 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. <p>For PPP over ATM (PPPoA) over ADSL interfaces, use this type of encapsulation.</p>	2. Click OK .	set encapsulation ethernet-over-atm
<ul style="list-style-type: none"> ■ ethernet-over-atm—Ethernet over ATM encapsulation. <p>For PPP over Ethernet (PPPoE) over ATM-over-ADSL interfaces that carry IPv4 traffic, use this type of encapsulation.</p>		

Table 31: Adding an ATM-over-ADSL Network Interface (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configuring Logical Properties		
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 143.)		
<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 31: Adding an ATM-over-ADSL Network Interface (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Add the Family protocol type—for example, <code>inet</code> .	Select the protocol type <code>inet</code> and then click Configure .	Enter <code>set unit 3 family inet</code> Commands vary depending on the protocol type.
Configure ATM virtual channel identifier (VCI) options for the interface.	<ol style="list-style-type: none"> From the Vci Type list, select vci. In the Vci box, type 35. 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 the VCI value, enter <code>set unit 3 vci 35</code> To configure OAM liveness values on a VCI, enter <code>set unit 3 vci 35 oam-liveness up-count 200 down-count 200</code> To configure the OAM period, enter <code>set unit 3 vci 35 oam-period 100</code>
<ul style="list-style-type: none"> ■ ATM VCI type—vci. ■ ATM VCI value—A number between 0 and 4089—for example, 35, with VCI's 0 through 31 reserved. ■ 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. 		

Configuring CHAP on the ATM-over-ADSL Interface (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 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 32.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying Interface Configuration” on page 143.

Table 32: Configuring CHAP

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 configuration hierarchy, select Configuration > Edit Configuration > View and Edit. 2. Next to Access, click Configure or Edit. 	<p>From the top of the configuration hierarchy, enter</p> <pre>set access profile A-ppp-client client client1 chap-secret my-secret</pre>
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 Configuration page. 	
Navigate to the at-3/0/0 unit 0 interface level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration hierarchy, select Interfaces. 2. In the Interface name box, click at-3/0/0. 3. In the Interface unit number box, click 0. 	<p>From the top of the configuration hierarchy, enter</p> <pre>edit interfaces at-3/0/0 unit 0</pre>
Configure CHAP on the ATM-over-ADSL 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. 	<p>Enter</p> <pre>set ppp-options chap access-profile A-ppp-client</pre>
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	<p>Enter</p> <pre>set ppp-options chap local-name A-at-3/0/0.0.</pre>
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> <pre>set ppp-options chap passive</pre>

Adding an ATM-over-SHDSL Interface

J4300 and J6300 Services Routers with G.SHDSL 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). ATM-over-SHDSL interfaces are not currently supported on J2300 Services Routers.



NOTE: You can configure J4300 and J6300 Services Routers with a G.SHDSL PIM for connections through SHDSL only, not for direct ATM connections.

The J-series Services Routers, with a 2-port G.SHDSL PIM installed, support the following modes. You can configure only one mode on each PIM.

- 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. In two-wire mode, the PIM has 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. In four-wire mode, the PIM has 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/0/port`. Multiple encapsulation types are supported on both the physical and logical ATM-over-SHDSL interface.

To configure ATM-over-SHDSL network interfaces 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 33.
3. Go on to configuring PPP over Ethernet (PPPoE) encapsulation on an ATM-over-SHDSL interface. See “Configuring Point-to-Point Protocol over Ethernet” on page 151.

Table 33: 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 configuration editor hierarchy, select Configuration > View and Edit > Edit Configuration. 2. Next to Chassis, click Configure. 	<p>From the top of the configuration hierarchy, add the interface to the chassis:</p> <pre>set chassis fpc 6 pic 0 shdsl pic-mode 4-port-atm</pre>
<p>Set the ATM-over-SHDSL mode on the G.SHDSL PIM, if required. By default, the G.SHDSL PIM is enabled in 2-wire Annex B mode. To configure the 4-wire mode on the G.SHDSL PIM, follow the tasks in this table.</p> <p>The configuration editor uses standard JUNOS names for Services Router hardware and interfaces. “PIC” indicates the PIM. The FPC value identifier identifies the chassis slot, from 1 through 6, in which the PIM is installed. The slot value is always 0 on J-series Services Routers. See “Interfaces Overview” on page 41.</p>	<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 4-port-atm. 7. Click OK. 	
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	<p>From the top of the configuration hierarchy, create and name the interface:</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 33: 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.</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. Next to Atm options, click Configure. 2. Next to Vpi, click Add new entry. 3. In the Vpi number box, type 25. 4. In the Actions box, click Edit. 5. Next to Oam liveness, click Configure. 6. In the Down count box, type 200. 7. In the Up count box, type 200. 8. Click OK. 9. Next to Oam period, click Configure. 10. From the Oam period choices list, select Oam period. 11. In the Oam period box, type 100. 12. Click OK. 	<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-ADSL 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. 	<ol style="list-style-type: none"> 1. From the Encapsulation list, select ethernet-over-atm. 2. Click OK. 	<p>Enter 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 set shdsl-options annex annex-a</p>

Table 33: 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 on two-wire G.SHDSL PIMs 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 G.SHDSL PIM has a default line rate value of 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 to between –10 dB and 10 dB—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. ■ 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"> 1. Next to Snr margin, select Yes, then click Configure. 2. From the Current list, select Enter Specific Value. 3. In the Value box, type 5. 4. From the Snext list, select Enter Specific Value. 5. In the Value box, type 5. 6. Click OK until you return to the Interfaces page. 	<ol style="list-style-type: none"> 1. Enter 2. Enter <p>set shdsl-options snr-margin current 5</p> <p>set shdsl-options snr-margin snext disabled</p>
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"> 1. Scroll down the page to Unit, and click Add new entry. 2. In the Interface unit number box, type 3. 3. Enter other values in the fields required by your network. 	<p>Enter</p> <p>set unit 3</p>

Table 33: Adding an ATM-over-SHDSL Network Interface (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure encapsulation for the ATM-for-SHDSL logical unit—for example, atm-nlpid .	From the Encapsulation list, select atm-nlpid .	Enter
The following encapsulations are supported on the ATM-over-SHDSL 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-SHDSL for PPP-over-ATM (PPPoA) interfaces only. (For a sample PPPoA configuration, see “Verifying Interface Configuration” on page 143.)		
<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-SHDSL 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 33: Adding an ATM-over-SHDSL Network Interface (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Add the Family protocol type—for example, <code>inet</code> .	Select <code>inet</code> , and then click Configure .	Enter <code>set unit 3 family inet</code> Commands vary depending on the protocol type.
Configure ATM virtual channel identifier (VCI) options for the interface.	<ol style="list-style-type: none"> From the Vci type list, select vci. In the Vci box, type 35. 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 box, click Configure. From the Oam period list, select Oam period. In the Oam period box, type 100. Click OK. 	<ol style="list-style-type: none"> To configure the VCI value, enter <code>set unit 3 vci 35</code> To configure OAM liveness values on a VCI, enter <code>set unit 3 vci 35 oam-liveness up-count 200 down-count 200</code> To configure the OAM period, enter <code>set unit 3 vci 35 oam-period 100</code>
<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 VCI's 0 through 31 reserved. 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. 		

Configuring Compressed Real-Time Transport Protocol (CRTP)

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 12-byte RTP header, the IP and UDP header, can be too large a payload on networks using low-speed lines such as dial-up modems. Compressed Real-Time Transport Protocol (CRTP) can be configured on a single link to reduce network overhead on low-speed links.

On the Services Router, CRTP can be configured on a T1 or E1 interface with PPP encapsulation and using the link services interface as a compression device.

To configure CRTP on the Services Router:

- 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 34.
3. If you are finished configuring the router, commit the configuration.

Table 34: Adding CRTP to an E1 or T1 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 configuration editor hierarchy, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	<p>From the top of the configuration hierarchy, create and name the interface:</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. 	<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"> 2. From the Encapsulation list, select ppp as the encapsulation type. 3. Under Unit, click Edit. 	<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. 	<ol style="list-style-type: none"> 1. Enter <pre>set compression-device ls-0/0/0.0</pre> 2. Enter <pre>exit</pre> <p>until you return to the edit interfaces hierarchy.</p>

Table 34: Adding CRTP to an E1 or T1 Interface (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Add the link services interface, ls-0/0/0, to the Services Router.	<ol style="list-style-type: none"> Next to Interface, click Add new entry. In the Interface name box, type ls-0/0/0. Click OK to return to the Interfaces page. On the main Interface page, next to ls-0/0/0, click Edit. Next to Unit, click Add new entry. In the Interface unit number box, type 0. 	<p>From the [edit interfaces] hierarchy level, enter</p> <p>edit interfaces ls-0/0/0 unit 0</p>
<p>Configure the link services interface, ls-0/0/0, properties.</p> <p>F-max period —Maximum number of compressed packets allowed between transmission of full headers. It has a range from 1 to 65535.</p> <p>Maximum and Minimum—UDP port values from 1 to 65536 to reserve these ports for RTP compression. CRTP is applied to network traffic on ports within this range. This is only applicable to voice services interfaces.</p>	<ol style="list-style-type: none"> Next to Compression, select yes, and then click Configure. Select RTP, and then click Configure. In the F-Max period box, type 2500. Select Port, then click Configure. In the Minimum value box, type 2000. In the Maximum value box, type 64009. Click OK. 	<p>Enter</p> <p>set compression rtp f-max-period 2500 port maximum 64009 minimum 2000</p>

For more information about configuring CRTP on a single link, see the *JUNOS Network Interfaces Configuration Guide* and the *JUNOS Services Interfaces Configuration Guide*.

Deleting a Network Interface with a Configuration Editor

To delete an interface on a Services Router:

- Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- Perform the configuration tasks described in Table 35.



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 35: 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 configuration editor hierarchy, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	From the top of the configuration hierarchy, enter edit interfaces
Select the interface you want to delete.	In the Interface table, under Interface name, select the name of the interface you want to delete.	Enter delete <i>interface-name</i>
Execute the selection.	<ol style="list-style-type: none"> 1. Click Discard. 2. In the page that appears, select the appropriate radio button. If you have not made any previous changes, the only selection available is Delete Configuration Below This Point. 	Commit the configuration change: commit

Verifying Interface Configuration

To verify an interface configuration, perform these tasks:

- Verifying the Link State of All Interfaces on page 143
- Verifying Interface Properties on page 144
- Verifying ADSL Interface Properties on page 145
- Displaying a PPPoA Configuration for an ATM-over-ADSL Interface on page 149

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.

Sample Output

```
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*.

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 Protocols, Class of Service, and System Basics Command Reference*.

Verifying Interface Properties

Purpose Verify that the interface properties are correct.

Action From the CLI, enter the show interfaces detail command.

Sample Output user@host> **show interfaces detail**

```
Physical interface: fe-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.

For more information about `show interfaces detail`, see the *JUNOS Interfaces Command Reference*.

Verifying ADSL Interface Properties

- Purpose** Verify that the interface properties are correct.
- Action** From the CLI, enter the `show interfaces interface-name extensive` command.
- Sample Output**
- ```
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
 Attenuation (dB) : 0.5
 Capacity used (%) : 81
 Noise margin (dB) : 9.0
 Output power (dBm) : 7.5
 ATU-C
 Attenuation (dB) : 0.0
 Capacity used (%) : 72
 Noise margin (dB) : 9.5
 Output power (dBm) : 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 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.

For more information about `show interfaces` extensive, 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.

**Sample Output**

```
[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 "9G4ikPu0ISyKP5cIKv7Nik.PT3"; ## SECRET-DATA
}
```

**What It Means** Verify that the output shows the intended configuration of PPPoA. For more information about the format of a configuration file, see “Viewing the Configuration Text” on page 12.



## Chapter 4

# 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:** J4300 and J6300 Services Routers with asymmetrical digital subscriber line (ADSL) Physical Interface Modules (PIMs) can use PPPoE over Asynchronous Transfer Mode (ATM) to connect through DSL lines only, not for direct ATM connections.

You can use either the J-Web configuration editor or CLI configuration editor to configure PPPoE.

This chapter contains the following topics:

- PPPoE Terms on page 151
- PPPoE Overview on page 152
- Before You Begin on page 155
- Configuring PPPoE with a Configuration Editor on page 155
- Verifying a PPPoE Configuration on page 162

## PPPoE Terms

Before configuring PPPoE on a Services Router, become familiar with the terms defined in Table 36.

**Table 36: PPPoE Terms**

| <b>Term</b>                                                      | <b>Definition</b>                                                                                                                                                                                                                                          |
|------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>access concentrator</b>                                       | Router that acts as a server in a PPPoE session—for example, an E-series router.                                                                                                                                                                           |
| <b>customer premises equipment (CPE)</b>                         | Router that acts as a PPPoE client in a PPPoE session—for example, a Services Router.                                                                                                                                                                      |
| <b>Logical Link Control (LLC)</b>                                | Encapsulation protocol that allows transport of multiple protocols over a single ATM virtual connection.                                                                                                                                                   |
| <b>Point-to-Point Protocol (PPP)</b>                             | Encapsulation protocol for transporting IP traffic over point-to-point links.                                                                                                                                                                              |
| <b>PPP over Ethernet (PPPoE)</b>                                 | Network protocol that encapsulates PPP frames in Ethernet frames and connects multiple hosts over a simple bridging access device to a remote access concentrator.                                                                                         |
| <b>PPPoE Active Discovery Initiation (PADI) packet</b>           | Initiation packet that is broadcast by the client to start the discovery process.                                                                                                                                                                          |
| <b>PPPoE Active Discovery Offer (PADO) packet</b>                | Offer packets sent to the client by one or more access concentrators in reply to a PADI packet.                                                                                                                                                            |
| <b>PPPoE Active Discovery Request (PADR) packet</b>              | Packet sent by the client to one selected access concentrator to request a session.                                                                                                                                                                        |
| <b>PPPoE Active Discovery Session-Confirmation (PADS) packet</b> | Packet sent by the selected access concentrator to confirm the session.                                                                                                                                                                                    |
| <b>PPPoE Active Discovery Termination (PADT) packet</b>          | Packet sent by either the client or the access concentrator to terminate a session.                                                                                                                                                                        |
| <b>PPPoE over ATM</b>                                            | Network protocol that encapsulates PPPoE frames in Asynchronous Transfer Mode (ATM) frames for asymmetrical digital subscriber line (ADSL) transmission, and connects multiple hosts over a simple bridging access device to a remote access concentrator. |
| <b>virtual path identifier (VPI)</b>                             | An identifier of the virtual path that establishes a route between two devices in a network.                                                                                                                                                               |
| <b>virtual channel identifier (VCI)</b>                          | 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 or Asynchronous Transfer Mode (ATM) for ADSL 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 153

- PPPoE Stages on page 154
- Optional CHAP Authentication on page 155

## PPPoE Interfaces

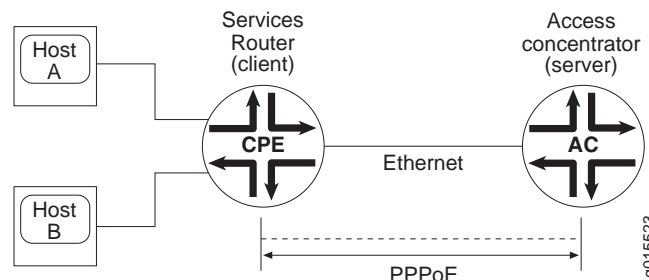
The PPPoE interface to the access concentrator can be either a Fast Ethernet interface on any Services Router or an ATM-over-ADSL interface on a J4300 or J6300 Services Router. The PPPoE configuration is the same for both interfaces. The only difference is the encapsulation for the underlying interface to the access concentrator:

- If the interface is Fast Ethernet, use a PPPoE encapsulation.
- If the interface is ATM for ADSL, use a PPPoE over ATM encapsulation.

## Fast Ethernet Interface

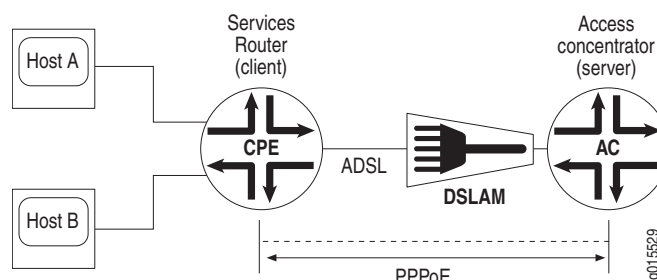
The Services Router encapsulates each PPP frame in an Ethernet frame and transports the frames over an Ethernet loop. Figure 29 shows a typical PPPoE session between a Services Router and an access concentrator on the Ethernet loop.

**Figure 29: PPPoE Session on the Ethernet Loop**



## ATM-over-ADSL 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 loop and a digital subscriber line access multiplexer (DSLAM). Figure 30 shows a typical PPPoE over ATM session between a Services Router and an access concentrator on an ADSL loop.

**Figure 30: 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. For more information about PPPoE stages, see “Interfaces Overview” on page 41.

### 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 *J-series Services Router Getting Started Guide*.
- Configure network interfaces. See “Configuring Network Interfaces” on page 103.

For a PPPoE over ATM interface, see “Adding an ATM-over-ADSL Network Interface with a Configuration Editor” on page 128.

## 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 155
- Configuring a PPPoE Interface (Required) on page 158
- Configuring CHAP (Optional) on page 161

### 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 interface, you must configure encapsulation on both the physical and logical interfaces. To configure encapsulation on an ATM-over-ADSL physical interface, use Ethernet over ATM encapsulation. To configure encapsulation on an ATM-over-ADSL logical interface, use PPPoE over AAL5 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 156
- Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface on page 157

## **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 37.
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 a PPPoE Interface (Required)” on page 158.
  - To enable authentication on the interface, see “Configuring CHAP (Optional)” on page 161.
  - To check the configuration, see “Verifying a PPPoE Configuration” on page 162.

**Table 37: Configuring PPPoE Encapsulation on an Ethernet Interface**

| <b>Task</b>                                                                      | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                        | <b>CLI Configuration Editor</b>                                                                                    |
|----------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy.          | In the configuration editor hierarchy select <b>Interfaces</b> .                                                                                                                                                                                                         | From the top of the configuration hierarchy, enter<br><br>edit interfaces                                          |
| Configure encapsulation on a logical Ethernet interface—for example, fe-0/0/1.0. | <ol style="list-style-type: none"> <li>1. In the Interface name box, click <b>fe-0/0/1</b>.</li> <li>2. In the Interface unit number box, click <b>0</b>.</li> <li>3. From the Encapsulation list, select <b>ppp-over-ether</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> | Set PPP encapsulation on unit 0 of the Ethernet interface:<br><br>set fe-0/0/1 unit 0 encapsulation ppp-over-ether |

### Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface

To configure PPPoE encapsulation on an ATM-over-ADSL 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 38.
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 a PPPoE Interface (Required)” on page 158.
  - To enable authentication on the interface, see “Configuring CHAP (Optional)” on page 161.
  - To check the configuration, see “Verifying a PPPoE Configuration” on page 162.

**Table 38: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface**

| <b>Task</b>                                                             | <b>J-Web Configuration Editor</b>                                | <b>CLI Configuration Editor</b>                                           |
|-------------------------------------------------------------------------|------------------------------------------------------------------|---------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy. | In the configuration editor hierarchy select <b>Interfaces</b> . | From the top of the configuration hierarchy, enter<br><br>edit interfaces |

**Table 38: Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface (continued)**

| <b>Task</b>                                                                                                               | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                              | <b>CLI Configuration Editor</b>                                                                    |
|---------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|
| Navigate to the ATM-over-ADSL interface—for example, <b>at-2/0/0</b> —and set the ATM virtual path identifier (VPI) to 0. | <ol style="list-style-type: none"> <li>1. In the Interface name box, click <b>at-2/0/0</b>.</li> <li>2. Next to ATM options, click <b>Configure</b>.</li> <li>3. Next to Vpi, click <b>Add new entry</b>.</li> <li>4. In the Vpi number box, type 0.</li> <li>5. Click <b>OK</b> twice.</li> </ol>             | <p>Enter</p> <p><b>set at-2/0/0 atm-options vpi 0</b></p>                                          |
| Configure the ADSL operating mode on the physical ATM interface—for example, autonegotiation.                             | <ol style="list-style-type: none"> <li>1. Next to Dsl options, click <b>Configure</b>.</li> <li>2. From the Operating mode list, select <b>auto</b>.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                   | <p>Enter</p> <p><b>set at-2/0/0 dsl-options operating-mode auto</b></p>                            |
| Configure Ethernet over ATM encapsulation on the physical ATM-over-ADSL interface.                                        | From the Encapsulation list, select <b>ethernet-over-atm</b> .                                                                                                                                                                                                                                                 | <p>Enter</p> <p><b>set at-2/0/0 encapsulation ethernet-over-atm</b></p>                            |
| Create an ATM-over-ADSL logical interface, configure LLC encapsulation, and specify a VCI number.                         | <ol style="list-style-type: none"> <li>1. Next to Unit, click <b>Add new entry</b>.</li> <li>2. In the Interface unit number box, type 0.</li> <li>3. From the Encapsulation list, select <b>ppp-over-ether-over-atm-llc</b>.</li> <li>4. In the Multicast vci box, type 0.120 and click <b>OK</b>.</li> </ol> | <p>Enter</p> <p><b>set at-2/0/0 unit 0 encapsulation ppp-over-ether-over-atm-llc vci 0.120</b></p> |

### **Configuring a PPPoE Interface (Required)**

To create and configure a PPPoE interface over the underlying Fast 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 39.
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 (Optional)” on page 161.
  - To check the configuration, see “Verifying a PPPoE Configuration” on page 162.



**Table 39: Configuring a PPPoE Interface**

| Task                                                                                                                                                                                         | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                  | CLI Configuration Editor                                                                                                                                                                              |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy.                                                                                                                      | In the configuration editor hierarchy select <b>Interfaces</b> .                                                                                                                                                                                                                                                                                            | From the top of the configuration hierarchy, enter<br><br>edit interfaces                                                                                                                             |
| Create a PPPoE interface with a logical interface unit 0.                                                                                                                                    | <ol style="list-style-type: none"> <li>Next to Interface, click <b>Add new entry</b>.</li> <li>In the Interface name box, type <b>pp0</b> and click <b>OK</b>.</li> <li>Under Interface name, click <b>pp0</b>.</li> <li>Next to Unit, click <b>Add new entry</b>.</li> <li>In the Interface unit number box, type 0.</li> </ol>                            | Enter<br><br>edit pp0 unit 0                                                                                                                                                                          |
| Configure an ISDN interface as the backup interface for the PPPoE interface—for example, <b>dl0.0</b> .                                                                                      | <ol style="list-style-type: none"> <li>Next to Backup options, click <b>Configure</b>.</li> <li>In the Interface box, type <b>dl0.0</b>.</li> <li>Click <b>OK</b>.</li> </ol>                                                                                                                                                                               | Enter<br><br>set backup-options interface dl0.0                                                                                                                                                       |
| Specify the logical Ethernet interface or the logical ATM interface as the underlying interface for the PPPoE session—for example, <b>fe-0/0/1.0</b> or <b>at-2/0/0.0</b> .                  | <ol style="list-style-type: none"> <li>Next to Pppoe options, click <b>Edit</b>.</li> <li>In the Underlying Interface box, type one of the following interface names: <ul style="list-style-type: none"> <li>For a logical Ethernet interface, type <b>fe-0/0/1.0</b>.</li> <li>For a logical ATM interface type, <b>at-2/0/0.0</b>.</li> </ul> </li> </ol> | Enter one of the following commands: <ul style="list-style-type: none"> <li>set pppoe-options underlying-interface fe-0/0/1.0.</li> <li>set pppoe-options underlying-interface at-2/0/0.0.</li> </ul> |
| Identify the access concentrator by a unique name—for example, <b>ispl.com</b> .                                                                                                             | In the Access concentrator box type <b>ispl.com</b> .                                                                                                                                                                                                                                                                                                       | Enter<br><br>set pppoe-options access-concentrator ispl.com                                                                                                                                           |
| Specify the time in seconds to reconnect after a PPPoE session is terminated—for example, <b>100 seconds</b> .                                                                               | In the Auto reconnect box, type <b>100</b> .                                                                                                                                                                                                                                                                                                                | Enter<br><br>set pppoe-options auto-reconnect 100                                                                                                                                                     |
| 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, <b>video@ispl.com</b> . | <ol style="list-style-type: none"> <li>In the Service name box, type <b>video@ispl.com</b>.</li> <li>Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                            | Enter<br><br>set pppoe-options service-name video@ispl.com                                                                                                                                            |

**Table 39: Configuring a PPPoE Interface (continued)**

| Task                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | CLI Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the maximum transmission unit (MTU) of the IPv4, IPv6, or Multiprotocol Label Switching (MPLS) protocol families—for example, 1492.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <ol style="list-style-type: none"> <li>Select one of the following protocol families: <ul style="list-style-type: none"> <li>For the IPv4 family, in the Inet box, select <b>Yes</b> and click <b>Configure</b>.</li> <li>For the IPv6 family, in the Inet6 box, select <b>Yes</b> and click <b>Configure</b>.</li> <li>For the MPLS family, in the Mpls box, select <b>Yes</b> and click <b>Configure</b>.</li> </ul> </li> <li>In the Mtu box, type 1492.</li> <li>Click <b>OK</b> until you return to the Unit page.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <p>Enter one of the following:</p> <ul style="list-style-type: none"> <li>set family inet mtu 1492</li> <li>set family inet6 mtu 1492</li> <li>set family mpls mtu 1492</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <p>Configure the PPPoE logical interface address in one of the following ways:</p> <p>Do one of the following:</p> <ul style="list-style-type: none"> <li>Assign IPv4 or IPv6 source and destination addresses—for example: <ul style="list-style-type: none"> <li>192.168.1.1/32 and 192.168.1.2 for IPv4</li> <li>2004:1/128 and 2004:2 for IPv6.</li> </ul> </li> <li>Derive the source address from a specified interface—for example, the loopback interface, lo0.0—and assign a destination address—for example, 192.168.1.2. The specified interface must include a logical unit number and have a configured IP address.</li> <li>Obtain an IP address by negotiation with the remote end. This method might require the access concentrator to use a RADIUS authentication server.</li> </ul> | <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"> <li>Next to Address, click <b>Add new entry</b>.</li> <li>In the Inet Source box, type 192.168.1.1/32, or in the Inet6 Source box, type 2004::1/128.</li> <li>In the Inet Destination box, type 192.168.1.2, or in the Inet6 Destination box, type 2004::2.</li> <li>Click <b>OK</b> until you return to the Unit page.</li> </ol> <p>To derive the source address and assign the destination address:</p> <ol style="list-style-type: none"> <li>Next to Unnumbered address, select the <b>Yes</b> check box and click <b>Configure</b>.</li> <li>In the Destination box, type 192.168.1.2.</li> <li>In the Source box, type lo0.0.</li> <li>Click <b>OK</b> until you return to the Unit page.</li> </ol> <p>To obtain an IP address from the remote end:</p> <ol style="list-style-type: none"> <li>Next to Negotiate address, select the <b>Yes</b> check box.</li> <li>Click <b>OK</b> until you return to the Unit page.</li> </ol> | <p>Do one of the following:</p> <ul style="list-style-type: none"> <li>To assign source and destination addresses enter one of the following sets of commands: <ul style="list-style-type: none"> <li>For IPv4 addresses, set family inet address 192.168.1.1/32 destination 192.168.1.2.</li> <li>For IPv6 addresses, set family inet6 address 2004::1/128 destination 2004::2.</li> </ul> </li> <li>To derive the source address and assign the destination address, enter set family inet unnumbered-address lo0.0 destination 192.168.1.2.</li> <li>To obtain an IP address from the remote end, enter set family inet negotiate-address.</li> </ul> |
| Disable the sending of keepalives on a logical interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <ol style="list-style-type: none"> <li>From the Keepalive choices list, select <b>no keepalives</b>.</li> <li>Click <b>OK</b> to apply your entries to the configuration.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <p>Enter</p> <p>set no-keepalives</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

## Configuring CHAP (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 40.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying a PPPoE Configuration” on page 162.

**Table 40: Configuring CHAP**

| Task                                                                                                                                                       | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                               | CLI Configuration Editor                                                                     |
|------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| Navigate to the <b>Profile</b> level in the configuration hierarchy.                                                                                       | <ol style="list-style-type: none"> <li>1. In the configuration hierarchy, select <b>Configuration &gt; Edit Configuration &gt; View and Edit</b>.</li> <li>2. Next to Access, click <b>Configure</b> or <b>Edit</b>.</li> </ol>                                                                                                                                                          | <p>Enter</p> <p>set access profile A-ppp-client client client1<br/>chap-secret my-secret</p> |
| Define a CHAP access profile—for example, <b>A-ppp-client</b> —with a client named <b>client 1</b> and the secret (password) <b>my-secret</b> .            | <ol style="list-style-type: none"> <li>1. Next to Profile, click <b>Add new entry</b>.</li> <li>2. In the Profile name box, type A-ppp-client.</li> <li>3. Next to Client, click <b>Add new entry</b>.</li> <li>4. In the Name box, type client1.</li> <li>5. In the Chap secret box, type my-secret.</li> <li>6. Click <b>OK</b> until you return to the Configuration page.</li> </ol> |                                                                                              |
| Navigate to the <b>pp0 unit 0</b> interface level in the configuration hierarchy.                                                                          | <ol style="list-style-type: none"> <li>1. In the configuration hierarchy, select <b>Interfaces</b>.</li> <li>2. In the Interface name box, click <b>pp0</b>.</li> <li>3. In the Interface unit number box, click <b>0</b>.</li> </ol>                                                                                                                                                    | <p>From the top of the configuration hierarchy, enter</p> <p>edit interfaces pp0 unit 0</p>  |
| Configure CHAP on the PPPoE interface, and specify a unique profile name containing a client list and access parameters—for example, <b>A-ppp-client</b> . | <ol style="list-style-type: none"> <li>1. Next to Ppp options, click <b>Configure</b>.</li> <li>2. Next to Chap, click <b>Configure</b>.</li> <li>3. In the Access profile box, type A-ppp-client.</li> </ol>                                                                                                                                                                            | <p>Enter</p> <p>set ppp-options chap access-profile A-ppp-client</p>                         |

**Table 40: Configuring CHAP (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-fe-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"> <li>■ For an Ethernet interface, type A-fe-0/0/1.0.</li> <li>■ For an ATM interface, type A-at-2/0/0.0.</li> </ul> | <p>Do one of the following:</p> <ul style="list-style-type: none"> <li>■ For the Ethernet interface, enter <b>set ppp-options chap local-name A-fe-0/0/1.0</b>.</li> <li>■ For the ATM interface, enter <b>set ppp-options chap local-name A-at-2/0/0.0</b>.</li> </ul> |
| Set the <b>passive</b> option to handle incoming CHAP packets only.                                                    | <ol style="list-style-type: none"> <li>1. In the Passive box, click <b>Yes</b>.</li> <li>2. Click <b>OK</b>.</li> </ol>                                                                                        | <p>Enter</p> <p><b>set ppp-options chap passive</b></p>                                                                                                                                                                                                                 |

## Verifying a PPPoE Configuration

To verify PPPoE configuration perform the following tasks:

- Displaying a PPPoE Configuration for an Ethernet Interface on page 162
- Displaying a PPPoE Configuration for an ATM-over-ADSL Interface on page 163
- Verifying PPPoE Interfaces on page 164
- Verifying PPPoE Sessions on page 165
- Verifying the PPPoE Version on page 166
- Verifying PPPoE Statistics on page 166

### 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.

**Sample Output**

```
[edit]
user@host# show interfaces
fe-0/0/1 {
 unit 0 {
 }
}
pp0 {
 mtu 1492;
 unit 0 {
```

```

ppp-options {
 chap {
 access-profile A-ppp-client;
 local-name A-fe-0/0/1.0;
 }
}
pppoe-options {
 underlying-interface fe-0/0/1;
 access-concentrator ispl.com;
 service-name "video@ispl.com";
 auto-reconnect 100;
}
no-keepalives;
family inet {
 address 192.168.1.1/32 {
 destination 192.168.1.2;
 }
}
family inet6 {
 address 2004:1/128 {
 destination 2004:2/128;
 }
}
family mpls;
}

```

**What It Means** Verify that the output shows the intended configuration of PPPoE. For more information about the format of a configuration file, see “Viewing the Configuration Text” on page 12.

### ***Displaying a PPPoE Configuration for an ATM-over-ADSL Interface***

**Purpose** Verify the PPPoE 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 command from the top level.

**Sample Output**

```

[edit]
user@host# show interfaces
at-2/0/0 {
 encapsulation ethernet-over-atm;
 atm-options {
 vpi 0;
 }
 dsl-options {
 operating-mode auto;
 }
 unit 0 {
 encapsulation ppp-over-ether-over-atm-llc;
 vci 0.120;
 }
}

```

```

}
pp0 {
 mtu 1492;
 unit 0 {
 ppp-options {
 chap {
 access-profile A-ppp-client;
 local-name A-at-2/0/0.0;
 }
 }
 pppoe-options {
 underlying-interface at-2/0/0;
 access-concentrator ispl.com;
 service-name "video@ispl.com";
 auto-reconnect 100;
 }
 no-keepalives;
 family inet {
 negotiate-address;
 }
 }
}

```

**What It Means** Verify that the output shows the intended configuration of PPPoE. For more information about the format of a configuration file, see “Viewing the Configuration Text” on page 12.

## Verifying PPPoE Interfaces

**Purpose** Verify that the PPPoE router interfaces are configured properly.

**Action** From the CLI, enter the show interfaces pp0 command.

**Sample Output** 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: ispl.com, AC MAC address: 00:90:1a:40:f6:4c,
 Service name: video@ispl.com, Configured AC name: ispl.com,
 Auto-reconnect timeout: 60 seconds
 Underlying interface: fe-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—for example, fe-5/0/0.0.
  - For an ATM-over-ADSL connection, the underlying interface is ATM—for example, at-2/0/0.0.

For more information about the `show interfaces pp0` command, 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.

**Sample Output** `user@host> show pppoe interfaces`

```

pp0.0 Index 67
 State: Session up, Session ID: 31,
 Service name: video@ispl.com, Configured AC name: ispl.com,
 Session AC name: belur, AC MAC address: 00:90:1a:40:f6:4e,
 Auto-reconnect timeout: 1 seconds,
 Underlying interface: fe-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—for example, **fe-0/0/1.0**.
    - For an ATM-over-ADSL connection, the underlying interface is ATM—for example, **at-2/0/0.0**.

For more information about the `show pppoe interfaces` command, 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.
- Sample Output**
- ```
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.

For more information about the `show pppoe version` command, 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.
- Sample Output**
- ```
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.

For more information about the `show pppoe statistics` command, see the *JUNOS Interfaces Command Reference*.



## Chapter 5

# Configuring ISDN

ISDN connectivity is supported on the J-series Services Routers as a backup for a primary Internet connection. This chapter contains the following topics:

- ISDN Terms on page 169
- ISDN Overview on page 170
- Before You Begin on page 172
- Configuring ISDN Interfaces with a Configuration Editor on page 172
- Verifying the ISDN Configuration on page 192

## ISDN Terms

Before configuring ISDN, become familiar with the terms defined in Table 41.

**Table 41: ISDN Terminology**

| Term                       | Definition                                                                                                                                                                                                                                                                      |
|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| bearer-channel (B-channel) | Channel that carries main data on an ISDN interface.                                                                                                                                                                                                                            |
| basic rate interface (BRI) | ISDN interface intended for home and small enterprise applications, BRI consists of two 64 Kbps B-channels and one 16 Kbps D-channel.                                                                                                                                           |
| bandwidth on-demand        | ISDN interface is activated when network activity reaches a pre-defined threshold and provides additional bandwidth on the network.                                                                                                                                             |
| delta-channel (D-channel)  | Channel that carries control and signaling information on an ISDN interface.                                                                                                                                                                                                    |
| dial backup                | Feature that allows one or more dialer interfaces to be used as a backup link for the primary interface.                                                                                                                                                                        |
| dialer interface (dl)      | Logical interface configured as the activation interface for an ISDN connection.                                                                                                                                                                                                |
| dial-on-demand routing     | Feature that allows the ISDN connection to appear “always on.” When an interesting packet arrives at the ISDN interface, connectivity is established over ISDN. Connectivity is lost after a configured period of inactivity and thus saves expensive inactive connection time. |

**Table 41: ISDN Terminology (continued)**

| <b>Term</b>                                       | <b>Definition</b>                                                                                                                                                                                                                                                                                                                                             |
|---------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>dialer profile</b>                             | 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.      |
| <b>dialer watch</b>                               | Feature that provides reliable connectivity without relying on “interesting” network traffic to activate the ISDN interface. Dialer watch integrates failover support with routing capabilities. The J-series Services Router monitors the existence of a route. If the route is absent, dialer watch initiates the ISDN interface for failover connectivity. |
| <b>floating static route</b>                      | Routes with an administrative distance greater than the administrative distance of dynamic routes.                                                                                                                                                                                                                                                            |
| <b>integrated services digital network (ISDN)</b> | 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.                                                                                                                                                        |
| <b>terminal endpoint identifier (TEI)</b>         | 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 range 0–63 are used for static TEI assignment; 64–126 are used for dynamic assignment; and 127 is used for group assignment.                                      |
| <b>service profile identifier (SPID)</b>          | 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.                                                                                                                                                                                                  |

## 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.

ISDN provides a Services Router with a backup connection for network interfaces.

## ISDN Interfaces

There are four types of interfaces available for ISDN connectivity:

- 1-port S/T interface supporting ITU-T I.430, ETSI TS 101080, and GR-1089-Core Type III built into a J2300 Services Router
- 1-port U interface supporting ANSI T.601 and GR-1089-Core built into a J2300 Services Router
- 4-port S/T PIM supporting ITU-T I.430, ETSI TS 101080, and GR-1089-Core Type III as a field-replaceable unit (FRU) on J4300 and J6300 Services Routers
- 4-port U PIM supporting ANSI T.601 and GR-1089-Core as a FRU on J4300 and J6300 Services Routers

Each ISDN interface uses the naming convention *br-pim /0/ port*.

Each B-channel is identified by *bc-pim /0/ port:channel*, where *channel* represents the B-channel ID and has a value of 1 or 2.

The D-channel is identified by *dc-pim /0/ port*.



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

---

The dialer interface, *dl n*, is a logical interface for configuring 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 cannot be configured simultaneously in the following modes:

- Backup interface and dialer filter
- Backup interface and dialer watch
- Dialer watch and dialer filter
- Backup interface for more than one primary 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 *J-series Services Router Getting Started Guide*.
- Establish basic connectivity. For more information, see the *J-series Services Router Getting Started Guide*.
- Order an ISDN line from your telecommunications service provider. Contact your service provider for more information.
- If you do not already have an understanding of physical and logical interfaces and Juniper Networks interface conventions, read “Configuring Network Interfaces” on page 103.

Although it is not a requirement, you may also want to plan how you are going to use the ISDN interfaces on your network before you begin configuring them. You can see a list of ISDN interfaces by displaying the **Configuration > View and Edit > Edit Configuration > Interfaces** page.

## Configuring ISDN Interfaces with a Configuration Editor

---

- Adding an ISDN Interface (Required) on page 172
- Configuring a Dialer Interface (Required) on page 175
- Enabling an ISDN Interface as a Secondary Connection (Optional) on page 178
- Configuring Dial-on-Demand Connectivity (Optional) on page 179
- Configuring Bandwidth-on-Demand (Optional) on page 181
- Configuring Dial-on-Demand Routing (Optional) on page 184
- Configuring Dialer Watch (Optional) on page 186
- Configuring Dial-on-Demand Routing with OSPF Support (Optional) on page 190
- Configuring Dialer Profiles (Optional) on page 191

### ***Adding an ISDN Interface (Required)***

To enable ISDN interfaces installed on your Services Router to work properly, you must configure the interface properties.

To configure an ISDN 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 42.
3. When you are finished configuring the interface, go to “Configuring a Dialer Interface (Required)” on page 175.
4. To verify that the network interface is configured correctly, see “Verifying the ISDN Configuration” on page 192.

**Table 42: Adding an ISDN Interface**

| Task                                                                                                                                                                                                                                                                                                                                                 | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                             | CLI Configuration Editor                                                                                                 |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy.                                                                                                                                                                                                                                                                              | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Configuration &gt; View and Edit &gt; Edit Configuration</b>.</li> <li>2. Next to Interfaces, click <b>Configure</b> or <b>Edit</b>.</li> </ol>                                                                                                                                                                             | <p>From the top of the command hierarchy, enter</p> <p>edit interfaces br-1/0/3</p>                                      |
| Create the new interface—for example, br-1/0/3.                                                                                                                                                                                                                                                                                                      | <ol style="list-style-type: none"> <li>1. Next to Interface, click <b>Add new entry</b>.</li> <li>2. In the Interface name box, type the name of the new interface, br-1/0/3.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                  |                                                                                                                          |
| Configure dialer options. <ul style="list-style-type: none"> <li>■ Name the dialer pool—for example, ISDN-dialer-group.</li> <li>■ Set the dialer pool priority—for example, 25.</li> </ul> <p>Dialer pool priority has a range from 0 to 255 with 0 designating lowest-priority interfaces and 255 designating the highest-priority interfaces.</p> | <ol style="list-style-type: none"> <li>1. In the Encapsulation column, next to the new interface, click <b>Edit</b>.</li> <li>2. Next to Dialer options, select <b>Yes</b>, and then click <b>Edit</b>.</li> <li>3. Next to Pool, click <b>Add new entry</b>.</li> <li>4. In the Pool identifier box, type <b>isdn-dialer-group</b>.</li> <li>5. In the Priority box, type 25.</li> <li>6. Click <b>OK</b>.</li> </ol> | <p>From the [edit interfaces br-1/0/3] hierarchy, enter</p> <p>set dialer-options pool ISDN-dialer-group priority 25</p> |

**Table 42: Adding an ISDN Interface (continued)**

| Task                                                                                                                                                                                                                                                                                                       | J-Web Configuration Editor                        | CLI Configuration Editor                               |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------|--------------------------------------------------------|
| Configure ISDN properties.                                                                                                                                                                                                                                                                                 | 1. Next to Isdn options, click <b>Configure</b> . | 1. To set the ISDN options, enter                      |
| <ul style="list-style-type: none"> <li>■ Calling number of your ISDN provider—for example, 18005555555.</li> </ul>                                                                                                                                                                                         | 2. In the Calling number box, type 18005555555.   | <pre>set isdn-options calling-number 18005555555</pre> |
| <ul style="list-style-type: none"> <li>■ Service provider ID (SPID)—for example, 00108005555555.</li> </ul>                                                                                                                                                                                                | 3. In the Spid1 box, type 00108005555555.         | 2. Enter                                               |
| <ul style="list-style-type: none"> <li>■ 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.</li> </ul>                                                                                                    | 4. In the Static tei val box, type 23.            | <pre>set isdn-options spid1 00108005555555</pre>       |
|                                                                                                                                                                                                                                                                                                            |                                                   | 3. Enter                                               |
|                                                                                                                                                                                                                                                                                                            |                                                   | <pre>set isdn-options static-tei-value 23</pre>        |
| <p>If you have a NTDMS-100 or NI1 switch, an additional box for a service provider ID is provided.</p> <p>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.</p> |                                                   |                                                        |
| Select the type of ISDN switch—for example, <b>att5e</b> . The following switches are compatible with Services Routers:                                                                                                                                                                                    | From the Switch type list, select <b>att5e</b> .  | To select the switch type, enter                       |
| <ul style="list-style-type: none"> <li>■ <b>att5e</b>—AT&amp;T 5ESS</li> <li>■ <b>etsi</b>—NET3 for the UK and Europe</li> <li>■ <b>ni1</b>—National ISDN-1</li> <li>■ <b>ntdms-100</b>—Northern Telecom DMS-100</li> <li>■ <b>ntt</b>—NTT Group switch for Japan</li> </ul>                               |                                                   | <pre>set isdn-options switch-type att5e</pre>          |



**Table 42: Adding an ISDN Interface (continued)**

| <b>Task</b>                                                                                                                                                                                                                                                                                     | <b>J-Web Configuration Editor</b>                                                                                                                                    | <b>CLI Configuration Editor</b>                                                                                                                                            |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 <b>10</b> seconds, but can be configured between <b>1</b> and <b>65536</b> seconds—for example, <b>15</b> .                                               | <ol style="list-style-type: none"> <li>1. In the T306 box, type <b>15</b>.</li> <li>2. In the T310 box, type <b>15</b>.</li> </ol>                                   | <ol style="list-style-type: none"> <li>1. Enter<br/><br/><code>set isdn-options t306 15</code></li> <li>2. Enter<br/><br/><code>set isdn-options t310 15</code></li> </ol> |
| Configure when the TEI negotiates with the ISDN provider. <ul style="list-style-type: none"> <li>■ <b>first-call</b>—Activation does not occur until a call is sent.</li> <li>■ <b>power-up</b>—Activation occurs when the Services Router is powered on. This is the default value.</li> </ul> | <ol style="list-style-type: none"> <li>1. From the Tei option list, select <b>power-up</b>.</li> <li>2. Click <b>OK</b> to return to the Interfaces page.</li> </ol> | To initiate activation at power-up, enter<br><br><code>set isdn-options tei-option power-up</code>                                                                         |

### Configuring a Dialer Interface (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.

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 43.
3. If you are finished configuring the router, commit the configuration.
4. Go on to any of the following optional tasks:
  - “Enabling an ISDN Interface as a Secondary Connection (Optional)” on page 178.
  - “Configuring Dial-on-Demand Connectivity (Optional)” on page 179.
  - “Configuring Bandwidth-on-Demand (Optional)” on page 181.
  - “Configuring Dial-on-Demand Routing (Optional)” on page 184.
  - “Configuring Dialer Watch (Optional)” on page 186.
  - “Configuring Dial-on-Demand Routing with OSPF Support (Optional)” on page 190.

- “Configuring Dialer Profiles (Optional)” on page 191.
5. To verify that the network interface is configured correctly, see “Verifying the ISDN Configuration” on page 192.

**Table 43: Adding a Dialer Interface to a Services Router**

| Task                                                                                                                                                                                                                                                                                                                                                                                                 | J-Web Configuration Editor                                                                                                                                                                                                                       | CLI Configuration Editor                                                                                                                                         |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy.                                                                                                                                                                                                                                                                                                                              | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Configuration &gt; View and Edit &gt; Edit Configuration</b>.</li> <li>2. Next to Interfaces, click <b>Configure</b> or <b>Edit</b>.</li> </ol>       | From the top of the configuration hierarchy, enter<br>edit interfaces                                                                                            |
| Create the new interface—for example, <b>d10</b> .<br><br>Adding a description can differentiate between different dialer interfaces—for example, <b>backup</b> .                                                                                                                                                                                                                                    | <ol style="list-style-type: none"> <li>1. Next to Interface, click <b>Add new entry</b>.</li> <li>2. In the Interface name box, type <b>d10</b>.</li> <li>3. In the Description box, type <b>backup</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> | Create and name the interface:<br><br><ol style="list-style-type: none"> <li>1. edit interfaces d10</li> <li>2. set interfaces d10 description backup</li> </ol> |
| Configure encapsulation options—for example, <b>cisco-hdlc</b> .<br><br><ul style="list-style-type: none"> <li>■ <b>cisco-hdlc</b>—Cisco-compatible High-Level Data Link Control is a group of protocols for transmitting data between network points.</li> <li>■ <b>ppp</b>—Point-to-Point Protocol is a protocol used for communication between two computers using a serial interface.</li> </ul> | <ol style="list-style-type: none"> <li>1. In the <b>Encapsulation</b> column, next to the new interface, click <b>Edit</b>.</li> <li>2. From the Encapsulation list, select <b>cisco-hdlc</b>.</li> </ol>                                        | Enter<br><br>set encapsulation cisco-hdlc                                                                                                                        |
| Enter a hold-time value in milliseconds—for example, <b>60</b> . 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 down for the hold-time period. Similarly, an interface is not advertised as up until it remains up for the hold-time period.                                 | <ol style="list-style-type: none"> <li>1. In the Hold time section, type <b>60</b> in the Down box.</li> <li>2. In the Up box, type <b>60</b>.</li> </ol>                                                                                        | <ol style="list-style-type: none"> <li>1. Enter<br/><br/>set hold-time down 60</li> <li>2. Enter<br/><br/>set hold-time up 60</li> </ol>                         |

**Table 43: Adding a Dialer Interface to a Services Router (continued)**

| Task                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | J-Web Configuration Editor                                                                                                                                                                                                                                              | CLI Configuration Editor                                                                                                                                                                                                                                                 |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create the logical unit—for example, 0.<br><br><b>NOTE:</b> You can only set the logical unit to 0 unless you are configuring the dialer interface for Multilink PPP encapsulation.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <ol style="list-style-type: none"> <li>Next to Unit, click <b>Add new entry</b>.</li> <li>In the Interface unit number box, type 0.</li> <li>Next to Dialer options, select <b>Yes</b>, and then click <b>Configure</b>.</li> </ol>                                     | Enter<br><br>set unit 0                                                                                                                                                                                                                                                  |
| Configure dialer options.<br><br><ul style="list-style-type: none"> <li><b>Activation delay</b>—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.</li> <li><b>Deactivation delay</b>—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.</li> <li><b>Idle timeout</b>—Time a connection is idle before disconnecting—for example, 30. Default value is 120 seconds with a range from 0 to 4294967295.</li> <li><b>Initial route check</b>—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.</li> <li><b>Load interval</b>—Interval of time between calculations of the average load on the network. Default value is 60 seconds and has a range of 20-180 seconds incremented in 10 seconds. Used only when configuring bandwidth-on-demand.</li> <li><b>Load threshold</b>—Percentage of load on all links—for example 90. Default value is 100 with a range from 0 to 100. Used only for configuring bandwidth on demand.</li> <li><b>Pool</b>—Name of a group of ISDN interfaces configured to use the dialer interface—for example, 3.</li> </ul> | <ol style="list-style-type: none"> <li>In the Activation delay box, type 60.</li> <li>In the Deactivation delay box, type 30.</li> <li>In the Idle timeout box, type 30.</li> <li>In the Initial route check box, type 30.</li> <li>In the Pool box, type 3.</li> </ol> | <ol style="list-style-type: none"> <li>Enter<br/><br/>edit dialer-options</li> <li>Enter<br/><br/>set activation-delay 30</li> <li>Enter<br/><br/>set deactivation-delay 30</li> <li>Enter<br/><br/>set idle-timeout 30<br/>initial-route-check 30<br/>pool 3</li> </ol> |
| Configure the remote destination to call—for example, 5551212.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | In the Dial string box, type 5551212.                                                                                                                                                                                                                                   | Enter<br><br>set dial-string 5551212                                                                                                                                                                                                                                     |

**Table 43: Adding a Dialer Interface to a Services Router (continued)**

| <b>Task</b>                                                                                                      | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                            | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure a list of routes to watch—for example, 192.1.1.0/24. Specify one or more IP address prefixes.          | <ol style="list-style-type: none"> <li>1. Next to Watch list, click <b>Add new entry</b>.</li> <li>2. In the Prefix box, type 192.1.1.0/24.</li> <li>3. Click <b>OK</b> until you return to the Unit page.</li> </ol>                                                        | Enter<br><br>set watch-list<br>192.1.1.0/24                                                                                                                                                                                                                                                                     |
| Configure source and destination IP addresses for the dialer interface—for example, 172.20.10.2 and 172.20.10.1. | <ol style="list-style-type: none"> <li>1. Select <b>Inet</b> under Family, and click <b>Edit</b>.</li> <li>2. Next to Address, click <b>Add new entry</b>.</li> <li>3. In the Source box, type 172.20.10.2.</li> <li>4. In the Destination box, type 172.20.10.1.</li> </ol> | <ol style="list-style-type: none"> <li>1. Enter<br/><br/>               set family<br/>               inet address<br/>               172.20.10.2</li> <li>2. Enter<br/><br/>               set family<br/>               inet address<br/>               destination<br/>               172.20.10.1</li> </ol> |

### **Enabling an ISDN Interface as a Secondary Connection (Optional)**

Continuous network connectivity is important to every network and crucial for businesses that depend on network connectivity for day-to-day operations and important business applications. The Services Router can be configured to fail over to the ISDN interface when the primary connection experiences interruptions in Internet connectivity.

ISDN backup connectivity is supported on all interfaces except ls-0/0/0.

To configure an 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 44.
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 192.

**Table 44: Configuring an Interface for ISDN Backup**

| <b>Task</b>                                                             | <b>J-Web Configuration Editor</b>                                                                                                                                                                                       | <b>CLI Configuration Editor</b>                                                                                               |
|-------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy. | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Configuration &gt; View and Edit &gt; Edit Configuration</b>.</li> <li>2. Next to Interfaces, click <b>Edit</b>.</li> </ol>  | <p>From the top of the configuration hierarchy, enter</p> <pre>edit interfaces <i>interface-name</i> unit <i>number</i></pre> |
| Select the interface for backup ISDN connectivity.                      | <ol style="list-style-type: none"> <li>1. In the Interface name column, click the interface name.</li> <li>1. Under Unit, in the Nested Configuration column, click <b>Edit</b>.</li> </ol>                             |                                                                                                                               |
| Configure the backup interface—for instance, d10.0.                     | <ol style="list-style-type: none"> <li>1. Next to Backup options, click <b>Configure</b>.</li> <li>2. In the Interface box, type d10.0.</li> <li>3. Click <b>OK</b> until you return to the Interfaces page.</li> </ol> | <p>Enter</p> <pre>set backup-options d10.0</pre>                                                                              |

### Configuring Dial-on-Demand Connectivity (Optional)

Dial-on-demand connectivity 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.

You define an interesting packet using the firewall filter feature of the Services Router. There are two steps to configuring dial-on-demand connectivity:

- Configuring a Dialer Filter on page 179
- Applying the Dial-on-Demand Dialer Filter to the Dialer Interface on page 180

#### Configuring a Dialer Filter

To configure dial-on-demand 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 45.
3. Go on to “Applying the Dial-on-Demand Dialer Filter to the Dialer Interface” on page 180.

**Table 45: Configuring the Dialer Filter for Interesting Packets**

| <b>Task</b>                                                                                                                                                                                                                                                                          | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>CLI Configuration Editor</b>                                                                                                                          |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Firewall</b> level in the configuration hierarchy.                                                                                                                                                                                                                | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Configuration &gt; View and Edit &gt; Edit Configuration</b>.</li> <li>2. Next to Firewall, click <b>Edit</b>.</li> </ol>                                                                                                                                                                                                                                          | <p>From the top of the configuration hierarchy, enter</p> <pre>edit firewall</pre>                                                                       |
| Configure the dialer filter name—for example, <b>int-packet</b> .                                                                                                                                                                                                                    | <ol style="list-style-type: none"> <li>1. Next to <b>Inet</b>, click <b>Edit</b>.</li> <li>2. Next to <b>Dialer filter</b>, click <b>Add new entry</b>.</li> <li>3. In the Filter name box, type <b>int-packet</b>.</li> </ol>                                                                                                                                                                                                                                | <ol style="list-style-type: none"> <li>1. Enter <pre>edit family inet</pre> </li> <li>2. Then enter <pre>edit dialer-filter int-packet</pre> </li> </ol> |
| <p>Configure the firewall rule name—for example, <b>term1</b>.</p> <p>Configure term behavior. For example, you might want to configure your interesting packet as an ICMP packet.</p> <p>To configure the term completely, include both <b>from</b> and <b>then</b> statements.</p> | <ol style="list-style-type: none"> <li>1. Next to Term, click <b>Add new entry</b>.</li> <li>2. In the Rule name box, type <b>term1</b>.</li> <li>3. Next to From, click <b>Configure</b>.</li> <li>4. From the Protocol choice list, select <b>Protocol</b>.</li> <li>5. Next to Protocol, click <b>Add new entry</b>.</li> <li>6. From the Value keyword list, select <b>icmp</b>.</li> <li>7. Click <b>OK</b> twice to return to the Term page.</li> </ol> | <p>Enter</p> <pre>set term term1 from protocol icmp</pre>                                                                                                |
| Configure the <b>then</b> part of the dialer filter.                                                                                                                                                                                                                                 | <ol style="list-style-type: none"> <li>1. Next to Then, click <b>Configure</b>.</li> <li>2. From the Designation list, select <b>Note</b>.</li> </ol>                                                                                                                                                                                                                                                                                                         | <p>Enter</p> <pre>set term1 then note</pre>                                                                                                              |

## Applying the Dial-on-Demand Dialer Filter to the Dialer Interface

To complete dial-on-demand connectivity 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 46.
3. When 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 192.

**Table 46: Applying the Dialer Filter to the Dialer Interface**

| <b>Task</b>                                                             | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                     | <b>CLI Configuration Editor</b>                                                                                                    |
|-------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy. | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Configuration &gt; View and Edit &gt; Edit Configuration</b>.</li> <li>2. Next to Interfaces, click <b>Edit</b>.</li> </ol>                                                                                                                                | From the top of the configuration editor hierarchy, enter<br>edit interfaces d10 unit 0                                            |
| Select the dialer interface to apply the filter—for example, d10.0      | <ol style="list-style-type: none"> <li>1. In the Interface name column, click <b>d10.0</b>.</li> <li>2. Under Unit, in the Nested Configuration column, click <b>Edit</b>.</li> </ol>                                                                                                                                                                 |                                                                                                                                    |
| Select the dialer filter and apply it to the dialer interface.          | <ol style="list-style-type: none"> <li>1. In the Family section, next to Inet, click <b>Edit</b>.</li> <li>2. Next to Filter, click <b>Configure</b>.</li> <li>3. In the Dialer box, type <b>int-packet</b>, the dialer-filter configured in “Configuring a Dialer Filter” on page 179, as the dialer-filter.</li> <li>4. Click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>1. Enter<br/>edit family inet filter</li> <li>2. Enter<br/>set dialer int-packet</li> </ol> |

### **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 and assigned a single dialer profile. Initially, only one ISDN link is active and all packets are sent through this interface. When a predefined threshold is reached, the dialer interface activates another ISDN link and initiates a data connection.

### **Configuring a Dialer Interface 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 47.
3. Go on to “Configuring an ISDN Interface for Bandwidth-on-Demand” on page 183.

**Table 47: Configuring a Dialer Interface for Bandwidth-on-Demand**

| Task                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                       | CLI Configuration Editor                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Configuration &gt; View and Edit &gt; Edit Configuration</b>.</li> <li>2. Next to Interfaces, click <b>Edit</b>.</li> <li>3. Next to the dialer interface name, click <b>Edit</b>.</li> </ol>                                                                                                                                                                                         | <p>From the top of the configuration editor, enter</p> <pre>edit interfaces d10</pre>                                                                                                                                                                                                                                                          |
| Configure multilink properties on the dialer interface.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <ol style="list-style-type: none"> <li>1. Select <b>multilink-ppp</b> as the encapsulation type.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                      | <p>From the [edit interfaces d10] hierarchy, enter</p> <pre>set encapsulation multilink-ppp</pre>                                                                                                                                                                                                                                              |
| <p>Configure the dialer options.</p> <ul style="list-style-type: none"> <li>■ <b>Dial string</b>—Telephone number for the interface to dial that establishes ISDN connectivity—for example, 4085551515.</li> <li>■ <b>Idle timeout</b>—Time a connection is idle before disconnecting—for example, 300. Default value is 120 seconds with a range from 0 to 4294967295.</li> <li>■ <b>Load interval</b>—Interval of time between average network load calculations—for example, 90. Default value is 60 seconds with a range of 20-180 seconds incremented in 10 seconds.</li> <li>■ <b>Load threshold</b>—Percentage of load on all links—for example 90. Default value is 100 with a range from 0 to 100.</li> <li>■ <b>Pool</b>—Name of a group of ISDN interfaces configured to use the dialer interface—for example, 3.</li> </ul> | <ol style="list-style-type: none"> <li>1. In the Unit section, click <b>Dialer options</b> under Encapsulation.</li> <li>2. Next to Dial string, click <b>Add new entry</b>.</li> <li>3. In the Value box, type 4085551515 and click <b>OK</b>.</li> <li>4. In the Idle timeout box, type 300.</li> <li>5. In the Load interval box, type 90.</li> <li>6. In the Load threshold box, type 95.</li> <li>7. In the Pool box, type bw-pool.</li> <li>8. Click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>1. Enter<br/>edit unit 0</li> <li>2. Enter<br/>edit dialer-options</li> <li>3. Enter<br/>set dial-string 4085551515</li> <li>4. Enter<br/>set idle-timeout 300</li> <li>5. Enter<br/>set load-interval 90</li> <li>6. Enter<br/>set load-threshold 95</li> <li>7. Enter<br/>set pool bw-pool</li> </ol> |
| <p>Configure unit properties.</p> <p>To configure a multiple dialer interfaces for bandwidth-on-demand, increment the Unit number—for example, d10.1, d10.2, and so on.</p> <p><b>F max period</b> is the maximum number of compressed packets between transmission of full packets. The value can be between 1 and 65535.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <ol style="list-style-type: none"> <li>1. Next to Compression, select <b>Yes</b>, and then <b>Configure</b>.</li> <li>2. Select <b>Rtp</b>, and then <b>Configure</b>.</li> <li>3. In the F max period box, type 100.</li> <li>4. Next to Queues, click <b>Add new entry</b>.</li> <li>5. From the Value list, select <b>q3</b>. Then click <b>OK</b> and <b>OK</b> again.</li> </ol>                                                                                            | <ol style="list-style-type: none"> <li>1. From the <b>edit interfaces dl</b> hierarchy, enter<br/>edit unit 0</li> <li>2. Enter<br/>set compression rtp f-max-period 500 queues q3</li> </ol>                                                                                                                                                  |



**Table 47: Configuring a Dialer Interface for Bandwidth-on-Demand (continued)**

| Task                                                                                                                                                                                                                                                                                                                                          | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                             | CLI Configuration Editor                                                                                                                                                                                                              |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure logical properties.<br><br>Maximum received reconstructed unit (MRRU) is expressed as a number between 1500 and 4500 bytes—for example, 1500.                                                                                                                                                                                       | <ol style="list-style-type: none"> <li>1. In the Fragment threshold box, type 1024.</li> <li>2. In the Mrru box, type 1500.</li> </ol>                                                                                                                                                                                                                                                                                                 | <ol style="list-style-type: none"> <li>1. Enter<br/>set fragment-threshold 1024</li> <li>2. Enter<br/>set mrru 1500</li> </ol>                                                                                                        |
| Configure PPP options.<br><br>You can also configure the following compression types:                                                                                                                                                                                                                                                         | <ol style="list-style-type: none"> <li>1. Next to Ppp options, click <b>Configure</b>.</li> <li>2. Next to Chap, click <b>Configure</b>.</li> <li>3. In the Access profile box, type <b>bw-profile</b>.</li> <li>4. Click <b>OK</b> and <b>OK</b> again.</li> <li>5. Under Compression, select <b>acfc</b>.</li> </ol>                                                                                                                 | <ol style="list-style-type: none"> <li>1. Enter<br/>edit ppp-options chap bw-profile</li> <li>2. Enter<br/>edit ppp-options compression acfc</li> </ol>                                                                               |
| <ul style="list-style-type: none"> <li>■ <b>acfc (address and control field compression)</b>—Conserves bandwidth by compressing the address and control fields of PPP-encapsulated packets.</li> <li>■ <b>pfc (protocol field compression)</b>—Conserves bandwidth by compressing the protocol field of a PPP-encapsulated packet.</li> </ul> |                                                                                                                                                                                                                                                                                                                                                                                                                                        |                                                                                                                                                                                                                                       |
| Configure the Family Inet properties.<br><br>You can also configure the Inet properties to use <b>unnumbered-address</b> with the source interface as lo-0/0/0 and then set an IP address—for example, 172.13.31.1, as the destination.                                                                                                       | <ol style="list-style-type: none"> <li>1. Next to Inet, select <b>Yes</b> and click <b>Configure</b>.</li> <li>2. Next to Negotiate address, select <b>Yes</b>.</li> <li>3. Select <b>Unnumbered address</b>, and then <b>Configure</b>.</li> <li>4. In the <b>Destination</b> box, type 172.31.13.1 as the destination.</li> <li>5. In the Source box, type lo-0/0/0 as the source interface.</li> <li>6. Click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>1. Enter<br/>set family inet negotiate-address</li> <li>2. To use the unnumbered-address option, enter<br/><br/>set family inet unnumbered address lo-0/0/0 destination 172.13.31.1</li> </ol> |

## 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 48.
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 192.

**Table 48: Configuring an ISDN Interface for Bandwidth-on-Demand**

| Task                                                                                                                                                                                                                                                                                                                                                                | J-Web Configuration Editor                                                                                                                                                                                                                                                                             | CLI Configuration Editor                                                         |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy.                                                                                                                                                                                                                                                                                             | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Configuration &gt; View and Edit &gt; Edit Configuration</b>.</li> <li>2. Next to Interfaces, click <b>Edit</b>.</li> <li>3. Next to the ISDN interface name, click <b>Edit</b>.</li> </ol>                 | <p>From the top of the configuration hierarchy, enter</p> <p>edit interfaces</p> |
| <p>Configure dialer options for each ISDN interface by following the instructions in Table 47.</p> <p>Each ISDN interface must have the same pool identifier to participate in bandwidth on demand.</p> <p>You can group up to four <b>br</b> interfaces together when configuring bandwidth-on-demand with a total of eight B-channels providing connectivity.</p> | <ol style="list-style-type: none"> <li>1. Next to the interface name, click <b>Dialer options</b>.</li> <li>2. Next to Pool, click <b>Add new entry</b>.</li> <li>3. In the Pool identifier box, type the name of the dialer pool—for example, <b>bw-pool</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> | <p>Enter</p> <p>edit interfaces br-1/0/3 dialer options pool bw-pool</p>         |

## Configuring Dial-on-Demand Routing (Optional)

Dial-on-demand routing (DDR) provides a way to link two sites over a public network and provide needed bandwidth by setting up an ISDN connection. The ISDN connections can provide secondary links to back up primary communication lines when they become overloaded or fail.

The dialer interface is configured as a passive static route with a lower priority than dynamic routes. When the dynamic route is lost, a packet destined for that IP address is received and the dialer interface initiates an ISDN connection and sends the packets over it. When no new packets are sent to the destination, the dialer interface initiates an inactivity timer and the ISDN connection is terminated when the timer expires.

### Configuring the Dial-on-Demand Dialer Filter

To configure dial-on-demand routing on the dialer 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 49.

3. Go on to “Applying the Dial-on-Demand Dialer Filter to the Dialer Interface” on page 185.

**Table 49: Configuring a Dialer Filter for Interesting Packets and Dial-on-Demand Routing**

| Task                                                                                                                                                                                                                                                                            | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                         | CLI Configuration Editor                                                                                    |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Firewall</b> level in the configuration hierarchy.                                                                                                                                                                                                           | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Configuration &gt; View and Edit &gt; Edit Configuration</b>.</li> <li>2. Next to Firewall, click <b>Edit</b>.</li> </ol>                                                                                                                                                                                                                                               | From the top of the configuration editor hierarchy, enter <code>edit firewall</code>                        |
| Configure the dialer filter name—for example, <code>ddr-packet</code> .                                                                                                                                                                                                         | <ol style="list-style-type: none"> <li>1. Next to <b>Inet</b>, click <b>Edit</b>.</li> <li>2. Next to <b>Dialer filter</b>, click <b>Add new entry</b>.</li> <li>3. In the Filter name box, type <code>ddr-packet</code>.</li> </ol>                                                                                                                                                                                                                               | From the <code>edit firewall</code> hierarchy, enter <code>edit family inet dialer-filter ddr-packet</code> |
| Configure the dialer filter—for example, <code>term1</code> .<br><br>Configure term behavior. For example, you might want to configure your interesting packet as an EBGp packet.<br><br>To configure the term completely, include both <b>from</b> and <b>then</b> statements. | <ol style="list-style-type: none"> <li>1. Next to Term, click <b>Add new entry</b>.</li> <li>2. In the Rule name box, type <code>term1</code></li> <li>3. Next to From, click <b>Configure</b>.</li> <li>4. From the Protocol choice list, select <b>Protocol</b>.</li> <li>5. Next to Protocol, click <b>Add new entry</b>.</li> <li>6. From the Value keyword list, select <b>ebgp</b>.</li> <li>7. Click <b>OK</b> twice to return to the Term page.</li> </ol> | <p>Enter</p> <p><code>set term term1 from protocol ebgp</code></p>                                          |
| Configure the <b>then</b> part of the dialer filter.                                                                                                                                                                                                                            | <ol style="list-style-type: none"> <li>1. Next to Then, click <b>Configure</b>.</li> <li>2. From the Designation list, select <b>Note</b>.</li> </ol>                                                                                                                                                                                                                                                                                                              | <p>Enter</p> <p><code>set term1 then note</code></p>                                                        |

## Applying the Dial-on-Demand Dialer Filter to the Dialer Interface

To complete dial-on-demand connectivity 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 50.
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 192.

**Table 50: Applying the Dialer Filter to the Dialer Interface**

| Task                                                                                | J-Web Configuration Editor                                                                                                                                                                                                                           | CLI Configuration Editor                                                      |
|-------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy.             | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Configuration &gt; View and Edit &gt; Edit Configuration</b>.</li> <li>2. Next to Interfaces, click <b>Edit</b>.</li> </ol>                               | From the top of the CLI configuration hierarchy, enter<br><br>edit interfaces |
| Select the dialer interface to apply the dialer filter—for example, <b>dl0</b> .    | <ol style="list-style-type: none"> <li>1. In the Interface name column, click <b>dl0</b>.</li> <li>2. Under Unit, in the Nested Configuration column, click <b>Edit</b>.</li> </ol>                                                                  | Enter<br><br>edit interfaces dl0 unit 0                                       |
| Select the dialer filter, <b>ddr-packet</b> , and apply it to the dialer interface. | <ol style="list-style-type: none"> <li>1. In the Family section, next to Inet, click <b>Edit</b>.</li> <li>2. Next to Filter, click <b>Configure</b>.</li> <li>3. In the Dialer box, type <b>ddr-packet</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> | Enter<br><br>edit family inet filter dialer ddr-packet                        |

### Configuring Dialer Watch (Optional)

Dialer watch is a feature that integrates backup dialing with routing capabilities and provides reliable connectivity without relying on “interesting” packets to trigger outgoing ISDN connections. With dialer watch, the Services Router monitors the existence of a specified route and if the route 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 51.
3. Go on to “Configuring the ISDN Interface for Dialer Watch” on page 189.

**Table 51: Adding a Dialer Watch Interface**

| <b>Task</b>                                                                                                                                                                                                                                                                                                                  | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                   | <b>CLI Configuration Editor</b>                                                                                                                |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy.                                                                                                                                                                                                                                                      | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Configuration &gt; View and Edit &gt; Edit Configuration</b>.</li> <li>2. Next to Interfaces, click <b>Configure</b> or <b>Edit</b>.</li> </ol>          | <p>From the top of the configuration hierarchy, enter</p> <pre>edit interfaces d10</pre>                                                       |
| Create the new interface—for example, d10.                                                                                                                                                                                                                                                                                   | <ol style="list-style-type: none"> <li>1. Next to Interface, click <b>Add new entry</b>.</li> </ol>                                                                                                                                                 | <p>Enter</p> <pre>set description dialer-watch</pre>                                                                                           |
| Adding a description, such as <b>dialer-watch</b> , can help you identify one dialer interface from another.                                                                                                                                                                                                                 | <ol style="list-style-type: none"> <li>2. In the Interface name box, type <b>d10</b>.</li> <li>3. In the Description box, type <b>dialer-watch</b>.</li> <li>4. Click <b>OK</b>.</li> </ol>                                                         |                                                                                                                                                |
| Configure encapsulation options.                                                                                                                                                                                                                                                                                             | <ol style="list-style-type: none"> <li>1. In the <b>Encapsulation</b> column, next to the interface, click <b>Edit</b>.</li> <li>2. From the Encapsulation list, select <b>cisco-hdlc</b>.</li> </ol>                                               | <p>Enter</p> <pre>set encapsulation cisco-hdlc</pre>                                                                                           |
| <ul style="list-style-type: none"> <li>■ <b>cisco-hdlc</b>—Cisco-compatible High-level Data Link Control is a group of protocols for transmitting data between network points.</li> <li>■ <b>ppp</b>—Point-to-Point Protocol is a protocol used for communication between two computers using a serial interface.</li> </ul> |                                                                                                                                                                                                                                                     |                                                                                                                                                |
| Set a hold-time value, in milliseconds, to be used when negotiating a connection with the peer—for example, <b>60</b> . The hold time is three times the interval at which keepalive messages are sent.                                                                                                                      | <ol style="list-style-type: none"> <li>1. Under the Hold time section, type <b>60</b> in the <b>Down</b> box.</li> <li>2. In the Up box, type <b>60</b>.</li> </ol>                                                                                 | <ol style="list-style-type: none"> <li>1. Enter <pre>set hold-time down 60</pre> </li> <li>2. Enter <pre>set hold-time up 60</pre> </li> </ol> |
| Create the logical unit properties—for example, <b>0</b> .                                                                                                                                                                                                                                                                   | <ol style="list-style-type: none"> <li>1. Next to Unit, click <b>Add new entry</b>.</li> <li>2. In the Interface unit number box, type <b>0</b>.</li> <li>3. Next to Dialer options, select <b>Yes</b>, and then click <b>Configure</b>.</li> </ol> | <p>Enter</p> <pre>edit interfaces d10 unit 0</pre>                                                                                             |

**Table 51: Adding a Dialer Watch Interface (continued)**

| Task                                                                                                                                                                                                                                                                                                               | J-Web Configuration Editor                     | CLI Configuration Editor                       |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------|------------------------------------------------|
| Configure dialer options.                                                                                                                                                                                                                                                                                          | 1. In the Activation delay box, type 60.       | 1. Enter<br>edit dialer-options                |
| <ul style="list-style-type: none"> <li>■ <b>Activation delay</b>—Time, in seconds, 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.</li> </ul>       | 2. In the Deactivation delay box, type 60.     | 2. Enter<br>set activation-delay 60            |
|                                                                                                                                                                                                                                                                                                                    | 3. In the Dialer string box, type 18005555555. | 3. Enter<br>set deactivation-delay 60          |
|                                                                                                                                                                                                                                                                                                                    | 4. In the Idle timeout box, type 30.           | 4. Enter<br>set dial-string 18005555555        |
| <ul style="list-style-type: none"> <li>■ <b>Deactivation delay</b>—Time, in seconds, to wait before deactivating the backup interface once the primary interface is up—for example, 30. The default value is 0 seconds with a maximum value of 60 seconds. Use only for dialer backup and dialer watch.</li> </ul> | 5. In the Initial route check box, type 30.    | 5. Enter<br>set dialer-options idle-timeout 30 |
|                                                                                                                                                                                                                                                                                                                    | 6. In the Pool box, type dw-group.             | 6. Enter<br>set initial-route-check 30         |
| <ul style="list-style-type: none"> <li>■ <b>Dialer string</b>—Telephone number for the interface to dial that establishes ISDN connectivity—for example, 8005555555.</li> </ul>                                                                                                                                    |                                                | 7. Enter<br>set pool dw-group                  |
| <ul style="list-style-type: none"> <li>■ <b>Idle timeout</b>—Time a connection is idle before disconnecting—for example, 30. Default value is 120 seconds with a range from 0 to 4294967295.</li> </ul>                                                                                                            |                                                |                                                |
| <ul style="list-style-type: none"> <li>■ <b>Initial route check</b>—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.</li> </ul>                                                                                         |                                                |                                                |
| <ul style="list-style-type: none"> <li>■ <b>Pool</b>—Name of a group of ISDN interfaces configured to use the dialer interface—for example, dw-group.</li> </ul>                                                                                                                                                   |                                                |                                                |

**Table 51: Adding a Dialer Watch Interface (continued)**

| <b>Task</b>                                                                | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                           | <b>CLI Configuration Editor</b>                                                            |
|----------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------|
| Configure the list of routes for dialer watch—for example, 172.27.27.0/24. | <ol style="list-style-type: none"> <li>1. Next to Watch list, click <b>Add new entry</b>.</li> <li>2. In the Prefix box, type 172.27.27.0/24.</li> </ol>                                                                                                                    | Enter<br><br>set watch-list 172.27.27.0/24                                                 |
| Configure an IP address for the dialer interface—for example, 10.1.1.2/24. | <ol style="list-style-type: none"> <li>1. Under Family, next to Inet, select <b>Yes</b>, and then click <b>Configure</b>.</li> <li>2. Next to Address, click <b>Add new entry</b>.</li> <li>3. In the Source box, type 10.1.1.2/24.</li> <li>4. Click <b>OK</b>.</li> </ol> | From the [edit interfaces dl0 unit 0] hierarchy, enter<br><br>edit family inet 10.1.1.2/24 |

### 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 52.
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 192.

**Table 52: Configuring an ISDN Interface for Dialer Watch**

| <b>Task</b>                                                                                                                                                                                      | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                      | <b>CLI Configuration Editor</b>                                                                                            |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy.                                                                                                                          | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Configuration &gt; View and Edit &gt; Edit Configuration</b>.</li> <li>2. Next to Interfaces, click <b>Edit</b>.</li> <li>3. Next to the ISDN interface name, click <b>Edit</b>.</li> </ol> | <p>From the top of the configuration hierarchy, enter</p> <pre>edit interfaces br-1/0/3 dialer options pool dw-group</pre> |
| Configure dialer watch options for each ISDN interface participating in the dialer watch feature.                                                                                                | <ol style="list-style-type: none"> <li>1. Next to the interface name, click <b>Dialer options</b>.</li> <li>2. Next to Pool, click <b>Add new entry</b>.</li> </ol>                                                                                                                    |                                                                                                                            |
| Each ISDN interface must have the same pool identifier to participate in dialer watch. Therefore, the dialer watch interface configured in Table 51 is used when configuring the ISDN interface. | <ol style="list-style-type: none"> <li>3. In the Pool identifier box, type <b>dw-group</b>.</li> <li>4. Click <b>OK</b>.</li> </ol>                                                                                                                                                    |                                                                                                                            |

### **Configuring Dial-on-Demand Routing 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 with OSPF support. For information on configuring OSPF, see “Configuring an OSPF Network” on page 265.

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 53.
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 192.



**Table 53: Configuring OSPF Demand Circuits**

| Task                                                                                                                             | J-Web Configuration Editor                                                                                                                                                                                                                                                       | CLI Configuration Editor                                                                                 |
|----------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Protocols</b> level in the configuration hierarchy.                                                           | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Configuration &gt; View and Edit &gt; Edit Configuration</b>.</li> <li>2. Next to Protocols, click <b>Edit</b>.</li> <li>3. Next to the OSPF, click <b>Edit</b>.</li> </ol>           | <p>From the top of the configuration editor hierarchy, enter</p> <p>edit protocols OSPF area 0.0.0.0</p> |
| Configure OSPF on-demand circuits for each ISDN interface participating as an on-demand routing interface—for example, br-5/0/0. | <ol style="list-style-type: none"> <li>1. Next to the Area id, click <b>Edit</b>.</li> <li>2. Next to Interfaces, click <b>Add new entry</b>.</li> <li>3. In the Interface box, type br-5/0/0.</li> <li>4. Select <b>Demand circuit</b>.</li> <li>5. Click <b>OK</b>.</li> </ol> | <p>Enter</p> <p>edit interface br-5/0/0</p> <p>Enter</p> <p>set demand-circuit</p>                       |

### Configuring Dialer Profiles (Optional)

You can configure multiple dialer interfaces to participate as part of a dialer profile. After configuring dialer interfaces, you configure an ISDN interface on the Services Router to participate as part of a dialer profile. In the configuration in Table 54, dialer interfaces dl0 and dl1 and ISDN interface br-1/0/3 are used as examples.

To configure an 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 54.
3. When you are finished configuring the router, commit the configuration.

**Table 54: Dialer Profile Configuration**

| Task                                                                    | J-Web Configuration Editor                                                                                                                                                                                                                                                                     | CLI Configuration Editor                                                                           |
|-------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy. | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Configuration &gt; View and Edit &gt; Edit Configuration</b>.</li> <li>2. Next to Interfaces, click <b>Configure</b> or <b>Edit</b>.</li> <li>3. Under Interface name, click <b>dl0</b>.</li> </ol> | <p>From the top of the configuration editor hierarchy, enter</p> <p>edit interfaces dl0 unit 0</p> |

**Table 54: Dialer Profile Configuration (continued)**

| <b>Task</b>                                                                                                                                                                                                                                                                                                                                                    | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Add a dialer pool, <b>pool1</b> , to a dialer interface.                                                                                                                                                                                                                                                                                                       | <ol style="list-style-type: none"> <li>1. In the Unit table, click <b>Dialer options</b>.</li> <li>2. In the Pool box, type <b>pool1</b>.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                     | Enter<br><br>set dialer-options pool pool1                                                                                                                                                                                                                           |
| Configure a source IP address—for example, <b>172.20.10.1</b> for the dialer interface.<br><br>Configure a destination IP address—for example, <b>172.20.10.2</b> for the dialer interface.                                                                                                                                                                    | <ol style="list-style-type: none"> <li>1. Select <b>Inet</b> under Family, and click <b>Edit</b>.</li> <li>2. Next to Address, click <b>Add new entry</b>.</li> <li>3. In the Source box, type <b>172.20.10.1</b>.</li> <li>4. In the Destination box, type <b>172.20.10.2</b>.</li> <li>5. Click <b>OK</b> until you return to the Interfaces page.</li> </ol>                                                                                                       | <ol style="list-style-type: none"> <li>1. Enter<br/><br/>               set family inet address 172.20.10.1</li> <li>2. Enter<br/><br/>               set family inet address destination 172.20.10.2</li> </ol>                                                     |
| Configure the ISDN interface—for example, <b>br-1/0/3</b> , with a dialer profile that uses either dialer interface to initiate an ISDN connection.<br><br><b>Priority</b> has a range from 0 to 255 with 255 having the highest priority.<br><br>The <b>br-1/0/3</b> interface now uses <b>pool2</b> to establish connectivity first, and then <b>pool1</b> . | <ol style="list-style-type: none"> <li>1. Next to <b>br-1/0/3</b>, click Dialer options.</li> <li>2. Next to Pool, click <b>Add new entry</b>.</li> <li>3. In the Pool identifier box, type <b>pool1</b>.</li> <li>4. In the Priority box, type <b>10</b>.</li> <li>5. Click <b>OK</b>.</li> <li>6. Click <b>Add new entry</b> again, and add <b>pool2</b> to the interface.</li> <li>7. In the Priority box, type <b>25</b>.</li> <li>8. Click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>1. Enter<br/><br/>               edit interfaces br-1/0/3 dialer-options</li> <li>2. Enter<br/><br/>               set pool pool1 priority 10</li> <li>3. Enter<br/><br/>               set pool pool2 priority 25</li> </ol> |

## Verifying the ISDN Configuration

To verify an ISDN configuration, perform the following tasks:

- Displaying the ISDN Status on page 193
- Verifying an ISDN Interface on page 193
- Checking B-Channel Statistics on page 194
- Checking D-Channel Interface Statistics on page 196

- Verifying Dialer Interface Configuration on page 197

## Displaying the ISDN Status

|                      |                                                                                                                                                                                                                                                                                                                                                     |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b>       | Display the status of ISDN parameters on the ISDN interface. For example, you can display ISDN parameters on the <code>br-6/0/0</code> interface.                                                                                                                                                                                                   |
| <b>Action</b>        | From the operational mode in the CLI, enter <code>show isdn status</code> .                                                                                                                                                                                                                                                                         |
| <b>Sample Output</b> | <pre>user@host&gt; show isdn status  Interface:br-6/0/0   Layer 1 status: active   Layer 2 status: Q.921: up, TEI:12   Layer 3 status: 1 Active calls     Switch Type           = ETSI     Interface Type        = USER     T306                   = 10 seconds     T310                   = 10 seconds     Tei Option             = Power Up</pre> |
| <b>What It Means</b> | The output shows a summary of interface information. Verify the following information:                                                                                                                                                                                                                                                              |

- **Interface**—ISDN interface currently active on the Services Router.
- **Layer 1 status**—Displays as active or inactive.
- **Layer 2 status**—Displays Q.921 as up or down.
- **TEI**—Displays the assigned TEI number.
- **Layer 3 status**—Displays the 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**—Displays the telephone number configured for dial out.
- **T306 and T310**—Q.931 specific timers.
- **TEI Option**—Determines when TEI negotiations occur on the interface.

## Verifying an ISDN Interface

|                      |                                                                         |
|----------------------|-------------------------------------------------------------------------|
| <b>Purpose</b>       | Verify that the ISDN interface is correctly configured.                 |
| <b>Action</b>        | From the CLI, enter the <code>show interfaces extensive</code> command. |
| <b>Sample Output</b> | <pre>user@host&gt; show interfaces br-6/0/0 extensive</pre>             |

```

Physical interface: br-6/0/0, Enabled, Physical link is Up
 Interface index: 144, SNMP ifIndex: 27, Generation: 25
 Type: Serial, Link-level type: Controller, MTU: 4092, Clocking: Internal, Speed: 144kbps
 Parent: None
 Device flags : Present Running
 Interface flags: Point-To-Point 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
 Last flapped : 2005-04-25 22:03:53 UTC (02:12:33 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, 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

```

**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.

## Checking B-Channel Statistics

**Purpose** Verify that the ISDN B-channel interface is correctly configured.

**Action** From the CLI, enter the show interfaces extensive command.

**Sample Output** user@host> **show interfaces bc-0/0/4 extensive**

```

Physical interface: bc-0/0/4:1, Administratively down, Physical link is Up
Interface index: 144, SNMP ifIndex: 51, Generation: 25
Type: Serial, Link-level type: Multilink-PPP, MTU: 1510, Clocking: Internal,
Speed: 64kbps,
Parent: br-0/0/4 Interface index 142
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
Last flapped : Never
Statistics last cleared: Never
Traffic statistics:
Input bytes : 149289226 0 bps
Output bytes : 166219636 0 bps
Input packets: 278442 0 pps
Output packets: 319599 0 pps
Input errors:
Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Giants: 0, Policed discards: 2481,
Resource errors: 44
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 314335 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 5264 5264 0
Packet Forwarding Engine configuration:
Destination slot: 0, PLP byte: 1 (0x00)
CoS transmit queue Bandwidth Buffer Priority Limit
 % bps % bytes
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 72) (SNMP ifIndex 55) (Generation 19)
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).
  - 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.

## Checking D-Channel Interface Statistics

**Purpose** Verify that the ISDN D-channel interface is correctly configured.

**Action** From the CLI, enter the show interfaces extensive command.

**Sample Output** user@host> **show interfaces dc-0/0/4 extensive**

```
Physical interface: dc-0/0/4, Enabled, Physical link is Up
Interface index: 143, SNMP ifIndex: 49, Generation: 24
Type: Serial, Link-level type: 55, MTU: 1504, Clocking: Internal, Speed: 16kbps,
Parent: br-0/0/4 Interface index 142
Device flags : Present Running
Interface flags: 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
Last flapped : 2005-04-23 13:07:53 PDT (2d 05:25 ago)
Statistics last cleared: Never
Traffic statistics:
Input bytes : 58975 0 bps
Output bytes : 73967 0 bps
Input packets : 14674 0 pps
Output packets : 14685 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: 5, Errors: 0, Drops: 0, MTU errors: 0, Resource errors: 0
ISDN alarms : None
ISDN media:
 Seconds Count State
LOF 0 2 OK
LOS 0 1 OK

Logical interface dc-0/0/4.32767 (Index 71) (SNMP ifIndex 50) (Generation 9)
Flags: Point-To-Point SNMP-Traps Encapsulation: 60
Traffic statistics:
Input bytes : 58975
Output bytes : 59282
Input packets: 14674
Output packets: 14685
Local statistics:
Input bytes : 58975
Output bytes : 59282
Input packets: 14674
Output packets: 14685

```

**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).
- 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.

## Verifying Dialer Interface Configuration

**Purpose** Verify that the dialer interface is correctly configured.

**Action** From the CLI, enter the `show interfaces extensive` command.

**Sample Output** `user@host> show interfaces dl0 extensive`

```

Physical interface: dl0, Enabled, Physical link is Up
Interface index: 146, SNMP ifIndex: 53, Generation: 27

```

```

Link-level type: LinkService, MTU: 1504
Device flags : Present Running
Interface flags: Point-To-Point SNMP-Traps
Last flapped : 2005-04-23 00:24:29 PDT (2d 18:04 ago)
Link flags : No-Keepalives ACFC PFC
Statistics last cleared: Never
Traffic statistics:
 Input bytes : 293333987 0 bps
 Output bytes : 328418548 0 bps
 Input packets : 365724 0 pps
 Output packets: 628595 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
Queue counters: Queued packets Transmitted packets Dropped packets
0 best-effort 628595 628595 0
1 expedited-fo 0 0 0
2 assured-forw 0 0 0
3 network-cont 0 0 0

Logical interface dl0.0 (Index 66) (SNMP ifIndex 54) (Generation 14)
Flags: Hardware-Down Point-To-Point SNMP-Traps Encapsulation: Multilink-PPP
Bandwidth: 128kbps
Bundle options:
 MRRU 1504
 Drop timer period 0
 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 202 (121378 bytes)
 MRRU exceeded 0
 Exception events 0
Statistics Frames fps Bytes bps
Bundle:
 Fragments:
 Input : 557280 0 297042248 0
 Output: 628583 0 332189058 0
 Packets:
 Input : 365718 0 293333755 0
 Output: 628581 0 328417288 0
Link:
 bc-0/0/4:1.0
 Input : 278289 0 149287251 0

```



```

Output: 314326 0 166146384 0
bc-0/0/4:2.0
Input : 278980 0 147754846 0
Output: 314257 0 166042674 0
Protocol inet, MTU: 1500, Generation: 15, Route table: 0
Flags: None
Addresses, Flags: Dest-route-down Is-Preferred Is-Primary
Destination: 10.2.0/24, Local: 10.2.0.1, Broadcast: Unspecified, Generation: 19

```

**What It Means**

The output shows a summary of dialer 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 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.

For complete descriptions of the interface output, see the *JUNOS Network and Services Interfaces Command Reference*.



## **Part 3**

# **Configuring Routing Protocols**

- Routing Overview on page 203
- Configuring Static Routes on page 237
- Configuring a RIP Network on page 249
- Configuring an OSPF Network on page 265
- Configuring the IS-IS Protocol on page 287
- Configuring BGP Sessions on page 295



## Chapter 6

# 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 203
- Routing Overview on page 207
- RIP Overview on page 213
- OSPF Overview on page 217
- IS-IS Overview on page 222
- BGP Overview on page 225

## Routing Terms

---

To understand routing, become familiar with the terms defined in Table 55 .

**Table 55: Routing Terms**

| <b>Term</b>                            | <b>Definition</b>                                                                                                                                                                                                                                                                                             |
|----------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>adjacency</b>                       | Portion of the local routing information that pertains to the reachability of a single neighbor over a single circuit or interface.                                                                                                                                                                           |
| <b>area</b>                            | 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.                                     |
| <b>area border router (ABR)</b>        | 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.                               |
| <b>AS path</b>                         | In BGP, the list of autonomous system (ASs) that a packet must traverse to reach a given set of destinations within a single AS.                                                                                                                                                                              |
| <b>autonomous system (AS)</b>          | Network, collection of routers, or portion of a large internetwork under a single administrative authority.                                                                                                                                                                                                   |
| <b>backbone area</b>                   | 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.                                                                                                                                |
| <b>bidirectional connectivity</b>      | Ability of directly connected devices to communicate with each other over the same link.                                                                                                                                                                                                                      |
| <b>Border Gateway Protocol (BGP)</b>   | Exterior gateway protocol used to exchange routing information among routers in different autonomous systems.                                                                                                                                                                                                 |
| <b>broadcast</b>                       | Operation of sending network traffic from one network node to all other network nodes.                                                                                                                                                                                                                        |
| <b>cluster</b>                         | 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. |
| <b>confederation</b>                   | In BGP, a group of autonomous systems (ASs) that appears to external ASs to be a single AS.                                                                                                                                                                                                                   |
| <b>confederation sequence</b>          | Ordered set of autonomous systems (ASs) for a confederation. The closest AS in the path is first in the sequence.                                                                                                                                                                                             |
| <b>convergence</b>                     | After a topology change, the time all the routers in a network take to receive the information and update their routing tables.                                                                                                                                                                               |
| <b>cost</b>                            | 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.                                               |
| <b>designated router (DR)</b>          | In OSPF, a node designated to process link-state advertisements (LSAs) and distribute topology updates for an autonomous system (AS).                                                                                                                                                                         |
| <b>distance vector</b>                 | Number of hops to a routing destination.                                                                                                                                                                                                                                                                      |
| <b>dynamic routing</b>                 | Routing method that enables the route of a message through a network to change as network conditions change. Compare <i>static routing</i> .                                                                                                                                                                  |
| <b>end systems</b>                     | Network entities that send and receive packets.                                                                                                                                                                                                                                                               |
| <b>exterior gateway protocol (EGP)</b> | Protocol that exchanges routing information between autonomous systems (ASs). BGP is an EGP. Compare <i>interior gateway protocol (IGP)</i> .                                                                                                                                                                 |
| <b>external BGP (EBGP)</b>             | BGP configuration in which sessions are established between routers in different autonomous systems (ASs).                                                                                                                                                                                                    |
| <b>external peer</b>                   | In BGP, a peer that resides in a different autonomous system (AS) from the Services Router.                                                                                                                                                                                                                   |
| <b>external route</b>                  | Route to an area outside the network.                                                                                                                                                                                                                                                                         |

**Table 55: Routing Terms (continued)**

| <b>Term</b>                                               | <b>Definition</b>                                                                                                                                                                                                                                                                                                                    |
|-----------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>flooding</b>                                           | 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.                     |
| <b>forwarding table</b>                                   | 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. |
| <b>full mesh</b>                                          | Network in which devices are organized in a mesh topology, with each node connected to every other network node.                                                                                                                                                                                                                     |
| <b>gateway router</b>                                     | Node on a network that serves as an entrance to another network.                                                                                                                                                                                                                                                                     |
| <b>global AS</b>                                          | Global autonomous system (AS). An AS consisting of multiple subautonomous systems (sub-ASs).                                                                                                                                                                                                                                         |
| <b>handshake</b>                                          | Process of exchanging signaling information between two communications devices to establish the method and transmission speed of a connection.                                                                                                                                                                                       |
| <b>hello packet</b>                                       | In OSPF, a packet sent periodically by a router to first establish and then maintain network adjacency, and to discover neighbor routers.                                                                                                                                                                                            |
| <b>hold time</b>                                          | Maximum number of seconds allowed to elapse between the time a BGP system receives successive keepalive or update messages from a peer.                                                                                                                                                                                              |
| <b>hop</b>                                                | 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.                                                                      |
| <b>intermediate systems</b>                               | Network entities that relay (forward) packets as well as send and receive them on the network. Intermediate systems are also known as routers.                                                                                                                                                                                       |
| <b>Intermediate System-to-Intermediate System (IS-IS)</b> | Link-state, interior gateway routing protocol for IP networks that also uses the shortest-path-first (SPF) algorithm to determine routes.                                                                                                                                                                                            |
| <b>interior gateway protocol (IGP)</b>                    | Protocol that exchanges routing information within autonomous systems (ASs). IS-IS, OSPF, and RIP are IGPs. Compare <i>exterior gateway protocol (EGP)</i> .                                                                                                                                                                         |
| <b>Internal BGP (IBGP)</b>                                | BGP configuration in which sessions are established between routers in the same autonomous systems (ASs).                                                                                                                                                                                                                            |
| <b>internal peer</b>                                      | In BGP, a peer that resides in the same autonomous system (AS) as the Services Router.                                                                                                                                                                                                                                               |
| <b>keepalive message</b>                                  | Periodic message sent by one BGP peer to another to verify that the session between them is still active.                                                                                                                                                                                                                            |
| <b>latency</b>                                            | Delay that occurs when a packet or signal is transmitted over a communications system.                                                                                                                                                                                                                                               |
| <b>link-state advertisement (LSA)</b>                     | Messages that announce the presence of OSPF-enabled interfaces to adjacent OSPF interfaces (neighbors). The exchange of LSAs establishes bidirectional connectivity between neighbors.                                                                                                                                               |
| <b>local preference</b>                                   | Optional BGP path attribute carried in internal BGP update packets that indicates the degree of preference for an external route.                                                                                                                                                                                                    |
| <b>mesh</b>                                               | 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> .                                                                                                                                                    |
| <b>metric</b>                                             | Numerical value that determines how quickly a packet can reach its destination. See also <i>cost</i> .                                                                                                                                                                                                                               |
| <b>multiple exit discriminator (MED)</b>                  | 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.                                                                                                                                                   |

**Table 55: Routing Terms (continued)**

| <b>Term</b>                                     | <b>Definition</b>                                                                                                                                                                                                                                                                                                                             |
|-------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>neighbor</b>                                 | Adjacent router interface. A node can directly route packets to its neighbors only. See also <i>peer</i> .                                                                                                                                                                                                                                    |
| <b>network</b>                                  | Series of nodes interconnected by communication paths.                                                                                                                                                                                                                                                                                        |
| <b>network diameter</b>                         | Maximum hop count in a network.                                                                                                                                                                                                                                                                                                               |
| <b>network topology</b>                         | Arrangement of nodes and connections in a network.                                                                                                                                                                                                                                                                                            |
| <b>node</b>                                     | Connection point that operates as a redistribution point or an end point in a network, recognizing data transmissions and either forwarding or processing them.                                                                                                                                                                               |
| <b>notification message</b>                     | 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.                                                                                                                                                                             |
| <b>not-so-stubby area (NSSA)</b>                | In OSPF, a type of stub area in which external route advertisements can be flooded.                                                                                                                                                                                                                                                           |
| <b>open message</b>                             | Message sent between BGP peers to establish communication.                                                                                                                                                                                                                                                                                    |
| <b>Open Shortest Path First protocol (OSPF)</b> | 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).                                                                                                                                                                      |
| <b>origin</b>                                   | 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.                                                                                                                                                                                           |
| <b>path-vector protocol</b>                     | 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.                                        |
| <b>peer</b>                                     | Immediately adjacent router with which a protocol relationship has been established. See also <i>neighbor</i> .                                                                                                                                                                                                                               |
| <b>peering</b>                                  | The practice of exchanging Internet traffic with directly connected peers according to commercial and contractual agreements.                                                                                                                                                                                                                 |
| <b>point of presence (POP)</b>                  | 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.                                                                                                                                                     |
| <b>poison reverse</b>                           | 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> . |
| <b>propagation</b>                              | 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.                                                                                                                                         |
| <b>reachability</b>                             | In BGP, the feasibility of a route.                                                                                                                                                                                                                                                                                                           |
| <b>round-robin</b>                              | Scheduling algorithm in which items have the same priority and are handled in a fixed cyclic order.                                                                                                                                                                                                                                           |
| <b>route advertisement</b>                      | 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> .                                                                                                               |
| <b>route aggregation</b>                        | 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.                                                                                                                                                            |
| <b>route reflection</b>                         | 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.                                                                                                      |



**Table 55: Routing Terms (continued)**

| Term                                      | Definition                                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Routing Information Protocol (RIP)</b> | Distance-vector routing protocol that keeps a database of routing information gathered from periodic broadcasts by each router in a network.                                                                                                                                                                                                                                  |
| <b>routing table</b>                      | 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.                                                                                                                                                                                                |
| <b>split horizon</b>                      | 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> .                                                                                      |
| <b>static routing</b>                     | 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> . |
| <b>stub area</b>                          | In OSPF, an area through which or into which autonomous system (AS) external route advertisements are not flooded.                                                                                                                                                                                                                                                            |
| <b>subautonomous system (sub-AS)</b>      | Autonomous system (AS) members of a BGP confederation.                                                                                                                                                                                                                                                                                                                        |
| <b>subnetwork</b>                         | Subdivision of a network, which functions exactly like a network except that it has a more specific address and subnet mask (destination prefix).                                                                                                                                                                                                                             |
| <b>three-way handshake</b>                | Process by which two routers synchronize protocols and establish a bidirectional connection.                                                                                                                                                                                                                                                                                  |
| <b>topology database</b>                  | Map of connections between the nodes in a network. The topology database is stored in each node.                                                                                                                                                                                                                                                                              |
| <b>triggered update</b>                   | In a network that uses RIP, a routing update that is automatically sent whenever routing information changes.                                                                                                                                                                                                                                                                 |
| <b>virtual link</b>                       | 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 208
- Autonomous Systems on page 208
- Interior and Exterior Gateway Protocols on page 208
- Routing Tables on page 209

- Forwarding Tables on page 209
- Dynamic and Static Routing on page 210
- Route Advertisements on page 211
- Route Aggregation on page 211

## **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 31 shows a simple network of routers.

**Figure 31: 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 31 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 31, 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 32 shows a network that uses static routes.

**Figure 32: Static Routing Example**




In Figure 32, 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 33.

**Figure 33: Route Advertisement**




In Figure 33, 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 34.

**Figure 34: Route Aggregation**



9015004

Figure 34 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 213
- Maximizing Hop Count on page 214
- RIP Packets on page 215
- Split Horizon and Poison Reverse Efficiency Techniques on page 215
- Limitations of Unidirectional Connectivity on page 216

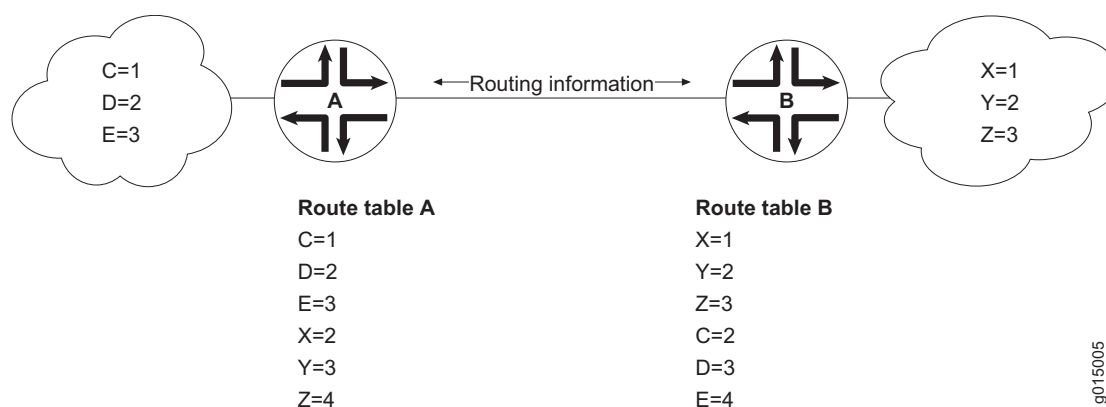


**NOTE:** The J-series Services Router supports both RIP version 1 and RIP version 2. In this guide, the term RIP refers to both versions of the protocol.

---

## 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 35 shows how distance-vector routing works.

**Figure 35: Distance-Vector Protocol**

In Figure 35, 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 36 shows an example of the split horizon technique.


**Figure 36: Split Horizon Example**



In Figure 36, 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 37 shows an example of the poison reverse technique.


**Figure 37: Poison Reverse Example**



In Figure 37, 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 38 shows, RIP networks are limited by their unidirectional connectivity.

**Figure 38: Limitations of Unidirectional Connectivity**

In Figure 38, 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 218.

## 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 218
- Role of the Designated Router on page 218
- Path Cost Metrics on page 219
- Areas and Area Border Routers on page 219
- Role of the Backbone Area on page 220
- Stub Areas and Not-So-Stubby Areas on page 221




**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 39.

**Figure 39: OSPF Three-Way Handshake**



In Figure 39, 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 40 shows an OSPF topology of three areas connected by two area border routers.


**Figure 40: 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 40, 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 41 shows a virtual link between a noncontiguous area and the backbone area through an area connected to both.

**Figure 41: OSPF Topology with a Virtual Link**

In the topology shown in Figure 41, 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 42 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 42: OSPF AS Network with Stub Areas and NSSAs**

g015012

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 42 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 42, 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 223
- Network Entity Titles and System Identifiers on page 223
- IS-IS Path Selection on page 224
- Protocol Data Units on page 224

## **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 IGPs 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 226
- BGP Messages for Session Establishment on page 226
- BGP Messages for Session Maintenance on page 227
- IBGP and EBGP on page 227
- Route Selection on page 228
- Local Preference on page 228
- AS Path on page 229

- Origin on page 230
- Multiple Exit Discriminator on page 230
- Scaling BGP for Large Networks on page 231

### 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 43 shows an example of a BGP peering session.

**Figure 43: BGP Peering Session**



In Figure 43, 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 EBG**

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 EBG 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 EBG.

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 231. For information about routing confederations, see “Scaling BGP for Large Networks” on page 231.

## **Route Selection**


A local BGP router uses the following primary criteria to select a route from the routing table for the forwarding table:

1. **Next-hop accessible**—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 228.)
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 229.)
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 230.)
5. **Lowest MED value**—The local router selects the route with the lowest multiple exit discriminator (MED) value. If multiple routes have the same MED value, route selection continues. (For more information, see “Multiple Exit Discriminator” on page 230.)

If more than one route remains after all these criteria are evaluated, the local BGP router evaluates a set of secondary criteria to select the single route to a destination for its forwarding table. The secondary criteria include whether the route was learned through an EBGp or Ibgp, the Igp route metric, and the router ID.

## **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 44 illustrates how to use local preference to determine BGP route selection.

**Figure 44: Local Preference**


The network in Figure 44 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 45 shows how BGP creates an AS path.

**Figure 45: BGP AS Path**

In the network shown in Figure 45, 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

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 neighbor 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 46 illustrates how to use an MED metric to determine route selection.




**Figure 46: MED Example**

Figure 46 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, and 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.

## 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 231
- Confederations—for Subdivision on page 234

### 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 47.



**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 47: Simple Route Reflector Topology (One Cluster)**




Figure 47 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 48).


**Figure 48: Basic Route Reflection (Multiple Clusters)**

Figure 48 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 49).

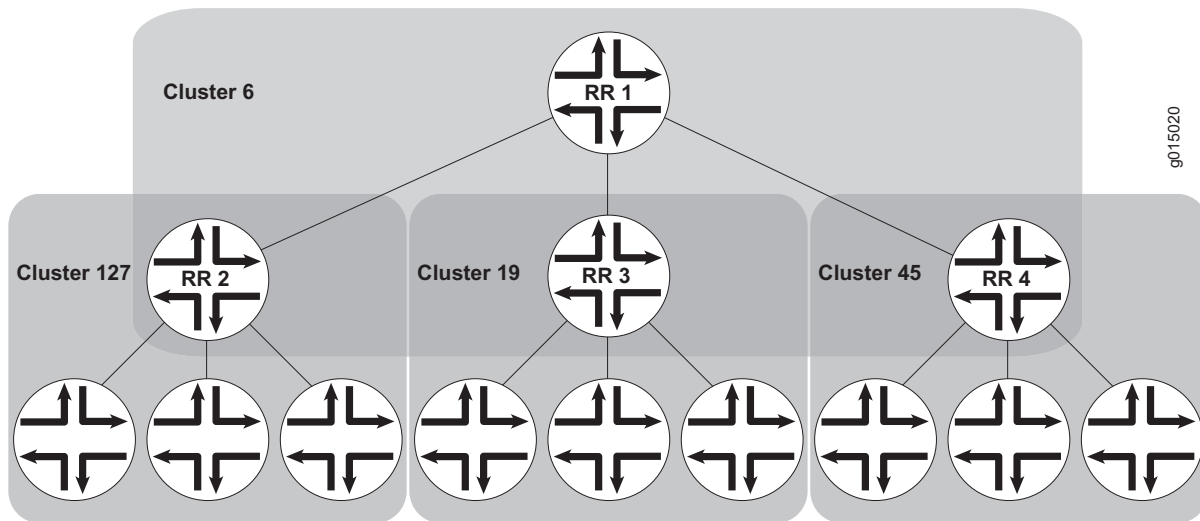
**Figure 49: Hierarchical Route Reflection (Clusters of Clusters)**

Figure 49 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 50 shows an AS divided into four confederations.


**Figure 50: BGP Confederations**

Figure 50 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 7

# 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 237
- Before You Begin on page 240
- Configuring Static Routes with Quick Configuration on page 240
- Configuring Static Routes with a Configuration Editor on page 242
- Verifying the Static Route Configuration on page 247

## 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 238
- Qualified Next Hops on page 238
- Control of Static Routes on page 238
- Default Properties on page 239

## 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 239.
- **no-readvertise**—Prevents the route from being readvertised to other routing protocols. For more information, see “Readvertisement Prevention” on page 239.
- **passive**—Rejects traffic destined for the route. For more information, see “Forced Rejection of Passive Route Traffic” on page 239.



## 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 *J-series Services Router Getting Started Guide*.
- Configure network interfaces. See “Configuring Network Interfaces” on page 103.

## Configuring Static Routes with Quick Configuration

J-Web Quick Configuration allows you to configure static routes. Figure 51 shows the Quick Configuration Routing page for static routing.

**Figure 51: Quick Configuration Routing Page for Static Routing**

Juniper NETWORKS

ROUTER - J4300

Logged in as: **regress**

[Help](#) [About](#) [Logout](#)

**Monitor** **Configuration** **Diagnose** **Manage**

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

**Quick Configuration**

Set Up

SSL

Interfaces

Users

SNMP

**Routing**

Firewall/NAT

IPSec Tunnels

Realtime Performance Monitoring

► **View and Edit**

► **History**

► **Rescue**

**Quick Configuration**

**Routing**

**Default Route**

**Default Route**

**Static Routes**

|                          | Static Route Address              | Next Hop        |
|--------------------------|-----------------------------------|-----------------|
| <input type="checkbox"/> | <a href="#">10.74.10.0/24</a>     |                 |
| <input type="checkbox"/> | <a href="#">172.16.0.0/12</a>     | 192.168.124.254 |
| <input type="checkbox"/> | <a href="#">192.168.0.0/18</a>    | 192.168.124.254 |
| <input type="checkbox"/> | <a href="#">192.168.64.0/18</a>   | 192.168.124.254 |
| <input type="checkbox"/> | <a href="#">207.17.136.192/32</a> | 192.168.124.254 |
| <input type="checkbox"/> | <a href="#">192.168.40.0/22</a>   | 192.168.124.254 |

To configure static routes with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Routing > Static Routing**.
2. Enter information into the Static Routing Quick Configuration page, as described in Table 56.
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 247.

**Table 56: Static Routing Quick Configuration Summary**

| Field                           | Function                                                                                         | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------|--------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Default Route</b>            |                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 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.                                                                                                                                                                                                                                                                                                                                                                   |
| <b>Static Routes</b>            |                                                                                                  |                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Static Route Address (required) | Specifies the static route to add to the routing table.                                          | <ol style="list-style-type: none"> <li>1. On the main static routing Quick Configuration page, click <b>Add</b>.</li> <li>2. In the Static Route Address box, type the 32-bit IP address of the static route in dotted decimal notation.</li> </ol>                                                                                                                                                                                                             |
| 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"> <li>1. In the Add box, type the 32-bit IP address of the next-hop host.</li> <li>2. Click <b>Add</b>.</li> <li>3. Add more next-hop addresses as necessary.</li> </ol> <p><b>NOTE:</b> 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"> <li>4. When you have finished adding next-hop addresses, click <b>OK</b>.</li> </ol> |

## 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 242
- Controlling Static Route Selection (Optional) on page 243
- Controlling Static Routes in the Routing and Forwarding Tables (Optional) on page 245
- Defining Default Behavior for All Static Routes (Optional) on page 246

For information about using the J-Web and CLI configuration editors, see “Using J-Web 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 52 shows a sample network.

**Figure 52: Customer Routes Connected to a Stub Network**



To configure customer routes as static routes, like the ones in Figure 52, 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 57.


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 243.
  - 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 245.
  - To define default properties for static routes, see “Defining Default Behavior for All Static Routes (Optional)” on page 246.
  - To check the configuration, see “Verifying the Static Route Configuration” on page 247.

**Table 57: Configuring Basic Static Routes**

| Task                                                                                               | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                 | CLI Configuration Editor                                                                                        |
|----------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Static</b> level in the configuration hierarchy.                                | In the configuration editor hierarchy, select <b>Routing options &gt; Static</b> .                                                                                                                                                                                                                                                                                         | From the top of the configuration hierarchy, enter<br><br>edit routing-options static                           |
| Add the static route <b>192.168.47.5/32</b> , and define the next-hop address <b>10.10.10.10</b> . | <ol style="list-style-type: none"> <li>1. Next to Route, click <b>Add new entry</b>.</li> <li>2. In the Destination box, type <b>192.168.47.5/32</b>.</li> <li>3. From the Next hop list, select <b>Next hop</b>.</li> <li>4. Next to Next hop, click <b>Add new entry</b>.</li> <li>5. In the Value box, type <b>10.10.10.10</b>.</li> <li>6. Click <b>OK</b>.</li> </ol> | Define the static route and set the next-hop address:<br><br><b>set route 192.168.47.5 next-hop 10.10.10.10</b> |

### Controlling Static Route Selection (Optional)

When multiple next hops exist for a single static route (see Figure 53), you can specify how traffic is to be routed to the destination.

**Figure 53: Controlling Static Routes in the Routing and Forwarding Tables**

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 58.
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 245.
  - To define default properties for static routes, see “Defining Default Behavior for All Static Routes (Optional)” on page 246.
  - To check the configuration, see “Verifying the Static Route Configuration” on page 247.

**Table 58: Controlling Static Route Selection**

| <b>Task</b>                                                                                        | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                           | <b>CLI Configuration Editor</b>                                                                                 |
|----------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Static</b> level in the configuration hierarchy.                                | In the configuration editor hierarchy, select <b>Routing options &gt; Static</b> .                                                                                                                                                                                                                                                                          | From the top of the configuration hierarchy, enter<br><br>edit routing-options static                           |
| Add the static route <b>192.168.47.5/32</b> , and define the next-hop address <b>10.10.10.10</b> . | <ol style="list-style-type: none"> <li>Next to Route, click <b>Add new entry</b>.</li> <li>In the Destination box, type <b>192.168.47.5/32</b>.</li> <li>From the Next hop list, select <b>Next hop</b>.</li> <li>In the Next hop box, click <b>Add new entry</b>.</li> <li>In the Value box, type <b>10.10.10.10</b>.</li> <li>Click <b>OK</b>.</li> </ol> | Define the static route and set the next-hop address:<br><br><b>set route 192.168.47.5 next-hop 10.10.10.10</b> |
| Set the preference for the <b>10.10.10.10</b> next hop to <b>7</b> .                               | <ol style="list-style-type: none"> <li>Next to Preference, select the <b>Yes</b> check box.</li> <li>Click <b>Configure</b>.</li> <li>In the Metric value box, type <b>7</b>.</li> <li>Click <b>OK</b>.</li> </ol>                                                                                                                                          | Set the preference to 7:<br><br><b>set route 192.168.47.5 next-hop 10.10.10.10 preference 7</b>                 |
| Define the qualified next-hop address <b>10.10.10.7</b> .                                          | <ol style="list-style-type: none"> <li>Next to Qualified next hop, click <b>Add new entry</b>.</li> <li>In the Nexthop box, type <b>10.10.10.7</b>.</li> </ol>                                                                                                                                                                                              | Set the qualified-next-hop address:<br><br><b>set route 192.168.47.5 qualified-next-hop 10.10.10.7</b>          |
| Set the preference for the <b>10.10.10.7</b> qualified next hop to <b>6</b> .                      | <ol style="list-style-type: none"> <li>In the Preference box, type <b>6</b>.</li> <li>Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                           | Set the preference to 6:<br><br><b>set route 192.168.47.5 qualified-next-hop 10.10.10.7 preference 6</b>        |

### **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:

- Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- Perform the configuration tasks described in Table 59.
- If you are finished configuring the router, commit the configuration.
- 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 246.
- To check the configuration, see “Verifying the Static Route Configuration” on page 247.

**Table 59: Controlling Static Routes in the Routing and Forwarding Tables**

| Task                                                                                                                                                                       | J-Web Configuration Editor                                                                                                                     | CLI Configuration Editor                                                                                    |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>192.168.47.5/32</b> level in the configuration hierarchy.                                                                                               | In the configuration editor hierarchy, select <b>Routing options &gt; Static</b> , then click <b>192.168.47.5/32</b> in the Destination field. | From the top of the configuration hierarchy, enter<br><br>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 <b>Yes</b> check box.                                                                                               | Set the <b>retain</b> attribute:<br><br>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 <b>No</b> check box.                                                                                           | Set the <b>no-readvertise</b> attribute:<br><br>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"> <li>From the Passive flag list, select <b>Passive</b>.</li> <li>Click <b>OK</b>.</li> </ol>                 | Set the <b>passive</b> attribute:<br><br>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:

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 60.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying the Static Route Configuration” on page 247.



**Table 60: Defining Static Route Defaults**

| Task                                                                                                                                                                       | J-Web Configuration Editor                                                                                                     | CLI Configuration Editor                                                                       |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------|
| Navigate to the <b>Defaults</b> level in the configuration hierarchy.                                                                                                      | In the configuration editor hierarchy, select <b>Protocols &gt; Static</b> , and then click <b>Configure</b> next to Defaults. | From the top of the configuration hierarchy, enter<br><br>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.                         | 1. Next to Retain, select the <b>Yes</b> check box.<br><br>2. Click <b>OK</b> .                                                | Set the <b>retain</b> attribute:<br><br>set retain                                             |
| Specify that the static route is not to be readadvertised. By default, static routes are eligible to be readadvertised.                                                    | 1. Next to Readvertise, select the <b>No</b> check box.<br><br>2. Click <b>OK</b> .                                            | Set the <b>no-readvertise</b> attribute:<br><br>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. | 1. From the Passive flag list, select <b>Passive</b> .<br><br>2. Click <b>OK</b> .                                             | Set the <b>passive</b> attribute:<br><br>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.

**Sample Output**

```

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 >fe-0/0/1.0
* 192.168.220.5/32 L 0 Local
* 192.168.220.8/30 D 0 >fe-0/0/2.0
* 192.168.220.9/32 L 0 Local
* 192.168.220.12/30 D 0 >fe-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.

## Chapter 8

# 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 249
- Before You Begin on page 250
- Configuring a RIP Network with Quick Configuration on page 250
- Configuring a RIP Network with a Configuration Editor on page 253
- Verifying the RIP Configuration on page 261

## 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 *J-series Services Router Getting Started Guide*.
- Configure network interfaces. See “Configuring Network Interfaces” on page 103.

## Configuring a RIP Network with Quick Configuration

---

J-Web Quick Configuration allows you to create RIP networks. Figure 54 shows the Quick Configuration Routing page for RIP.

**Figure 54: Quick Configuration Routing Page for RIP**

Juniper NETWORKS **ROUTER - J4300** Logged in as: **regress**  
[Help](#) [About](#) [Logout](#)

**Monitor Configuration Diagnose Manage**

**Quick Configuration**

Set Up  
SSL  
Interfaces  
Users  
SNMP  
**Routing**  
Firewall/NAT  
IPSec Tunnels  
Realtime Performance Monitoring  
View and Edit  
History  
Rescue

**Quick Configuration**

**Routing**

**RIP**

**Enable RIP** ☐ ?

**Advertise Default Route** ☐ ?

**RIP Interfaces**

**RIP-Enabled Interfaces**

**Logical Int**

fe-0/0/0.0  
fxp0.0  
lo0.0

OK Cancel Apply

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. [Trademark Notice](#)

To configure a RIP network with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Routing > RIP Routing**.
2. Enter information into the Quick Configuration page for RIP, as described in Table 61.
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 261.

**Table 61: RIP Routing Quick Configuration Summary**

| Field                   | Function                                                                   | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------|----------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>RIP</b>              |                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| Enable RIP              | Enables or disables RIP.                                                   | <ul style="list-style-type: none"> <li>■ To enable RIP, select the check box.</li> <li>■ To disable RIP, clear the check box.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Advertise Default Route | Advertises the default route using RIPv2.                                  | <ul style="list-style-type: none"> <li>■ To advertise the default route using RIPv2, select the check box.</li> <li>■ To disable the default route advertisement, clear the check box.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| RIP-Enabled Interfaces  | Designates one or more Services Router interfaces on which RIP is enabled. | <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"> <li>■ 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.</li> <li>■ 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.</li> <li>■ To enable RIP on all logical interfaces except the special <b>fxp0</b> management interface, select <b>All Interfaces</b> in the Logical Interfaces list and click the left arrow.</li> <li>■ To enable RIP on all the interfaces displayed in the Logical Interfaces list, click <b>All</b> to highlight every interface. Then click the left arrow to add the interfaces to the RIP interfaces list.</li> <li>■ 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.</li> </ul> |

## 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 253
- Controlling Traffic in a RIP Network (Optional) on page 256
- Enabling Authentication for RIP Exchanges (Optional) on page 259

For information about using the J-Web and CLI configuration editors, see “Using J-Web 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 55.

**Figure 55: 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 55, 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 62.
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 256.
  - To authenticate RIP exchanges, see “Enabling Authentication for RIP Exchanges (Optional)” on page 259.
  - To check the configuration, see “Verifying the RIP Configuration” on page 261.

**Table 62: Configuring a RIP Network**

| Task                                                             | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                   | CLI Configuration Editor                                                                                                                                                            |
|------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Rip</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Protocols &gt; Rip</b> .                                                                                                                                                                                                                                                                                                                    | From the top of the configuration hierarchy, enter<br><br>edit protocols rip                                                                                                        |
| Create the RIP group <b>alpha1</b> .                             | <ol style="list-style-type: none"> <li>1. Next to Group, click <b>Add new entry</b>.</li> <li>2. In the Group name box, type <b>alpha1</b>.</li> </ol>                                                                                                                                                                                                                                       | <ol style="list-style-type: none"> <li>1. Create the RIP group <b>alpha1</b>, and add an interface:<br/><br/><b>set group alpha1 neighbor fe-0/0/0.0</b></li> </ol>                 |
| Add interfaces to the RIP group <b>alpha1</b> .                  | <ol style="list-style-type: none"> <li>1. Next to Neighbor, click <b>Add new entry</b>.</li> <li>2. In the Neighbor name box, type the name of an interface on the Services Router—for example, <b>fe-0/0/0.0</b>—and click <b>OK</b>.</li> <li>3. Repeat Step 2 for each interface on this Services Router that you are adding to the RIP group. Only one interface is required.</li> </ol> | <ol style="list-style-type: none"> <li>2. Repeat Step 1 for each interface on this Services Router that you are adding to the RIP group. Only one interface is required.</li> </ol> |



**Table 62: Configuring a RIP Network (continued)**

| <b>Task</b>                                                                 | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure a routing policy to advertise directly connected routes.          | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Policy Options</b>.</li> <li>2. Next to Policy statement, click <b>Add new entry</b>.</li> <li>3. In the Policy name box, type the name of the policy statement—for example, <b>advertise-rip-routes</b>.</li> <li>4. Next to Term, click <b>Add new entry</b>.</li> <li>5. In the Term name box, type the name of the policy statement—for example, <b>from-direct</b>.</li> <li>6. Next to From, click <b>Configure</b>.</li> <li>7. Next to Protocol, click <b>Add new entry</b>.</li> <li>8. From the Value list, select <b>Direct</b>.</li> <li>9. Click <b>OK</b> until you return to the Policy statement page.</li> <li>10. Next to Then, click <b>Configure</b>.</li> <li>11. From the Accept reject list, select <b>Accept</b>.</li> <li>12. Click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><code>edit policy-options</code></li> <li>2. Set the match condition to match on direct routes:<br/><code>set policy-statement advertise-rip-routes term from-direct from protocol direct</code></li> <li>3. Set the match action to accept these routes:<br/><code>set policy-statement advertise-rip-routes term from-direct then accept</code></li> </ol> |
| Configure the previous routing policy to advertise routes learned from RIP. | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Policy Options</b>.</li> <li>2. Next to Policy statement, click <b>advertise-rip-routes</b>.</li> <li>3. Next to Term, click <b>Add new entry</b>.</li> <li>4. In the Term name box, type the name of the policy statement—for example, <b>from-rip</b>.</li> <li>5. Next to From, click <b>Configure</b>.</li> <li>6. Next to Protocol, click <b>Add new entry</b>.</li> <li>7. From the Value list, select <b>rip</b>.</li> <li>8. Click <b>OK</b> until you return to the Policy statement page.</li> <li>9. Next to Then, click <b>Configure</b>.</li> <li>10. From the Accept reject list, select <b>Accept</b>.</li> <li>11. Click <b>OK</b>.</li> </ol>                                                                                                                     | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><code>edit policy-options</code></li> <li>2. Set the match condition to match on direct routes:<br/><code>set policy-statement advertise-rip-routes term from-rip from protocol rip</code></li> <li>3. Set the match action to accept these routes:<br/><code>set policy-statement advertise-rip-routes term from-rip then accept</code></li> </ol>          |

## 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 256
- Controlling Traffic with the Outgoing Metric on page 257

## 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 56 shows a network with alternate routes between routers A and D.

### Figure 56: 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 63.
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 259.
  - To check the configuration, see “Verifying the RIP Configuration” on page 261.

Table 63: 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 <b>alpha1</b> RIP group. | <ol style="list-style-type: none"><li>1. In the configuration editor hierarchy, select <b>Protocols &gt; Rip</b>, and click <b>alpha1</b> in the Group name field.</li><li>2. Click the interface name—for example, <b>fe-0/0/0.0</b>—in the Neighbor name field.</li></ol> | From the top of the configuration hierarchy, enter<br><br>edit protocols rip group alpha1 neighbor fe-0/0/0 |
| Increase the incoming metric to <b>3</b> .                                                            | In the Metric in box, type <b>3</b> , and click <b>OK</b> .                                                                                                                                                                                                                 | Set the incoming metric to <b>3</b> :<br><br>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 57 shows a network with alternate routes between routers A and D.

**Figure 57: 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 64.
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 259.
  - To check the configuration, see “Verifying the RIP Configuration” on page 261.

**Table 64: Modifying the Outgoing Metric**

| <b>Task</b>                                                         | <b>J-Web Configuration Editor</b>                                                                                               | <b>CLI Configuration Editor</b>                                                           |
|---------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| Navigate to the <b>alpha1</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Protocols &gt; Rip</b> , and then click <b>alpha1</b> in the Group name field. | From the top of the configuration hierarchy, enter<br><br>edit protocols rip group alpha1 |
| Increase the outgoing metric to 3.                                  | In the Metric out box, type <b>3</b> , and click <b>OK</b> .                                                                    | Set the outgoing metric to 3:<br><br>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 259
- Enabling Authentication with MD5 Authentication on page 260

#### **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 65.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying the RIP Configuration” on page 261.

**Table 65: Configuring Simple RIP Authentication**

| <b>Task</b>                                                  | <b>J-Web Configuration Editor</b>                                         | <b>CLI Configuration Editor</b>                                              |
|--------------------------------------------------------------|---------------------------------------------------------------------------|------------------------------------------------------------------------------|
| Navigate to <b>Rip</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Protocols &gt; Rip</b> . | From the top of the configuration hierarchy, enter<br><br>edit protocols rip |

**Table 65: Configuring Simple RIP Authentication (continued)**

| <b>Task</b>                                                                                                                                                      | <b>J-Web Configuration Editor</b>                                                 | <b>CLI Configuration Editor</b>                                                                     |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|
| Set the authentication type to <b>simple</b> .                                                                                                                   | From the Authentication type list, select <b>simple</b> .                         | Set the authentication type to <b>simple</b> :<br><br>set authentication-type simple                |
| Set the authentication key to a simple-text password.<br><br>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 <b>OK</b> . | Set the authentication key to a simple-text password:<br><br>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 66.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying the RIP Configuration” on page 261.

**Table 66: Configuring MD5 RIP Authentication**

| <b>Task</b>                                                                                                                                      | <b>J-Web Configuration Editor</b>                                                    | <b>CLI Configuration Editor</b>                                                |
|--------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|
| Navigate to <b>Rip</b> level in the configuration hierarchy.                                                                                     | In the configuration editor hierarchy, select <b>Protocols &gt; Rip</b> .            | From the top of the configuration hierarchy, enter<br><br>edit protocols rip   |
| Set the authentication type to <b>MD5</b> .                                                                                                      | From the Authentication type list, select <b>md5</b> .                               | Set the authentication type to <b>md5</b> :<br><br>set authentication-type md5 |
| Set the MD5 authentication key (password).<br><br>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 <b>OK</b> . | Set the MD5 authentication key:<br><br>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 261
- Verifying the Exchange of RIP Messages on page 261
- Verifying Reachability of All Hosts in the RIP Network on page 262

### 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.

**Sample Output** `user@host> show rip neighbor`

| Source Neighbor | Destination State | Send Address  | Receive Address | In | Mode  | Mode | Met |
|-----------------|-------------------|---------------|-----------------|----|-------|------|-----|
| fe-0/0/0.0      | Dn                | (null)        | (null)          |    | mcast | both | 1   |
| fe-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.

### 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.

**Sample Output** `user@host> show rip statistics`

```

RIPv2 info: port 520; update interval 30s; holddown 180s; timeout 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
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 |

| fe-0/0/1.0: 10 routes learned; 3 routes advertised |       |            |             |
|----------------------------------------------------|-------|------------|-------------|
| 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 might indicate an authentication error.

## 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.

**Sample Output**

```
1 172.17.40.254 (172.17.40.254) 0.362 ms 0.284 ms 0.251 ms
2 router-a-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.

For information about the `traceroute` command and its output, see the *JUNOS System Basics and Services Command Reference*.



## Chapter 9

# 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 265
- Before You Begin on page 267
- Configuring an OSPF Network with Quick Configuration on page 267
- Configuring an OSPF Network with a Configuration Editor on page 269
- Tuning an OSPF Network for Efficient Operation on page 277
- Verifying an OSPF Configuration on page 281

## 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 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 169. For information about configuring OSPF demand circuits, see “Configuring Dial-on-Demand Routing with OSPF Support (Optional)” on page 190.

## Before You Begin

Before you begin configuring an OSPF network, complete the following tasks:

- Establish basic connectivity. See the *J-series Services Router Getting Started Guide*.
- Configure network interfaces. See “Configuring Network Interfaces” on page 103.

## Configuring an OSPF Network with Quick Configuration

J-Web Quick Configuration allows you to create single-area OSPF networks. Figure 58 shows the Quick Configuration Routing page for OSPF.

**Figure 58: Quick Configuration Routing Page for OSPF**

Juniper NETWORKS **ROUTER - J4300** Logged in as: **regress**  
[Help](#) [About](#) [Logout](#)

**Monitor Configuration Diagnose Manage**

**Quick Configuration**

Set Up  
SSL  
Interfaces  
Users  
SNMP

**Routing**

Firewall/NAT  
IPSec Tunnels  
Realtime Performance Monitoring

► **View and Edit**  
 ► **History**  
 ► **Rescue**

**Quick Configuration**

**Routing**

**Router Identification**

**Router Identifier**  ?

**OSPF**

**Enable OSPF** ☒

**OSPF Area ID**

**Area Type**  ?

**Enable OSPF on All Interfaces** ☒

**OSPF-Enabled Interfaces**

fe-0/0/0.0  
lo0.0

**OSP**

fxp0.

OK Cancel Apply

To configure a single-area OSPF network with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Routing > OSPF Routing**.
2. Enter information into the Quick Configuration Routing page for OSPF, as described in Table 67.
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 281.

**Table 67: OSPF Routing Quick Configuration Summary**

| Field                        | Function                                    | Your Action                                                                                                                                                                                                          |
|------------------------------|---------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Router Identification</b> |                                             |                                                                                                                                                                                                                      |
| Router Identifier (required) | Uniquely identifies the router.             | Type the Services Router’s 32-bit IP address, in dotted decimal notation.                                                                                                                                            |
| <b>OSPF</b>                  |                                             |                                                                                                                                                                                                                      |
| Enable OSPF                  | Enables or disables OSPF.                   | <ul style="list-style-type: none"> <li>■ To enable OSPF, select the check box.</li> <li>■ To disable OSPF, clear the check box.</li> </ul>                                                                           |
| OSPF Area ID                 | Uniquely identifies the area within its AS. | Type a 32-bit numeric identifier for the area, or an integer.<br><br>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. |

**Table 67: OSPF Routing Quick Configuration Summary (continued)**

| Field                   | Function                                                                    | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------|-----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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"> <li>■ <b>regular</b>—A regular OSPF area, including the backbone area</li> <li>■ <b>stub</b>—A stub area</li> <li>■ <b>nssa</b>—A not-so-stubby area (NSSA)</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| OSPF-Enabled Interfaces | Designates one or more Services Router interfaces on which OSPF is enabled. | <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"> <li>■ 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.</li> <li>■ 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.</li> <li>■ To enable OSPF on all logical interfaces except the special <b>fxp0</b> management interface, select <b>All Interfaces</b> in the Logical Interfaces list and click the left arrow.</li> <li>■ To enable OSPF on all the interfaces displayed in the Logical Interfaces list, click <b>All</b> to highlight every interface. Then click the left arrow to add the interfaces to the OSPF interfaces list.</li> <li>■ 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.</li> </ul> |

## 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 270
- Configuring a Single-Area OSPF Network (Required) on page 270
- Configuring a Multiarea OSPF Network (Optional) on page 272
- Configuring Stub and Not-So-Stubby Areas (Optional) on page 275

To configure an OSPF demand circuit, see “Configuring Dial-on-Demand Routing with OSPF Support (Optional)” on page 190. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 169.)

For information about using the J-Web and CLI configuration editors, see “Using J-Web 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 68.
3. Go on to “Configuring a Single-Area OSPF Network (Required)” on page 270.


**Table 68: Configuring the Router Identifier**

| Task                                                                                        | J-Web Configuration Editor                                                                                                 | CLI Configuration Editor                                                   |
|---------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------|
| Navigate to the <b>Routing-options</b> level in the configuration hierarchy.                | In the configuration editor hierarchy, select <b>Routing-options</b> .                                                     | From the top of the configuration hierarchy, enter<br>edit routing-options |
| Set the router ID value to the IP address of the Services Router—for example, 177.162.4.24. | <ol style="list-style-type: none"> <li>1. In the Router Id box, type 177.162.4.24.</li> <li>2. Click <b>OK</b>.</li> </ol> | Enter<br>set router-id 177.162.4.24                                        |

## 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 59.



**Figure 59: Typical Single-Area OSPF Network Topology**

To configure a single-area OSPF network with a backbone area, like the one in Figure 59, 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 69.
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 272.
  - To control external route advertisement in the AS, see “Configuring Stub and Not-So-Stubby Areas (Optional)” on page 275.
  - To configure an OSPF demand circuit, see “Configuring Dial-on-Demand Routing with OSPF Support (Optional)” on page 190. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 169.)
  - To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 277.
  - To check the configuration, see “Verifying an OSPF Configuration” on page 281.

**Table 69: Configuring a Single-Area OSPF Network**

| Task                                                              | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                              | CLI Configuration Editor                                                                                                                                                                |
|-------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Ospf</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Protocols &gt; Ospf</b> .                                                                                                                                                                                                                                                                              | From the top of the configuration hierarchy, enter<br><br>edit protocols ospf                                                                                                           |
| Create the backbone area with area ID 0.0.0.0.                    | <ol style="list-style-type: none"> <li>1. In the Area box, click <b>Add new entry</b>.</li> <li>2. In the Area ID box, type 0.0.0.0.</li> </ol>                                                                                                                                                                                                         | <ol style="list-style-type: none"> <li>1. Set the backbone area ID to 0.0.0.0 and add an interface:<br/><br/>set area 0.0.0.0 interface fe-0/0/0</li> </ol>                             |
| Add interfaces as needed to the OSPF area—for example, fe-0/0/0.  | <ol style="list-style-type: none"> <li>1. In the Interface box, click <b>Add new entry</b>.</li> <li>2. In the Interface name box, type fe-0/0/0.</li> <li>3. Click <b>OK</b>.</li> <li>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.</li> </ol> | <ol style="list-style-type: none"> <li>2. Repeat Step 1 for each interface on this Services Router that you are adding to the backbone area. Only one interface is required.</li> </ol> |

### 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 60.

**Figure 60: Typical Multiarea OSPF Network Topology**

To configure a multiarea OSPF network shown in Figure 60, 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 273
- Creating Additional OSPF Areas on page 273
- Configuring Area Border Routers on page 274

## 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 270.

## 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 70.
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 274.
  - To control external route advertisement in the AS, see “Configuring Stub and Not-So-Stubby Areas (Optional)” on page 275.
  - To configure an OSPF demand circuit, see “Configuring Dial-on-Demand Routing with OSPF Support (Optional)” on page 190. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 169.)
  - To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 277.
  - To check the configuration, see “Verifying an OSPF Configuration” on page 281.

**Table 70: Configuring a Multiarea OSPF Network**

| <b>Task</b>                                                                                        | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                       | <b>CLI Configuration Editor</b>                                                                                                                                                |
|----------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Ospf</b> level in the configuration hierarchy.                                  | In the configuration editor hierarchy, select <b>Protocols &gt; Ospf</b> .                                                                                                                                                                                                                                                                              | From the top of the configuration hierarchy, enter<br><br>edit protocols ospf                                                                                                  |
| 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"> <li>1. In the Area box, click <b>Add new entry</b>.</li> <li>2. In the Area ID box, type 0.0.0.2.</li> </ol>                                                                                                                                                                                                         | <ol style="list-style-type: none"> <li>1. Set the area ID to 0.0.0.2 and add an interface:<br/><br/>set area 0.0.0.2 interface fe-0/0/0</li> </ol>                             |
| Add interfaces as needed to the OSPF area—for example, fe-0/0/0.                                   | <ol style="list-style-type: none"> <li>1. In the Interface box, click <b>Add new entry</b>.</li> <li>2. In the Interface name box, type fe-0/0/0.</li> <li>3. Click <b>OK</b>.</li> <li>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.</li> </ol> | <ol style="list-style-type: none"> <li>2. Repeat Step 1 for each interface on this Services Router that you are adding to the area. Only one interface is required.</li> </ol> |

## 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 60 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 71.
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 275.
  - To configure an OSPF demand circuit, see “Configuring Dial-on-Demand Routing with OSPF Support (Optional)” on page 190. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 169.)

- To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 277.
- To check the configuration, see “Verifying an OSPF Configuration” on page 281.

**Table 71: Configuring Area Border Routers**


| Task                                                                       | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | CLI Configuration Editor                                                                                                                                                                                                                                                                                                                                                                            |
|----------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Ospf</b> level in the configuration hierarchy.          | In the configuration editor hierarchy, select <b>Protocols &gt; Ospf</b> .                                                                                                                                                                                                                                                                                                                                                                                                                        | From the top of the configuration hierarchy, enter<br><br>edit protocols ospf                                                                                                                                                                                                                                                                                                                       |
| Verify that the backbone area has at least one interface enabled for OSPF. | Click <b>0.0.0.0</b> to display the Area ID <b>0.0.0.0</b> page, and verify that the backbone area has at least one interface enabled for OSPF.<br><br>For example, Services Router B in Figure 60 has the following interfaces enabled for OSPF in the backbone area: <ul style="list-style-type: none"> <li>■ Interface fe-0/0/0.0</li> <li>■ Interface fe-0/0/1.0</li> </ul> To enable an interface on the backbone area, see “Configuring a Single-Area OSPF Network (Required)” on page 270. | View the configuration using the <b>show</b> command:<br><br><b>show</b><br><br>For example, Services Router B in Figure 60 has the following interfaces enabled for OSPF in the backbone area:<br><br><b>area 0.0.0.0 { interface fe-0/0/0.0; interface fe-0/0/1.0; }</b><br><br>To enable an interface on the backbone area, see “Configuring a Single-Area OSPF Network (Required)” on page 270. |
| Create the additional area with a unique area ID—for example, 0.0.0.2.     | 1. In the Area box, click <b>Add new entry</b> .<br><br>2. In the Area ID box, type 0.0.0.2.                                                                                                                                                                                                                                                                                                                                                                                                      | 1. Set the area ID to 0.0.0.2 and add an interface:<br><br><b>set area 0.0.0.2 interface fe-0/0/0</b>                                                                                                                                                                                                                                                                                               |
| Add interfaces as needed to the OSPF area—for example, fe-0/0/0.           | 1. In the Interface box, click <b>Add new entry</b> .<br><br>2. In the Interface name box, type fe-0/0/0.<br><br>3. Click <b>OK</b> .<br><br>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.                                                                                                                                                                                                 | 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 61, 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 61: OSPF Network Topology with Stub Areas and NSSAs**



To configure stub areas and NSSAs in an OSPF network like the one shown in Figure 61:

1. Create the area and enable OSPF on the interfaces within that area.  
For instructions, see “Creating Additional OSPF Areas” on page 273.
2. Configure an area border router to bridge the areas.  
For instructions, see “Configuring Area Border Routers” on page 274.
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 72.
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 with OSPF Support (Optional)” on page 190. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 169.)

- To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 277.
- To check the configuration, see “Verifying an OSPF Configuration” on page 281.

**Table 72: Configuring Stub Area and Not-So-Stubby Area Routers**

| Task                                                                     | J-Web Configuration Editor                                                                                                                                                                                                                                | CLI Configuration Editor                                                                                                                                                                                                                  |
|--------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>0.0.0.7</b> level in the configuration hierarchy.     | In the configuration editor hierarchy, select <b>Protocols &gt; Ospf &gt; Area id 0.0.0.7</b> .                                                                                                                                                           | From the top of the configuration hierarchy, enter<br><br>edit protocols ospf area 0.0.0.7                                                                                                                                                |
| Configure each Services Router in area <b>0.0.0.7</b> as a stub router.  | <ol style="list-style-type: none"> <li>1. In the Stub option list, select <b>Stub</b> and click <b>OK</b>.</li> <li>2. Repeat Step 1 for every Services Router in the stub area to configure them with the <b>stub</b> parameter for the area.</li> </ol> | <ol style="list-style-type: none"> <li>1. Set the stub attribute:<br/><br/><b>set stub</b></li> <li>2. Repeat Step 1 for every Services Router in the stub area to configure them with the <b>stub</b> parameter for the area.</li> </ol> |
| Navigate to the <b>0.0.0.9</b> level in the configuration hierarchy.     | In the configuration editor hierarchy, select <b>Protocols &gt; Ospf &gt; Area &gt; 0.0.0.9</b> .                                                                                                                                                         | From the top of the configuration hierarchy, enter<br><br>edit protocols ospf area 0.0.0.9                                                                                                                                                |
| Configure each Services Router in area <b>0.0.0.9</b> as an NSSA router. | <ol style="list-style-type: none"> <li>1. In the Stub option list, select <b>Nssa</b> and click <b>OK</b>.</li> <li>2. Repeat Step 1 for every Services Router in the NSSA to configure them with the <b>nssa</b> parameter for the area.</li> </ol>      | <ol style="list-style-type: none"> <li>1. Set the nssa attribute:<br/><br/><b>set nssa</b></li> <li>2. Repeat Step 1 for every Services Router in the NSSA to configure them with the <b>nssa</b> parameter for the area.</li> </ol>      |

## 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 278
- Controlling the Cost of Individual Network Segments on page 278
- Enabling Authentication for OSPF Exchanges on page 279
- Controlling Designated Router Election on page 280

## 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 (SPF) 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 73.

**Table 73: Controlling Route Selection in the Forwarding Table by Setting Preferences**

| Task                                                              | J-Web Configuration Editor                                                                                                                                                                              | CLI Configuration Editor                                                                                                                                                        |
|-------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Ospf</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Protocols &gt; Ospf</b> .                                                                                                                              | From the top of the configuration hierarchy, enter<br><br>edit protocols ospf                                                                                                   |
| Set the external and internal route preferences.                  | <ol style="list-style-type: none"> <li>1. In the External preference box, type 130.</li> <li>2. In the Preference box, type the internal preference value of 7.</li> <li>3. Click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>1. Set the external preference:<br/>set external-preference 130</li> <li>2. Set the internal preference:<br/>set preference 7</li> </ol> |

## 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 74.

**Table 74: Controlling the Cost of Individual Network Segments by Modifying the Metric**

| Task                                                                    | J-Web Configuration Editor                                                                                                           | CLI Configuration Editor                                                                                           |
|-------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>fe-0/0/0.0</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Protocols &gt; Ospf &gt; Area id 0.0.0.0 &gt; Interface name fe-0/0/0.0</b> .       | From the top of the configuration hierarchy, enter<br><br>edit protocols ospf area 0.0.0.0<br>interface fe-0/0/0.0 |
| Set the interface metric.                                               | <ul style="list-style-type: none"><li>1. In the Metric box, type the interface metric value 5.</li><li>2. Click <b>OK</b>.</li></ul> | Set the interface metric:<br><br>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 75.

**Table 75: Enabling OSPF Authentication**

| <b>Task</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                   | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                                                  |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>0.0.0.0</b> level in the configuration hierarchy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | In the configuration editor hierarchy, select <b>Protocols &gt; Ospf &gt; Area id 0.0.0.0</b> .                                                                                                                                                                                                                     | From the top of the configuration hierarchy, enter<br><br><code>edit protocols ospf area 0.0.0.0</code>                                                                                                                                                                                                          |
| Set the authentication type for the stub area to either simple or MD5—for example, MD5.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <ol style="list-style-type: none"> <li>From the Authentication type list, select <b>md5</b>.</li> <li>Click <b>OK</b>.</li> </ol>                                                                                                                                                                                   | Set the authentication type:<br><br><code>set authentication-type md5</code>                                                                                                                                                                                                                                     |
| Navigate to the <i>interface-name</i> level in the configuration hierarchy.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | In the configuration editor hierarchy under <b>Protocols &gt; Ospf &gt; Area &gt; 0.0.0.0 &gt; interface</b> , click an interface name.                                                                                                                                                                             | From the top of the configuration hierarchy, enter<br><br><code>edit protocols ospf area 0.0.0.0 interface interface-name</code>                                                                                                                                                                                 |
| Set the authentication password (key) and, for MD5 authentication only, the key identifier to associate with the MD5 password: <ul style="list-style-type: none"> <li>■ For simple authentication, set a password of from 1 through 8 ASCII characters—for example, <b>Chey3nne</b>.</li> <li>■ For MD5 authentication:               <ul style="list-style-type: none"> <li>■ Set a password of from 1 through 16 ASCII characters—for example, <b>Chey3nne</b>.</li> <li>■ Set a key identifier between 0 (the default) and 255—for example, 2.</li> </ul> </li> </ul> | <ol style="list-style-type: none"> <li>In the Key name box, type <b>Chey3nne</b>.</li> <li>For MD5 authentication only, in the Key ID box, type <b>2</b>.</li> <li>Click <b>OK</b>.</li> <li>Repeat Step 1 through Step 3 for each interface in the stub area for which you are enabling authentication.</li> </ol> | <ol style="list-style-type: none"> <li>Set the authentication password and, for MD5 authentication only, set the key identifier:<br/><br/><code>set authentication-key Chey3nne key-id 2</code></li> <li>Repeat Step 1 for each interface in the stub area for which you are enabling authentication.</li> </ol> |

## 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:

- Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
- Perform the configuration tasks described in Table 76.

**Table 76: 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>fe-0/0/1</code> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Protocols &gt; Ospf &gt; area id 0.0.0.3 &gt; Interface name fe-0/0/1</b> . | From the top of the configuration hierarchy, enter<br><br><code>edit protocols ospf area 0.0.0.3 interface fe-0/0/1</code> |
| Set the Services Router priority to a value between 0 and 255—for example, 200. The default value is 128.                                                | 1. In the Priority box, type 200.<br>2. Click <b>OK</b> .                                                                    | Set the priority value:<br><br><code>set priority 200</code>                                                               |

## Verifying an OSPF Configuration

To verify an OSPF configuration, perform these tasks:

- Verifying OSPF-Enabled Interfaces on page 281
- Verifying OSPF Neighbors on page 282
- Verifying the Number of OSPF Routes on page 283
- Verifying Reachability of All Hosts in an OSPF Network on page 284

### 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.

**Sample Output** `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**.

For more information about `show ospf interface`, 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.

**Sample Output** `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.

For more information about `show ospf neighbor`, 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 62 shows a sample network with an OSPF topology.

Figure 62: 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.

|                      |                                            |           |            |         |        |                   |                    |
|----------------------|--------------------------------------------|-----------|------------|---------|--------|-------------------|--------------------|
| <b>Sample Output</b> | <code>user@host&gt; show ospf route</code> |           |            |         |        |                   |                    |
|                      | Prefix                                     | Path Type | Route Type | NH Type | Metric | NextHop Interface | Nexthop addr/label |
|                      | 10.10.10.1/24                              | Intra     | Network    | IP      | 1      | fe-0/0/2.0        | 10.0.21.1          |
|                      | 10.10.10.2/24                              | Intra     | Network    | IP      | 1      | fe-0/0/2.0        | 10.0.21.1          |
|                      | 10.10.10.4/24                              | Intra     | Network    | IP      | 1      | fe-0/0/1.0        | 10.0.13.1          |

|                |       |         |    |   |            |           |
|----------------|-------|---------|----|---|------------|-----------|
| 10.10.10.5/24  | Intra | Network | IP | 1 | fe-0/0/2.0 | 10.0.21.1 |
| 10.10.10.6/24  | Intra | Network | IP | 1 | fe-0/0/1.0 | 10.0.13.1 |
| 10.10.10.10/24 | Intra | Network | IP | 1 | fe-0/0/2.0 | 10.0.21.1 |
| 10.10.10.11/24 | Intra | Network | IP | 1 | fe-0/0/1.0 | 10.0.13.1 |
| 10.10.10.13/24 | Intra | Network | IP | 1 | fe-0/0/1.0 | 10.0.13.1 |
| 10.10.10.16/24 | Intra | Network | IP | 1 | fe-0/0/1.0 | 10.0.13.1 |
| 10.10.10.19/24 | Intra | Network | IP | 1 | fe-0/0/1.0 | 10.0.13.1 |
| 10.10.10.20/24 | Intra | Network | IP | 1 | fe-0/0/2.0 | 10.0.21.1 |
| 10.10.10.21/24 | Intra | Network | IP | 1 | fe-0/0/2.0 | 10.0.21.1 |
| 192.168.5.1    | Intra | Router  | IP | 1 | fe-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 | fe-0/0/1.0 | 10.0.13.1 |
| 192.168.5.4    | Intra | Router  | IP | 1 | fe-0/0/1.0 | 10.0.13.1 |
| 192.168.5.5    | Intra | Router  | IP | 1 | fe-0/0/1.0 | 10.0.13.1 |
| 192.168.5.6    | Intra | Router  | IP | 1 | fe-0/0/2.0 | 10.0.21.1 |
| 192.168.5.7    | Intra | Router  | IP | 1 | fe-0/0/2.0 | 10.0.21.1 |
| 192.168.5.8    | Intra | Router  | IP | 1 | fe-0/0/2.0 | 10.0.21.1 |
| 192.168.5.9    | Intra | Router  | IP | 1 | fe-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 62, verify that the OSPF routing table has 21 entries, one for each network segment and one for each router's loopback address.

For more information about `show ospf route`, 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.

**Sample Output**

```
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 `ospf routeshow`, see “Verifying the Number of OSPF Routes” on page 283.

For information about the `traceroute` command and its output, see the *JUNOS System Basics and Services Command Reference*.





## Chapter 10

# 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 287
- Before You Begin on page 289
- Configuring IS-IS with a Configuration Editor on page 289
- Verifying IS-IS on a Services Router on page 290

### 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 287
- System Identifier Mapping on page 288

### ISO Network Addresses

IS-IS uses ISO network address. 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.

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 *J-series Services Router Getting Started Guide*.
- If you do not already have an understanding of IS-IS, read “IS-IS Overview” on page 222 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 77.
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 77: Configuring the IS-IS Protocol**

| Task                                                                    | J-Web Configuration Editor                                                                                                                                                        | CLI Configuration Editor                                               |
|-------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Interfaces</b> .                                                                                                                 | From the top of the configuration hierarchy, enter<br>edit interfaces. |
| Configure the loopback interface lo0.                                   | <ol style="list-style-type: none"> <li>1. Next to Interface, click <b>Add new entry</b>.</li> <li>2. In the Interface name box, type lo0.</li> <li>3. Click <b>OK</b>.</li> </ol> | Enter<br>edit interfaces lo0                                           |

**Table 77: Configuring the IS-IS Protocol (continued)**

| <b>Task</b>                                                                                             | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                    | <b>CLI Configuration Editor</b>                                                                                          |
|---------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------|
| Configure the logical unit on the loopback interface—for example 0.                                     | 1. Next to lo0, click <b>Edit</b> under Encapsulation.                                                                                                                                                                                                                                                                                                                                                                               | 1. Enter<br>edit unit 0                                                                                                  |
| Add the NET address to the loopback interface—for example, 49.0001.00a0.c96b.c490.00.                   | 2. Next to Unit, click <b>Add new entry</b> .<br>3. In the Interface unit number box, type 0.<br>4. Under Family, select <b>Iso</b> .<br>5. Next to Address, click <b>Add new entry</b> .<br>6. In the Source box, type 49.0001.00a0.c96b.c490.00.<br>7. Click <b>OK</b> until you return to the Interfaces page.                                                                                                                    | 2. Enter<br>set family iso address<br>49.0001.00a0.c96b.c490.00                                                          |
| Configure a physical interface—for example, fe-0/0/1—with the NET address, and add the Family type iso. | 1. Next to fe-0/0/1, click <b>Edit</b> under Encapsulation.<br>2. Next to Unit, click <b>Add new entry</b> .<br>3. In the Interface unit number box, type 0.<br>4. Under Family, select <b>Iso</b> .<br>5. Next to Iso, click <b>Configure</b> .<br>6. Next to Address, click <b>Add new entry</b> .<br>7. In the Source box, type 49.0001.00a0.c96b.c490.00.<br>8. Click <b>OK</b> until you return to the Edit Configuration page. | Enter<br>edit interfaces fe-0/0/1<br>Enter<br>set unit 0<br>Enter<br>set family iso address<br>49.0001.00a0.c96b.c490.00 |
| Navigate to the <b>Protocols</b> level in the configuration hierarchy.                                  | In the configuration editor hierarchy, select <b>Protocols</b> .                                                                                                                                                                                                                                                                                                                                                                     | From the top of the configuration hierarchy, enter<br>edit protocols                                                     |
| Add the IS-IS protocol to all interfaces on the Services Router.                                        | 1. Next to Protocols, click <b>Edit</b> .<br>2. Next to Isis, click <b>Edit</b> .<br>3. In the Interface name box, type all.<br>4. Click <b>OK</b> .                                                                                                                                                                                                                                                                                 | Enter<br>set isis interface all                                                                                          |

## Verifying IS-IS on a Services Router

To verify IS-IS, perform these tasks:

- Displaying IS-IS Interface Configuration on page 291
- Displaying IS-IS Interface Configuration Detail on page 291
- Displaying IS-IS Adjacencies on page 292
- Displaying IS-IS Adjacencies in Detail on page 293

## 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.

**Sample Output** `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
fe-0/0/1.0 2 0x9 Disabled router.03
fe-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.

For more information about the `show isis interface brief` command, 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.

**Sample Output** `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
fe-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

For more information about the `show isis interface detail` command, 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.

**Sample Output**

```
user@host> show isis adjacency brief

IS-IS adjacency database:
 Interface System L State Hold (secs) SNPA
```

```

fe-0/0/0.0 1921.6800.5067 2 Up 13
fe-0/0/1.0 1921.6800.5067 2 Up 25
fe-0/0/2.0 1921.6800.5067 2 Up 19

```

**What It Means** Verify adjacent routers in the IS-IS database.

For more information about the `show isis adjacency brief` command, 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.

**Sample Output** `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 Seenself

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 Seenself

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 Seenself

```

**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

For more information about the `show isis adjacency extensive` command, see the *JUNOS Routing Protocols and Policies Command Reference*.



## Chapter 11

# 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 295
- Before You Begin on page 297
- Configuring BGP Sessions with Quick Configuration on page 297
- Configuring BGP Sessions with a Configuration Editor on page 299
- Verifying a BGP Configuration on page 308

### 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.

### **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 231

### **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 234

## Before You Begin

Before you begin configuring a BGP network, complete the following tasks:

- Establish basic connectivity. See the *J-series Services Router Getting Started Guide*.
- Configure network interfaces. See “Configuring Network Interfaces” on page 103.

## Configuring BGP Sessions with Quick Configuration

J-Web Quick Configuration allows you to create BGP peering sessions. Figure 63 shows the Quick Configuration Routing page for BGP.

**Figure 63: Quick Configuration Routing Page for BGP**

Juniper NETWORKS **ROUTER - J4300** Logged in as: **regress**  
[Help](#) [About](#) [Logout](#)

**Monitor** **Configuration** **Diagnose** **Manage**

**Quick Configuration**

- Set Up
- SSL
- Interfaces
- Users
- SNMP
- Routing**
- Firewall/NAT
- IPSec Tunnels
- Realtime Performance Monitoring
- **View and Edit**
- **History**
- **Rescue**

**Configuration > Quick Configuration > Routing**

### Quick Configuration

### Routing

#### Router Identification

\* **Router Identifier**  ?

#### BGP

**Enable BGP** ☐

**Autonomous System Number**  ?

**Peer Autonomous System Number**  ?

**Peer Address**

**Local Address**  ?

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. [Trademark Notice](#).

To configure a BGP peering session with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Routing > BGP Routing**.
2. Enter information into the Quick Configuration page for BGP, as described in Table 78.
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 308.

**Table 78: BGP Routing Quick Configuration Summary**

| Field                         | Function                                                                                                | Your Action                                                                                                                                                                                                                                        |
|-------------------------------|---------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Router Identification</b>  |                                                                                                         |                                                                                                                                                                                                                                                    |
| Router Identifier (required)  | Uniquely identifies the router                                                                          | Type the Services Router's 32-bit IP address, in dotted decimal notation.                                                                                                                                                                          |
| <b>BGP</b>                    |                                                                                                         |                                                                                                                                                                                                                                                    |
| Enable BGP                    | Enables or disables BGP.                                                                                | <ul style="list-style-type: none"> <li>■ To enable BGP, select the check box.</li> <li>■ To disable BGP, clear the check box.</li> </ul>                                                                                                           |
| 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 <b>3</b>, the value assigned to the AS is <b>0.0.0.3</b>.</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 <b>3</b>, the value assigned to the AS is <b>0.0.0.3</b>.</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)*.

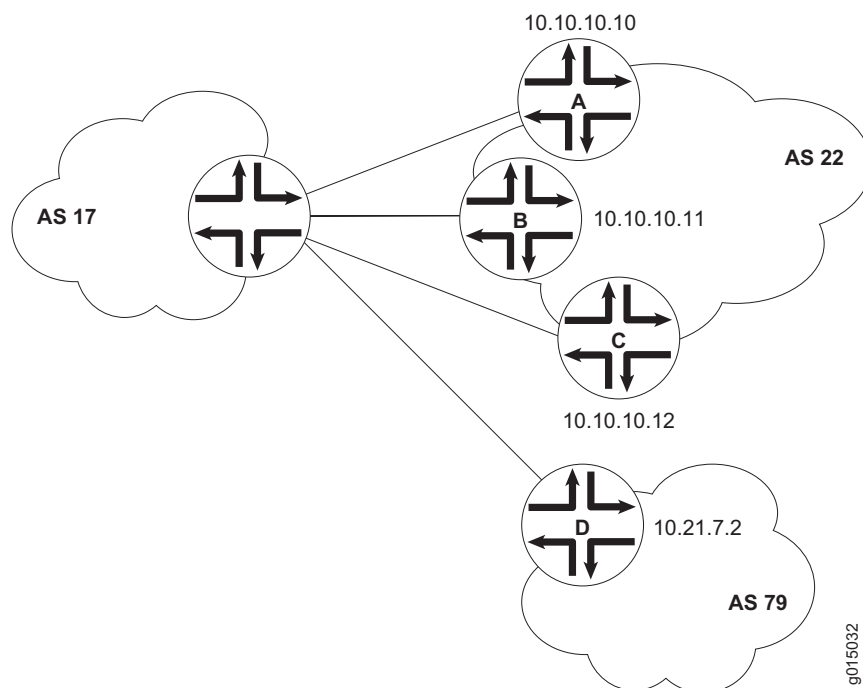
- Configuring a Point-to-Point Peering Session (Required) on page 299
- Configuring BGP Within a Network (Required) on page 302
- Configuring a Route Reflector (Optional) on page 303
- Configuring BGP Confederations (Optional) on page 306

For information about using the J-Web and CLI configuration editors, see “Using J-Web Configuration Tools” on page 3.

### Configuring a Point-to-Point Peering Session (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 64 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 64: Typical Network with BGP Peering Sessions**

To configure the BGP peering sessions shown in Figure 64:

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 79.
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 302.
  - To configure route reflector clusters, see “Configuring a Route Reflector (Optional)” on page 303.
  - To subdivide autonomous systems (ASs), see “Configuring BGP Confederations (Optional)” on page 306.
  - To check the configuration, see “Verifying a BGP Configuration” on page 308.

**Table 79: Configuring BGP Peering Sessions**


| <b>Task</b>                                                                                                                                                                                                                                                                       | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Routing-options</b> level in the configuration hierarchy.                                                                                                                                                                                                      | In the configuration editor hierarchy, select <b>Routing-options</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                               | From the top of the configuration hierarchy, enter<br><br>edit routing-options                                                                                                                                                                                                                            |
| Set the network's AS number to <b>17</b> .                                                                                                                                                                                                                                        | <ol style="list-style-type: none"> <li>1. In the AS Number box, enter <b>17</b>.</li> <li>2. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                             | Set the AS number to <b>17</b> :<br><br>set autonomous-system 17                                                                                                                                                                                                                                          |
| Navigate to the <b>Bgp</b> level in the configuration hierarchy.                                                                                                                                                                                                                  | In the configuration editor hierarchy, select <b>Protocols &gt; Bgp</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                            | From the top of the configuration hierarchy, enter<br><br>edit protocols bgp                                                                                                                                                                                                                              |
| Create the BGP group <b>external-peers</b> , and add the external neighbor addresses to the group.                                                                                                                                                                                | <ol style="list-style-type: none"> <li>1. In the Group box, click <b>Add new entry</b>.</li> <li>2. In the Group name box, type the name of the group of external BGP peers—<b>external-peers</b> in this case.</li> <li>3. In the Neighbor box, click <b>Add new entry</b>.</li> <li>4. In the Address box, type the IP address of an external BGP peer, in dotted decimal notation, and click <b>OK</b>.</li> <li>5. Repeat Step 3 and Step 4 for each BGP neighbor within the external group that you are configuring.</li> </ol> | <ol style="list-style-type: none"> <li>1. Create the group <b>external-peers</b>, and add the address of an external neighbor:<br/><br/>set group external-peers neighbor 10.10.10.10</li> <li>2. Repeat Step 1 for each BGP neighbor within the external peer group that you are configuring.</li> </ol> |
| At the group level, set the AS number for the group <b>external-peers</b> to <b>22</b> .<br><br>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"> <li>1. In the Peer as box, type the number of the AS in which most peers in the <b>external-peers</b> group reside.</li> <li>2. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                       | From the [edit protocols bgp] hierarchy level:<br><br>set group external-peers peer-as 22                                                                                                                                                                                                                 |
| At the individual neighbor level, set the AS number for peer D to <b>79</b> .<br><br>Because peer D is a member of the group <b>external-peers</b> , 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"> <li>1. Under Neighbor, in the Address column, click the IP address of peer D—<b>10.21.7.2</b> in this case.</li> <li>2. In the Peer as box, type the AS number of the peer.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                               | From the [edit protocols bgp group external-peers] hierarchy level:<br><br>set neighbor 10.21.7.2 peer-as 79                                                                                                                                                                                              |
| Set the group type to <b>external</b> .                                                                                                                                                                                                                                           | <ol style="list-style-type: none"> <li>1. From the Type list, select <b>external</b>.</li> <li>2. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                        | From the [edit protocols bgp group external-peers] hierarchy level:<br><br>set type external                                                                                                                                                                                                              |

## 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 65 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 65: Typical Network with EBGP External Sessions and IBGP Internal Sessions**



To configure IBGP in the network shown in Figure 65:

1. Configure all external peering sessions as described in “Configuring a Point-to-Point Peering Session (Required)” on page 299.
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 80.
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 303.
  - To subdivide autonomous systems (ASs), see “Configuring BGP Confederations (Optional)” on page 306.



- To check the configuration, see “Verifying a BGP Configuration” on page 308.

**Table 80: Configuring IBGP Peering Sessions**

| Task                                                                                                                                                                                                                                         | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | CLI Configuration Editor                                                                                                                                                                                                                                                             |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Bgp</b> level in the configuration hierarchy.                                                                                                                                                                             | In the configuration editor hierarchy, select <b>Protocols &gt; Bgp</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                         | From the top of the configuration hierarchy, enter<br><br>edit protocols bgp                                                                                                                                                                                                         |
| Create the BGP group <b>internal-peers</b> , and add the internal neighbor addresses to the group.<br><br>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"> <li>1. In the Group box, click <b>Add new entry</b>.</li> <li>2. In the Group name box, type the name of the group of internal BGP peers—<b>internal-peers</b> in this case.</li> <li>3. In the Neighbor box, click <b>Add new entry</b>.</li> <li>4. In the Address box, type the IP address of an internal BGP peer, in dotted decimal notation.</li> <li>5. Click <b>OK</b>.</li> <li>6. Repeat Step 3 and Step 4 for each internal BGP peer within the network.</li> </ol> | <ol style="list-style-type: none"> <li>1. Create the group <b>internal-peers</b>, and add the address of an internal neighbor:<br/><br/><b>set group internal-peers neighbor 192.168.6.4</b></li> <li>2. Repeat Step 1 for each internal BGP neighbor within the network.</li> </ol> |
| Set the group type to <b>internal</b> .                                                                                                                                                                                                      | <ol style="list-style-type: none"> <li>1. From the Type list, select <b>internal</b>.</li> <li>2. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                     | From the [edit protocols bgp group internal-peers] hierarchy level:<br><br><b>set type internal</b>                                                                                                                                                                                  |
| Configure a routing policy to advertise BGP routes.                                                                                                                                                                                          | See “Injecting OSPF Routes into the BGP Routing Table (Optional)” on page 464.                                                                                                                                                                                                                                                                                                                                                                                                                                    |                                                                                                                                                                                                                                                                                      |

### 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 66 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 66: 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 a Point-to-Point Peering Session (Required)” on page 299.
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 81.
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 306.

- To check the configuration, see “Verifying a BGP Configuration” on page 308.

**Table 81: 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 <b>Bgp</b> level in the configuration hierarchy.                                                                          | In the configuration editor hierarchy, select <b>Protocols &gt; Bgp</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                | From the top of the configuration hierarchy, enter<br><br>edit protocols bgp                                                                                                                                                                                                                       |
| On the Services Router that you are using as a route reflector, create the BGP group <b>cluster-peers</b> , and add to the group the IP addresses of the internal neighbors that you want in the cluster. | <ol style="list-style-type: none"> <li>1. In the Group box, click <b>Add new entry</b>.</li> <li>2. In the Group name box, type the name of the group in which the BGP peer is configured—<b>cluster-peers</b> in this case.</li> <li>3. In the Neighbor box, click <b>Add new entry</b>.</li> <li>4. In the Address box, type the IP address of a BGP peer, in dotted decimal notation.</li> <li>5. Click <b>OK</b>.</li> <li>6. Repeat Step 3 and Step 4 for each BGP neighbor within the cluster that you are configuring.</li> </ol> | <ol style="list-style-type: none"> <li>1. Create the group <b>cluster-peers</b>, and add the address of an internal neighbor:<br/><br/><b>set group cluster-peers neighbor 192.168.6.4</b></li> <li>2. Repeat Step 1 for each BGP neighbor within the cluster that you are configuring.</li> </ol> |
| On the Services Router that you are using as a route reflector, set the group type to <b>internal</b> .                                                                                                   | From the Type list, select <b>internal</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | From the [edit protocols bgp group internal-peers] hierarchy level:<br><br><b>set type internal</b>                                                                                                                                                                                                |
| 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"> <li>1. In the Cluster box, enter the unique numeric cluster identifier.</li> <li>2. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                       | Set the cluster identifier:<br><br><b>set cluster 2.3.4.5</b>                                                                                                                                                                                                                                      |

**Table 81: Configuring a Route Reflector (continued)**


| <b>Task</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                                                                                             |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>On the other routers in the cluster, create the BGP group <b>cluster-peers</b>, 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><b>NOTE:</b> 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"> <li>1. In the configuration editor hierarchy, select <b>Protocols &gt; Bgp</b>.</li> <li>2. In the Group box, click <b>Add new entry</b>.</li> <li>3. In the Group name box, type the name of the group in which the BGP peer is configured—<b>cluster-peers</b> in this case.</li> <li>4. In the Neighbor box, click <b>Add new entry</b>.</li> <li>5. In the Address box, type the IP address of the route reflector, in dotted decimal notation—in this case, <b>192.168.40.4</b>.</li> <li>6. Click <b>OK</b>.</li> </ol> | <p>On a client Services Router in the cluster:</p> <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><b>edit protocols bgp</b></li> <li>2. Create the group <b>cluster-peers</b>, and add only the route reflector address to the group:<br/><b>set group cluster-peers neighbor 192.168.40.4</b></li> </ol> |
| Configure a routing policy to advertise BGP routes.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | See “Injecting OSPF Routes into the BGP Routing Table (Optional)” on page 464.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |                                                                                                                                                                                                                                                                                                                                                             |

## 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 67 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 67: Typical Network Using BGP Confederations



To configure the BGP confederations shown in Figure 67:

- 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 82.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To check the configuration, see “Verifying a BGP Configuration” on page 308.

Table 82: Configuring BGP Confederations

| Task                                                                                                                                                                                                | J-Web Configuration Editor                                                                                                    | CLI Configuration Editor                                                                     |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| Navigate to the <b>Routing-options</b> level in the configuration hierarchy.                                                                                                                        | In the configuration editor hierarchy, select <b>Routing-options</b> .                                                        | From the top of the configuration hierarchy, enter<br><br>edit routing-options               |
| Set the AS number to the sub-AS number <b>64512</b> .<br><br>The sub-AS number is a unique AS number that is usually taken from the pool of private AS numbers— <b>64512</b> through <b>65535</b> . | <ul style="list-style-type: none"><li>1. In the AS Number box, enter the sub-AS number.</li><li>2. Click <b>OK</b>.</li></ul> | Set the sub-AS number:<br><br>set autonomous-system 64512                                    |
| Navigate to the <b>Confederation</b> level in the configuration hierarchy.                                                                                                                          | In the configuration editor hierarchy, select <b>Routing-options &gt; Confederation</b> .                                     | From the top of the configuration hierarchy, enter<br><br>edit routing-options confederation |
| Set the confederation number to the AS number <b>17</b> .                                                                                                                                           | In the Confederation as box, enter <b>17</b> .                                                                                | Set the confederation AS number:<br><br>set 17                                               |

**Table 82: Configuring BGP Confederations (continued)**

| Task                                                                                                                                  | J-Web Configuration Editor                                                                                                                                                                                                                                | CLI Configuration Editor                                            |
|---------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------|
| 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"> <li>1. Next to Members, click <b>Add new entry</b>.</li> <li>2. In the Value box, enter the sub-ASs that are members of this confederation. Separate multiple sub-ASs with a space.</li> </ol>                         | Add members to the confederation:<br><br>set 17 members 64512 64513 |
| Using EBGp, configure the peering session between the confederations (from router A to router B in this example).                     | See “Configuring a Point-to-Point Peering Session (Required)” on page 299.                                                                                                                                                                                |                                                                     |
| When setting the peer AS number for these sessions, use the sub-AS number rather than the AS number.                                  |                                                                                                                                                                                                                                                           |                                                                     |
| 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"> <li>■ To configure an IBGP full mesh, see “Configuring BGP Within a Network (Required)” on page 302.</li> <li>■ To configure a route reflector, see “Configuring a Route Reflector (Optional)” on page 303.</li> </ul> |                                                                     |

## Verifying a BGP Configuration

To verify a BGP configuration, perform these tasks:

- Verifying BGP Neighbors on page 308
- Verifying BGP Groups on page 309
- Verifying BGP Summary Information on page 310
- Verifying Reachability of All Peers in a BGP Network on page 311

### 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.

**Sample Output** 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/bgppgr 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.

For more information about `show bgp neighbor`, 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.

**Sample Output** `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
4.4.4.4+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.

For more information about `show bgp group`, 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.

**Sample Output** `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/Rece
10.0.0.2 65002 88675 88652 0 2 42:38 2/4/0
10.0.0.3 65002 54528 54532 0 1 2w4d22h 0/0/0
10.0.0.4 65002 51597 51584 0 0 2w3d22h 2/2/0

```



|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>What It Means</b> | <p>The output shows a summary of BGP session information. Verify the following information:</p> <ul style="list-style-type: none"> <li>■ For <b>Groups</b>, the total number of configured groups is shown.</li> <li>■ For <b>Peers</b>, the total number of BGP peers is shown.</li> <li>■ For <b>Down Peers</b>, the total number of unestablished peers is 0. If this value is not zero, one or more peering sessions are not yet established.</li> <li>■ Under <b>Peer</b>, the IP address for each configured peer is shown.</li> <li>■ Under <b>AS</b>, the peer AS for each configured peer is correct.</li> <li>■ Under <b>Up/Dwn State</b>, 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 <b>Active</b>, it indicates a problem in the establishment of the BGP session.</li> </ul> |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

For more information about `show bgp summary`, see the *JUNOS Routing Protocols and Policies Command Reference*.

## Verifying Reachability of All Peers in a BGP Network

|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b>       | 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Action</b>        | <p>For each Services Router in the BGP network:</p> <ol style="list-style-type: none"> <li>1. In the J-Web interface, select <b>Diagnose &gt; Ping Host</b>.</li> <li>2. In the Remote Host box, type the name of a host for which you want to verify reachability from the Services Router.</li> <li>3. Click <b>Start</b>. Output appears on a separate page.</li> </ol>                                                                                                                                                                                                                            |
| <b>Sample Output</b> | <pre> 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 </pre>                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>What It Means</b> | <p>If a host is active, it generates an ICMP response. If this response is received, the round-trip time is listed in the <code>time</code> field. For more information about the ping output, see the <i>J-series Services Router Administration Guide</i>.</p> <p>For more information about using the J-Web interface to ping a host, see the <i>J-series Services Router Administration Guide</i>.</p> <p>For information about the <code>ping</code> command, see the <i>J-series Services Router Administration Guide</i> or the <i>JUNOS System Basics and Services Command Reference</i>.</p> |



## **Part 4**

# **Configuring Private Communications over Public Networks with MPLS**

- Multiprotocol Label Switching Overview on page 315
- Configuring Signaling Protocols for Traffic Engineering on page 331
- Configuring Virtual Private Networks on page 343
- Configuring CLNS VPNs on page 367
- Configuring IPSec for Secure Packet Exchange on page 379



## Chapter 12

# Multiprotocol Label Switching Overview

Multiprotocol Label Switching (MPLS) provides a framework for controlling traffic patterns across a network. The MPLS framework allows Services Routers to pass traffic through transit networks on paths that are independent of the individual routing protocols enabled throughout the network.

The MPLS framework supports traffic engineering and the creation of virtual private networks (VPNs). Traffic is engineered (controlled) primarily by the use of signaling protocols to establish label-switched paths (LSPs). VPN support includes Layer 2 and Layer 3 VPNs and Layer 2 circuits.

This chapter contains the following topics. For more information, see the *JUNOS Routing Protocols Configuration Guide*, *JUNOS MPLS Applications Configuration Guide*, and *JUNOS VPNs Configuration Guide*.

- MPLS and VPN Terms on page 315
- MPLS Overview on page 317
- Signaling Protocols Overview on page 322
- VPN Overview on page 326

## MPLS and VPN Terms

To understand MPLS and VPNs, become familiar with the terms defined in Table 83 .

**Table 83: MPLS and VPN Terms**

| Term                                   | Definition                                                                                                                                                                                 |
|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| color                                  | See <i>link coloring</i> .                                                                                                                                                                 |
| Constrained Shortest Path First (CSPF) | MPLS algorithm that has been modified to include specific restrictions for calculating the shortest path across the network.                                                               |
| customer edge (CE) device              | Services Router or switch in the customer's network that is connected to a service provider's provider edge (PE) router and participates in a Layer 3 VPN.                                 |
| Explicit Route Object (ERO)            | Extension to the Resource Reservation Protocol (RSVP) that allows an RSVP PATH message to traverse an explicit sequence of routers independently of conventional shortest-path IP routing. |

**Table 83: MPLS and VPN Terms (continued)**

| <b>Term</b>                                 | <b>Definition</b>                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>inbound router</b>                       | Entry point for a label-switched path (LSP). Each LSP must have exactly one inbound router that is different from the outbound router. Inbound routers are also known as ingress routers. See also <i>outbound router</i> .                                                                                                                                                                                           |
| <b>label</b>                                | In Multiprotocol Label Switching (MPLS), a 20-bit unsigned integer in the range 0 through 1,048,575, used to identify a packet traveling along a label-switched path (LSP).                                                                                                                                                                                                                                           |
| <b>Label Distribution Protocol (LDP)</b>    | Protocol for distributing labels in non-traffic-engineered applications. LDP allows Services Routers to establish label-switched paths (LSPs) through a network by mapping Network layer routing information directly to Data Link layer switched paths.                                                                                                                                                              |
| <b>label-switched path (LSP)</b>            | Sequence of Services Routers that cooperatively perform Multiprotocol Label Switching (MPLS) operations for a packet stream. The first router in an LSP is called the inbound router, and the last router in the path is called the outbound router. An LSP is a point-to-point, half-duplex connection from the inbound router to the outbound router. (The inbound and outbound routers cannot be the same router.) |
| <b>label-switching router (LSR)</b>         | Any Services Router that is part of an LSP.                                                                                                                                                                                                                                                                                                                                                                           |
| <b>Layer 2 circuit</b>                      | Point-to-point Layer 2 connection transported by means of Multiprotocol Label Switching (MPLS) or another tunneling technology on a service provider's network. Multiple Layer 2 circuits can be transported over a single label-switched path (LSP) tunnel between two provider edge (PE) routers.                                                                                                                   |
| <b>Layer 2 VPN</b>                          | Private network service among a set of customer sites that use a service provider's existing Multiprotocol Label Switching (MPLS) and IP network. One customer's data is separated from another's by software rather than hardware. In a Layer 2 VPN, the Layer 3 routing of customer traffic occurs within the <i>customer</i> network.                                                                              |
| <b>Layer 3 VPN</b>                          | Private network service among a set of customer sites that use a service provider's existing Multiprotocol Label Switching (MPLS) and IP network. One customer's routes and data are separated from another customer's routes and data by software rather than hardware. In a Layer 3 VPN, the Layer 3 routing of customer traffic occurs within the <i>service provider</i> network.                                 |
| <b>link coloring</b>                        | In Constrained Shortest Path First (CSPF) routing, a way to group Multiprotocol Label Switching (MPLS) interfaces for CSPF path selection by assigning a color identifier and number to each administrative group.                                                                                                                                                                                                    |
| <b>Multiprotocol Label Switching (MPLS)</b> | Method for engineering network traffic patterns by assigning short labels to network packets that describe how to forward the packets through the network.                                                                                                                                                                                                                                                            |
| <b>multiple push</b>                        | Addition by a Services Router of up to three labels to a packet as it enters a Multiprotocol Label Switching (MPLS) domain.                                                                                                                                                                                                                                                                                           |
| <b>outbound router</b>                      | Exit point for a label-switched path (LSP). Each LSP must have exactly one outbound router that is different from the inbound router. Outbound routers are also called egress routers. See also <i>inbound router</i> .                                                                                                                                                                                               |
| <b>penultimate hop popping (PHP)</b>        | Using the penultimate router rather than the outbound router in a label-switched path (LSP) to remove the Multiprotocol Label Switching (MPLS) label from a packet.                                                                                                                                                                                                                                                   |
| <b>penultimate router</b>                   | Second-to-last Services Router in an LSP. The penultimate router is responsible for label popping when penultimate hop popping (PHP) is configured.                                                                                                                                                                                                                                                                   |
| <b>pop</b>                                  | Removal by a Services Router of the top label from a packet as it exits the Multiprotocol Label Switching (MPLS) domain.                                                                                                                                                                                                                                                                                              |
| <b>provider edge (PE) router</b>            | Services Router in the service provider network that is connected to a customer edge (CE) device and participates in a virtual private network (VPN).                                                                                                                                                                                                                                                                 |

**Table 83: MPLS and VPN Terms (continued)**

| <b>Term</b>                                      | <b>Definition</b>                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>provider router</b>                           | Services Router in the service provider's network that does not attach to a customer edge (CE) device.                                                                                                                                                                                                                                                                  |
| <b>push</b>                                      | Addition of a label or stack of labels by a Services Router to a packet as it enters a Multiprotocol Label Switching (MPLS) domain.                                                                                                                                                                                                                                     |
| <b>Resource Reservation Protocol (RSVP)</b>      | Resource reservation setup protocol that interacts with integrated services on the Internet.                                                                                                                                                                                                                                                                            |
| <b>route distinguisher</b>                       | A 6-byte virtual private network (VPN) identifier that is prefixed to an IPv4 address to make it unique. The new address is part of the VPN-IPv4 address family, which is a Border Gateway Protocol (BGP) extension. A route distinguisher allows you to configure private addresses within the VPN by preventing any overlap with the private addresses in other VPNs. |
| <b>routing instance</b>                          | Collection of routing tables, their interfaces, and the routing protocol parameters that control the information they contain.                                                                                                                                                                                                                                          |
| <b>swap</b>                                      | Replacement by a Services Router of a label or stack of labels on a packet as it travels through a Multiprotocol Label Switching (MPLS) domain.                                                                                                                                                                                                                         |
| <b>swap and push</b>                             | Replacement and subsequent push by a Services Router of a label or stack of labels on a packet as it travels through a Multiprotocol Label Switching (MPLS) domain.                                                                                                                                                                                                     |
| <b>traffic engineering database (TED)</b>        | Database populated by label-switched path (LSP) information such as the network topology, current reservable bandwidth of links, and link colors. The traffic engineering database is used to determine Constrained Shortest Path First (CSPF) path selection.                                                                                                          |
| <b>transit router</b>                            | Any label-switching router (LSR) between the inbound and outbound Services Router of a label-switched path (LSP).                                                                                                                                                                                                                                                       |
| <b>virtual private network (VPN)</b>             | Private data network that uses a public TCP/IP network, typically the Internet, while maintaining privacy with a tunneling protocol, encryption, and security procedures.                                                                                                                                                                                               |
| <b>VPN routing and forwarding (VRF) instance</b> | Routing instance for a Layer 3 VPN implementation that consists of one or more routing tables, a derived forwarding table, the interfaces that use the forwarding table, and the policies and routing protocols that determine what goes into the forwarding table.                                                                                                     |

## MPLS Overview

Multiprotocol Label Switching (MPLS) is a method for engineering traffic patterns by assigning short labels to network packets that describe how to forward them through the network. MPLS is independent of routing tables or any routing protocol and can be used for unicast packets.

This overview contains the following topics:

- Label Switching on page 318
- Label-Switched Paths on page 318
- Label-Switching Routers on page 319
- Labels on page 320

- Label Operations on page 320
- Penultimate Hop Popping on page 321
- LSP Establishment on page 321

## ***Label Switching***

In a traditional IP network, packets are transmitted with an IP header that includes a source and destination address. When a router receives such a packet, it examines its forwarding tables for the next-hop address associated with the packet's destination address and forwards the packet to the next-hop location.

In an MPLS network, each packet is encapsulated with an MPLS header. When a router receives the packet, it copies the header as an index into a separate MPLS forwarding table. The MPLS forwarding table consists of pairs of inbound interfaces and path information. Each pair includes forwarding information that the router uses to forward the traffic and modify, when necessary, the MPLS header.

Because the MPLS forwarding table has far fewer entries than the more general forwarding table, the lookup consumes less processing time and processing power. The resultant savings in time and processing are a significant benefit for traffic that uses the network to transit between outside destinations only.


## ***Label-Switched Paths***

Label-switched paths (LSPs) are unidirectional routes through a network or autonomous system (AS). In normal IP routing, the packet has no predetermined path. Instead, each router forwards a packet to the next-hop address stored in its forwarding table, based only on the packet's destination address. Each subsequent router then forwards the packet using its own forwarding table.

In contrast, MPLS routers within an AS determine paths through a network through the exchange of MPLS traffic engineering information. Using these paths, the routers direct traffic through the network along an established route. Rather than selecting the next hop along the path as in IP routing, each router is responsible for forwarding the packet to a predetermined next-hop address.

Figure 68 shows a typical LSP topology.



**Figure 68: Typical LSP Topology**

In the topology shown in Figure 68, traffic is forwarded from host C1 to the transit network with standard IP forwarding. When the traffic enters the transit network, it is switched across a preestablished LSP through the network. In this example, an LSP might switch the traffic from router R4 to router R2 to router R1. When the traffic exits the network, it is forwarded to its destination by IP routing protocols.

### Label-Switching Routers

Routers that are part of the LSP are label-switching routers (LSRs). Each LSR must be configured with MPLS so that it can interpret MPLS headers and perform the MPLS operations required to pass traffic through the network. An LSP can include four types of LSRs:

- Inbound router—The only entry point for traffic into MPLS. Native IPv4 packets are encapsulated into the MPLS protocol by the inbound router. Each LSP can have only one inbound router.
- Transit router—Any router in the middle of an LSP. An individual LSP can contain between 0 and 253 transit routers. Transit routers forward MPLS traffic along the LSP, using only the MPLS header to determine how the packet is routed.
- Penultimate router—The second-to-last router in the LSP. The penultimate router in an LSP is responsible for stripping the MPLS header from the packet before forwarding it to the outbound router.
- Outbound router—The endpoint for the LSP. The outbound router receives MPLS packets from the penultimate router and performs an IP route lookup.

The router then forwards the packet to the next hop of the route. Each LSP can have only one outbound router.

## **Labels**

To forward traffic through an MPLS network, MPLS routers encapsulate packets and assign and manage headers known as labels. The routers use the labels to index the MPLS forwarding tables that determine how packets are routed through the network.

When a network's inbound router receives traffic, it inserts an MPLS label between the IP packet and the appropriate Layer 2 header for the physical link. The label contains an index value that identifies a next-hop address for the particular LSP. When the next-hop transit router receives the packet, it uses the index in the MPLS label to determine the next-hop address for the packet and forwards the packet to the next router in the LSP.

As each packet travels through the transit network, every router along the way performs a lookup on the MPLS label and forwards the packet accordingly. When the outbound router receives a packet, it examines the header to determine that it is the final router in the LSP. The outbound router then removes the MPLS header, performs a regular IP route lookup, and forwards the packet with its IP header to the next-hop address.

## **Label Operations**

Each LSR along an LSP is responsible for examining the MPLS label, determining the LSP next hop, and performing the required label operations. LSRs can perform five label operations:

- **Push**—Adds a new label to the top of the packet. For IPv4 packets arriving at the inbound router, the new label is the first label in the label stack. For MPLS packets with an existing label, this operation adds a label to the stack and sets the stacking bit to 0, indicating that more MPLS labels follow the first.

When it receives the packet, the inbound router performs an IP route lookup on the packet. Because the route lookup yields an LSP next hop, the inbound router performs a label push on the packet, and then forwards the packet to the LSP next hop.

- **Swap**—Replaces the label at the top of the label stack with a new label.

When a transit router receives the packet, it performs an MPLS forwarding table lookup. The lookup yields the LSP next hop and the path index of the link between the transit router and the next router in the LSP.

- **Pop**—Removes the label from the top of the label stack. For IPv4 packets arriving at the penultimate router, the entire MPLS label is removed from the label stack. For MPLS packets with an existing label, this operation removes the top label from the label stack and modifies the stacking bit as necessary—sets it to 1, for example, if only a single label remains in the stack.

If multiple LSPs terminate at the same outbound router, the router performs MPLS label operations for all outbound traffic on the LSPs. To share the operations among multiple routers, most LSPs use penultimate hop popping (PHP).

- Multiple push—Adds multiple labels to the top of the label stack. This action is equivalent to performing multiple push operations.

The multiple push operation is used with label stacking, which is beyond the scope of this guide.

- Swap and push—Replaces the top label with a new label and then pushes a new label to the top of the stack.

The swap and push operation is used with label stacking, which is beyond the scope of this guide.

## **Penultimate Hop Popping**

Multiple LSPs terminating at a single outbound router load the router with MPLS label operations for all their outbound traffic. Penultimate hop popping (PHP) transfers the operation from the outbound router to penultimate routers.

With PHP, the penultimate router is responsible for popping the MPLS label and forwarding the traffic to the outbound router. The outbound router then performs an IP route lookup and forwards the traffic. For example, if four LSPs terminate at the same outbound router and each has a different penultimate router, label operations are shared across four routers.

## **LSP Establishment**

An MPLS LSP is established by one of two methods: static LSPs and dynamic LSPs.

### **Static LSPs**

Like a static route, a static LSP requires each router along the path to be configured explicitly. You must manually configure the path and its associated label values. Static LSPs require less processing by the LSRs because no signaling protocol is used. However, because paths are statically configured, they cannot adapt to network conditions. Topology changes and network outages can create black holes in the LSP that exist until you manually reconfigure the LSP.

### **Dynamic LSPs**

Dynamic LSPs use signaling protocols to establish themselves and propagate LSP information to other LSRs in the network. You configure the inbound router with LSP information that is transmitted throughout the network when you enable the signaling protocols across the LSRs. Because the LSRs must exchange and process signaling packets and instructions, dynamic LSPs consume more resources than static LSPs. However, dynamic LSPs can avoid

the network black holes of static LSPs by detecting topology changes and outages and propagating them throughout the network.

## Signaling Protocols Overview

---

Two MPLS signaling protocols are used to dynamically establish and maintain LSPs within a network:

- Label Distribution Protocol on page 322
- Resource Reservation Protocol on page 322

### ***Label Distribution Protocol***

LDP is a simple, fast-acting signaling protocol that automatically establishes LSP adjacencies within an MPLS network. Routers then share LSP updates such as hello packets and LSP advertisements across the adjacencies.

#### **LDP Operation**

Because LDP runs on top of an interior gateway protocol (IGP) such as IS-IS or OSPF, you must configure LDP and the IGP on the same set of interfaces. After both are configured, LDP begins transmitting and receiving LDP messages through all LDP-enabled interfaces.

Because of LDP's simplicity, it cannot perform true traffic engineering like RSVP. LDP does not support bandwidth reservation or traffic constraints.

#### **LDP Messages**

When you configure LDP on an LSR, the router begins sending LDP discovery messages out all LDP-enabled interfaces. When an adjacent LSR receives LDP discovery messages, it establishes an underlying TCP session. An LDP session is then created on top of the TCP session. The TCP three-way handshake ensures that the LDP session has bidirectional connectivity. After they establish the LDP session, the LDP neighbors maintain, and terminate, the session by exchanging messages.

LDP advertisement messages allow LSRs to exchange label information to determine the next hops within a particular LSP.

Any topology changes, such as a router failure, generate LDP notifications that can terminate the LDP session or generate additional LDP advertisements to propagate an LSP change.

### ***Resource Reservation Protocol***

Resource Reservation Protocol (RSVP) is a signaling protocol that handles bandwidth allocation and true traffic engineering across an MPLS network. Like LDP, RSVP uses discovery messages and advertisements to exchange LSP path

information between all hosts. However, RSVP also includes a set of features that control the flow of traffic through an MPLS network.

This section contains the following topics:

- “RSVP Fundamentals” on page 323
- “Bandwidth Reservation Requirement” on page 323
- “Explicit Route Objects” on page 323
- “Constrained Shortest Path First” on page 325
- “Link Coloring” on page 325

## **RSVP Fundamentals**

RSVP uses unidirectional and simplex flows through the network to perform its function. The inbound router initiates an RSVP path message and sends it downstream to the outbound router. The path message contains information about the resources needed for the connection. Each router along the path begins to maintain information about a reservation.

When the path message reaches the outbound router, resource reservation begins. The outbound router sends a reservation message upstream to the inbound router. Each router along the path receives the reservation message and sends it upstream, following the path of the original path message. When the inbound router receives the reservation message, the unidirectional network path is established.

The established path remains open as long as the RSVP session is active. The session is maintained by the transmission of additional path and reservation messages that report the session state every 30 seconds. If a router does not receive the maintenance messages for 3 minutes, it terminates the RSVP session and reroutes the LSP through another active router.

## **Bandwidth Reservation Requirement**

When a bandwidth reservation is configured, reservation messages propagate the bandwidth value throughout the LSP. Routers must reserve the bandwidth specified across the link for the particular LSP. If the total bandwidth reservation exceeds the available bandwidth for a particular LSP segment, the LSP is rerouted through another LSR. If no segments can support the bandwidth reservation, LSP setup fails and the RSVP session is not established.

## **Explicit Route Objects**

Explicit Route Objects (EROs) limit LSP routing to a specified list of LSRs. By default, RSVP messages follow a path that is determined by the network IGP's shortest path. However, in the presence of a configured ERO, the RSVP messages follow the path specified.


EROs consist of two types of instructions: loose hops and strict hops. When a loose hop is configured, it identifies one or more transit LSRs through which the LSP must be routed. The network IGP determines the exact route from the inbound router to the first loose hop, or from one loose hop to the next. The loose hop specifies only that a particular LSR be included in the LSP.

When a strict hop is configured, it identifies an exact path through which the LSP must be routed. Strict-hop EROs specify the exact order of routers through which the RSVP messages are sent.

You can configure loose-hop and strict-hop EROs simultaneously. In this case, the IGP determines the route between loose hops, and the strict-hop configuration specifies the exact path for particular LSP path segments.

Figure 69 shows a typical RSVP-signaled LSP that uses EROs.

**Figure 69: Typical RSVP-Signaled LSP with EROs**



In the topology shown in Figure 69, traffic is routed from host C1 to host C2. The LSP can pass through router R4 or router R7. To force the LSP to use R4, you can set up either a loose-hop or strict-hop ERO that specifies R4 as a hop in the LSP. To force a specific path through routers R4, R3, and R6, configure a strict-hop ERO through the three LSRs.

## Constrained Shortest Path First

Whereas IGPs use the Shortest Path First (SPF) algorithm to determine how traffic is routed within a network, RSVP uses the Constrained Shortest Path First (CSPF) algorithm to calculate traffic paths that are subject to the following constraints:

- LSP attributes—Administrative groups such as link coloring, bandwidth requirements, and EROs
- Link attributes—Colors on a particular link and available bandwidth

These constraints are maintained in the traffic engineering database (TED). The database provides CSPF with up-to-date topology information, the current reservable bandwidth of links, and the link colors.

In determining which path to select, CSPF follows these rules:

1. Computes LSPs one at a time, beginning with the highest-priority LSP—the one with the lowest setup priority value. Among LSPs of equal priority, CSPF starts with those that have the highest bandwidth requirement.
2. Prunes the traffic engineering database of links that are not full duplex and do not have sufficient reservable bandwidth.
3. If the LSP configuration includes the `include` statement, prunes all links that do not share any included colors.
4. If the LSP configuration includes the `exclude` statement, prunes all links that contain excluded colors. If a link does not have a color, it is accepted.
5. Finds the shortest path toward the LSP's outbound router, taking into account any EROs. For example, if the path must pass through **router A**, two separate SPF algorithms are computed: one from the inbound router to **router A** and one from **router A** to the outbound router.
6. If several paths have equal cost, chooses the one with a last-hop address the same as the LSP's destination.
7. If several equal-cost paths remain, selects the path with the fewest number of hops.
8. If several equal-cost paths remain, applies CSPF load-balancing rules configured on the LSP.

## Link Coloring

RSVP allows you to configure administrative groups for CSPF path selection. An administrative group is typically named with a color, assigned a numeric value, and applied to the RSVP interface for the appropriate link. Lower numbers indicate higher priority.

After configuring the administrative group, you can either exclude, include, or ignore links of that color in the traffic engineering database:

- If you exclude a particular color, all segments with an administrative group of that color are excluded from CSPF path selection.
- If you include a particular color, only those segments with the appropriate color are selected.
- If you neither exclude nor include the color, the metrics associated with the administrative groups and applied on the particular segments determine the path cost for that segment.

The LSP with the lowest total path cost is selected into the traffic engineering database.

## VPN Overview

---

Virtual private networks (VPNs) are private networks that use a public network to connect two or more remote sites. In place of dedicated connections between networks, VPNs use virtual connections routed (tunneled) through public networks that are typically service provider networks. The type of the VPN is determined by the connections it uses and whether the customer network or the provider network performs the virtual tunneling.

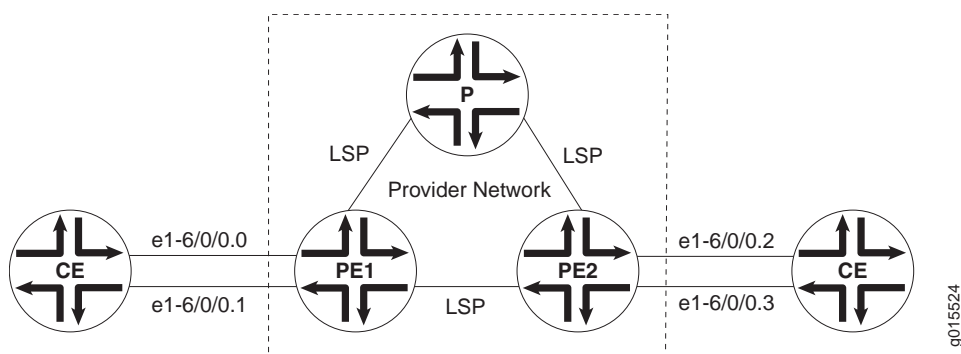
This overview contains the following topics:

- VPN Components on page 326
- VPN Routing Requirements on page 327
- VPN Routing Information on page 328
- Types of VPNs on page 329

## VPN Components

All types of VPNs share certain components. Figure 70 shows a typical VPN topology.



**Figure 70: Typical VPN Topology**

The provider edge (PE) routers in the provider's network connect to the customer edge (CE) devices located at customer sites. PE routers support VPN and MPLS label functionality. Within a single VPN, pairs of PE routers are connected through a virtual tunnel, typically an LSP.

Provider routers within the core of the provider's network are not connected to any routers at a customer site but are part of the tunnel between pairs of PE routers. Provider routers support LSP functionality as part of the tunnel support, but do not support VPN functionality.

Customer edge (CE) devices are the routers or switches located at the customer site that connect to the provider's network. CE devices are typically IP routers, but they can also be Asynchronous Transfer Mode (ATM), Frame Relay, or Ethernet switches.

All VPN functions are performed by the PE routers. Neither CE devices nor provider routers are required to perform any VPN functions.

## VPN Routing Requirements

VPNs tunnel traffic as follows from one customer site to another customer site, using a public network as a transit network, when certain requirements are met:

1. Traffic is forwarded by standard IP forwarding from the CE devices to the PE routers.

The CE devices require only a BGP connection to the PE routers.

2. The PE routers establish an LSP through the provider network.

The provider network must be running either OSPF or IS-IS as an IGP, as well as IBGP sessions through either a full mesh or route reflector. IBGP is required so that the PE routers can exchange route information for routes that originate or terminate in the VPN.

3. When the inbound PE router receives traffic, it performs a route lookup. The lookup yields an LSP next hop, and the traffic is forwarded along the LSP.

Either LDP or RSVP must be configured to dynamically set up LSPs through the provider network.

4. When the traffic reaches the outbound PE router, the PE router pops the MPLS label and forwards the traffic with standard IP routing.

Because the tunnel information is maintained at both PE routers, neither the provider core routers nor the CE devices need to maintain any VPN information in their configuration databases.

## **VPN Routing Information**

Routing information, including routes, route distinguishers, and routing policies, is stored in a VPN routing and forwarding (VRF) table. Routers must maintain separate VRF tables for each VPN.

### **VRF Instances**

A routing instance is a collection of routing tables, interfaces, and routing protocol parameters. The interfaces belong to the routing tables, and the routing protocol parameters control the information in the routing tables. In the case of VPNs, each VPN has a VPN routing and forwarding (VRF) instance.

A VRF instance consists of one or more routing tables, a derived forwarding table, the interfaces that use the forwarding table, and the policies and routing protocols that determine what goes into the forwarding table. Because each instance is configured for a particular VPN, each VPN has separate tables, rules, and policies that control its operation.

A separate VRF table is created for each VPN that has a connection to a CE router. The VRF table is populated with routes received from directly connected CE sites associated with the VRF instance, and with routes received from other PE routers in the same VPN.

### **Route Distinguishers**

Because a typical transit network is configured to handle more than one VPN, the provider routers are likely to have multiple VRF instances configured. As a result, depending on the origin of the traffic and any filtering rules applied to the traffic, the BGP routing tables can contain multiple routes for a particular destination address. Because BGP requires that exactly one BGP route per destination be imported into the forwarding table, BGP must have a way to distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPNs.

A route distinguisher is a locally unique number that identifies all route information for a particular VPN. Unique numeric identifiers allow BGP to distinguish between routes that are otherwise identical.

## Route Targets to Control the VRF Table

On each PE router, you must define routing policies that specify how routes are imported into and exported from the router's VRF table. Each advertisement must have an associated route target that uniquely identifies the VPN for which the advertisement is valid. The route target allows you to keep routing and signaling information for each VPN separate.

## Types of VPNs

There are three primary types of VPNs: Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs.

### Layer 2 VPNs

In a Layer 2 VPN, traffic is forwarded to the PE router in Layer 2 format, carried by MPLS through an LSP over the service provider network, and then converted back to Layer 2 format at the receiving CE device.

On a Layer 2 VPN, routing occurs on the customer routers, typically on the CE router. The CE router connected to a service provider on a Layer 2 VPN must select the appropriate circuit on which to send traffic. The PE router receiving the traffic sends it across the network to the PE router on the outbound side. The PE routers need no information about the customer's routes or routing topology, and need only to determine the virtual tunnel through which to send the traffic.

### Layer 2 Circuits

A Layer 2 circuit is a point-to-point Layer 2 connection that transports traffic by MPLS or another tunneling technology on a service provider network. The Layer 2 circuit creates a virtual connection to direct traffic between two CE routers. The primary difference between a Layer 2 circuit and an Layer 2 VPN is the method of setting up the virtual connection. Like a leased line, a Layer 2 circuit forwards all packets received from the local interface to the remote interface.

### Layer 3 VPNs

In a Layer 3 VPN, routing occurs on the service provider's routers. As a result, Layer 3 VPNs require information about customer routes and a more extensive VRF policy configuration to share and filter routes that originate or terminate in the VPN.

Because Layer 3 VPNs require the provider routers to route and forward VPN traffic at the entry and exit points of the transit network, the routes must be advertised and filtered throughout the provider network.

Route advertisements originate at the CE devices and are shared with the inbound PE routers through standard IP routing protocols, typically BGP. Based on the source address, the PE router filters route advertisements and imports them into the appropriate VRF table.

The PE router then exports the route in IBGP sessions to the other provider routers. Route export is governed by any routing policy that has been applied to the particular VRF table. To propagate the routes through the provider network, the PE router must also convert the route to VPN format, which includes the route distinguisher.

When the outbound PE router receives the route, it strips off the route distinguisher and advertises the route to the connected CE device, typically through standard BGP IPv4 route advertisements.

## Chapter 13

# Configuring Signaling Protocols for Traffic Engineering

Signaling protocols are used within a Multiprotocol Label Switching (MPLS) environment to establish label-switched paths (LSPs) for traffic across a transit network. J-series Services Routers support the Label Distribution Protocol (LDP) and the Resource Reservation Protocol (RSVP) as part of their suite of traffic engineering features.

You can use either the J-Web configuration editor or CLI configuration editor to configure signaling protocols.

This chapter contains the following topics. For more information about MPLS traffic engineering, see the *JUNOS MPLS Applications Configuration Guide*.

- Signaling Protocol Overview on page 331
- Before You Begin on page 332
- Configuring LDP and RSVP with a Configuration Editor on page 333
- Verifying an MPLS Configuration on page 338

## Signaling Protocol Overview

---

When transit traffic is routed through an IP network, MPLS is often used to engineer its passage. Although the exact path through the transit network is of little importance to either the sender or the receiver of the traffic, network administrators often want to route traffic more efficiently between certain source and destination address pairs. By adding a short label with specific routing instructions to each packet, MPLS switches packets from router to router through the network rather than forwarding packets based on next-hop lookups. The resulting routes are called label-switched paths (LSPs). LSPs control the passage of traffic through the network and speed traffic forwarding.

You can create LSPs manually, or through the use of signaling protocols. Services Routers support two signaling protocols—the Label Distribution Protocol (LDP) and the Resource Reservation Protocol (RSVP).

## LDP Signaling Protocol

The Label Distribution Protocol (LDP) is a signaling protocol that runs on a Services Router configured for MPLS support. The LDP configuration is added to the existing interior gateway protocol (IGP) configuration and included in the MPLS configuration. To configure a network to use LDP for LSP establishment, you first enable MPLS on all transit interfaces in the MPLS network and then enable LDP sessions on the interfaces.

The successful configuration of both MPLS and LDP initiates the exchange of TCP packets across the LDP interfaces. The packets establish TCP-based LDP sessions for the exchange of MPLS information within the network. Enabling both MPLS and LDP on the appropriate interfaces is sufficient to establish LSPs.

## RSVP Signaling Protocol

The Resource Reservation Protocol (RSVP) is a more flexible and powerful way to engineer traffic through a transit network. Like LDP, RSVP establishes LSPs within an MPLS network when you enable both MPLS and RSVP on the appropriate interfaces. However, whereas LDP is restricted to using the configured IGP's shortest path as the transit path through the network, RSVP uses a combination of the Constrained Shortest Path First (CSPF) algorithm and Explicit Route Objects (EROs) to determine how traffic is routed through the network.

Basic RSVP sessions are established in exactly the same way that LDP sessions are established. By configuring both MPLS and RSVP on the appropriate transit interfaces, you enable the exchange of RSVP packets and the establishment of LSPs. However, RSVP also lets you configure link authentication, explicit LSP paths, and link coloring. For more information about these topics, see the *JUNOS MPLS Applications Configuration Guide*.

## Before You Begin

---

Before you begin configuring signaling protocols for traffic engineering, complete the following tasks:

- Establish basic connectivity. See the *J-series Services Router Getting Started Guide*.
- Configure network interfaces. See “Configuring Network Interfaces” on page 103.
- Configure an interior gateway protocol (IGP) across your network. See “Configuring an OSPF Network” on page 265 or “Configuring a RIP Network” on page 249. For information about the IS-IS IGP, see the *JUNOS Routing Protocols Configuration Guide*.

## Configuring LDP and RSVP with a Configuration Editor

To configure either LDP or RSVP as a signaling protocol on the Services Router to establish LSPs through an IP network, perform one of the following tasks:


- Configuring LDP-Signaled LSPs on page 333
- Configuring RSVP-Signaled LSPs on page 335

For information about using the J-Web and CLI configuration editors, see “Using J-Web Configuration Tools” on page 3.

### Configuring LDP-Signaled LSPs

Using LDP as a signaling protocol, you create LSPs between Services Routers in an IP network. A sample network is shown in Figure 71.

**Figure 71: Typical LDP-Signaled LSP**



To establish an LSP between Services Routers R6 and R7, you must configure LDP on Services Routers R5, R6, and R7. This configuration ensures that hosts C1 and C2 use the LDP-signaled LSP when the entry (ingress) router is R6 or R7.

To configure LDP to establish the LSP shown in Figure 71, 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 84.

3. If you are finished configuring the router, commit the configuration.
4. Go on to “Verifying an LDP-Signaled LSP” on page 338.

**Table 84: Configuring an LDP-Signaled LSP**

| <b>Task</b>                                                                          | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level of the configuration hierarchy               | In the configuration editor hierarchy, select <b>Interfaces</b> .                                                                                                                                                                                                                                                                                                                                                | From the top of the configuration hierarchy, enter<br><br>edit interfaces                                                                                                                                                                                                                       |
| Enable the MPLS family on all transit interfaces on each router in the MPLS network. | <ol style="list-style-type: none"> <li>1. Click the transit interface on which you want to configure MPLS.</li> <li>2. In the Unit table, click the unit number for which you want to enable MPLS.</li> <li>3. In the Family area, select the <b>Mpls</b> check box.</li> <li>4. Click <b>OK</b>.</li> <li>5. Repeat Steps 1 through 4 for each transit interface on the routers in the MPLS network.</li> </ol> | <ol style="list-style-type: none"> <li>1. Add the MPLS family to all transit interfaces. For example:<br/><br/><b>set fe-0/0/0 unit 0 family mpls</b></li> <li>2. Repeat Step 1 for each transit interface on the routers in the MPLS network.</li> </ol>                                       |
| Enable the MPLS process on all MPLS interfaces for each router in the MPLS network.  | <ol style="list-style-type: none"> <li>1. Navigate to the <b>Protocols &gt; Mpls</b> level in the configuration hierarchy.</li> <li>2. Next to Interface, click <b>Add new entry</b>.</li> <li>3. In the Interface name box, type <b>all</b>.</li> <li>4. Click <b>OK</b>.</li> <li>5. Repeat Steps 1 through 4 for each transit interface on the routers in the MPLS network.</li> </ol>                        | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><br/><b>edit protocols mpls</b></li> <li>2. Enter<br/><br/><b>set interface all</b></li> <li>3. Repeat Steps 1 and 2 for each transit interface on the routers in the MPLS network.</li> </ol> |




**Table 84: Configuring an LDP-Signaled LSP (continued)**

| <b>Task</b>                                                                                                                                                               | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                          | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                                                                     |
|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create the LDP instance on each Services Router in the MPLS network.                                                                                                      | <ol style="list-style-type: none"> <li>1. Navigate to the <b>Protocols &gt; Ldp</b> level in the configuration hierarchy.</li> <li>2. Next to Interface, click <b>Add new entry</b>.</li> <li>3. In the Interface name box, type the name of a transit interface—for example, <b>fe-0/0/0</b>.</li> <li>4. Click <b>OK</b>.</li> <li>5. Repeat Steps 1 through 4 for each transit interface on the routers in the MPLS network.</li> </ol> | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><b>edit protocols ldp</b></li> <li>2. Enable LDP on a transit interface. For example:<br/><b>set interface fe-0/0/0</b></li> <li>3. Repeat Steps 1 and 2 for each transit interface on the routers in the MPLS network.</li> </ol> |
| Set the keepalive interval to 5 seconds.<br><br>The keepalive interval specifies the number of seconds between the transmission of keepalive messages along the LDP link. | <ol style="list-style-type: none"> <li>1. In the Keepalive interval box, type <b>5</b>.</li> <li>2. Click <b>OK</b>.</li> <li>3. Repeat Steps 1 and 2 for each router in the MPLS network.</li> </ol>                                                                                                                                                                                                                                      | <p>On each router in the MPLS network, enter</p> <p><b>set keepalive-interval 5</b></p>                                                                                                                                                                                                                                             |

### Configuring RSVP-Signaled LSPs

Using RSVP as a signaling protocol, you create LSPs between Services Routers in an IP network. A sample network is shown in Figure 72.

**Figure 72: Typical RSVP-Signaled LSP**

To establish an LSP between Services Routers R1 and R7, you must configure RSVP on all MPLS transit interfaces in the network. This configuration ensures that hosts C1 and C2 use the RSVP-signaled LSP corresponding to the network IGP's shortest path. Additionally, this configuration reserves 10 Mbps of bandwidth.

To configure RSVP to establish the LSP shown in Figure 72, 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 85.
3. If you are finished configuring the router, commit the configuration.
4. Go on to “Verifying an RSVP-Signaled LSP” on page 341.

**Table 85: Configuring an RSVP-Signaled LSP**

| Task                                                                   | J-Web Configuration Editor                                        | CLI Configuration Editor                                                  |
|------------------------------------------------------------------------|-------------------------------------------------------------------|---------------------------------------------------------------------------|
| Navigate to the <b>Interfaces</b> level of the configuration hierarchy | In the configuration editor hierarchy, select <b>Interfaces</b> . | From the top of the configuration hierarchy, enter<br><br>edit interfaces |

**Table 85: Configuring an RSVP-Signaled LSP (continued)**

| <b>Task</b>                                                                                                                              | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                           | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enable the MPLS family on all transit interfaces on each router in the MPLS network.                                                     | <ol style="list-style-type: none"> <li>1. Click the transit interface on which you want to configure MPLS.</li> <li>2. In the Unit table, click the unit number for which you want to enable MPLS.</li> <li>3. In the Family area, select the <b>Mpls</b> check box.</li> <li>4. Click <b>OK</b>.</li> <li>5. Repeat Steps 1 through 4 for each transit interface on the routers in the MPLS network.</li> </ol>                            | <ol style="list-style-type: none"> <li>1. Add the MPLS family to all transit interfaces. For example:<br/><br/><b>set fe-0/0/0 unit 0 family mpls</b></li> <li>2. Repeat Step 1 for each transit interface on the routers in the MPLS network.</li> </ol>                                                                                       |
| Enable the MPLS process on all MPLS interfaces for each router in the MPLS network.                                                      | <ol style="list-style-type: none"> <li>1. Navigate to the <b>Protocols &gt; Mpls</b> level in the configuration hierarchy.</li> <li>2. Next to Interface, click <b>Add new entry</b>.</li> <li>3. In the Interface name box, type <b>all</b>.</li> <li>4. Click <b>OK</b>.</li> <li>5. Repeat Steps 1 through 4 for each transit interface on the routers in the MPLS network.</li> </ol>                                                   | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><br/><b>edit protocols mpls</b></li> <li>2. Enter:<br/><br/><b>set interface all</b></li> <li>3. Repeat Steps 1 and 2 for each transit interface on the routers in the MPLS network.</li> </ol>                                                |
| Create the RSVP instance on each Services Router in the MPLS network.                                                                    | <ol style="list-style-type: none"> <li>1. Navigate to the <b>Protocols &gt; Rsvp</b> level in the configuration hierarchy.</li> <li>2. Next to Interface, click <b>Add new entry</b>.</li> <li>3. In the Interface name box, type the name of a transit interface—for example, <b>fe-0/0/0</b>.</li> <li>4. Click <b>OK</b>.</li> <li>5. Repeat Steps 1 through 4 for each transit interface on the routers in the MPLS network.</li> </ol> | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><br/><b>edit protocols rsvp</b></li> <li>2. Enable RSVP on a transit interface. For example:<br/><br/><b>set interface fe-0/0/0</b></li> <li>3. Repeat Steps 1 and 2 for each transit interface on the routers in the MPLS network.</li> </ol> |
| On the entry (ingress) router, <b>R1</b> , define the LSP <b>r1-r7</b> , using router <b>R7</b> 's loopback address ( <b>10.0.9.7</b> ). | <ol style="list-style-type: none"> <li>1. Navigate to the <b>Protocols &gt; Mpls</b> level in the configuration hierarchy.</li> <li>2. Next to Label switched path, click <b>Add new entry</b>.</li> <li>3. In the Path name box, type <b>r1-r7</b>.</li> <li>4. In the To box, type <b>10.0.9.7</b>.</li> </ol>                                                                                                                            | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><br/><b>edit protocols mpls</b></li> <li>2. Enter<br/><br/><b>set label-switched-path r1-r7 to 10.0.9.7</b></li> </ol>                                                                                                                         |

**Table 85: Configuring an RSVP-Signaled LSP (continued)**

| <b>Task</b>                                                                                                                                                                                                      | <b>J-Web Configuration Editor</b>                                                                                                                                     | <b>CLI Configuration Editor</b>                          |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------|
| Reserve 10 Mbps of bandwidth on the LSP.                                                                                                                                                                         | <ol style="list-style-type: none"> <li>1. In the Bandwidth box, click <b>Configure</b>.</li> <li>2. In the Ct0 box, type 10m.</li> <li>3. Click <b>OK</b>.</li> </ol> | Enter<br><br>set label-switched-path r1-r7 bandwidth 10m |
| Disable the use of the Constrained Shortest Path First (CSPF) algorithm.<br><br>By disabling the CSPF algorithm, you specify that traffic through the LSP is to be routed along the network IGP's shortest path. | <ol style="list-style-type: none"> <li>1. Select the <b>No cspf</b> check box.</li> <li>2. Click <b>OK</b>.</li> </ol>                                                | Enter<br><br>set label-switched-path r1-r7 no-cspf       |

## Verifying an MPLS Configuration

The tasks required to verify your MPLS configuration depend on the signaling protocol used. To validate the configuration, perform the appropriate set of tasks:

- Verifying an LDP-Signaled LSP on page 338
- Verifying an RSVP-Signaled LSP on page 341

### Verifying an LDP-Signaled LSP

Suppose that LDP is configured to establish an LSP as shown in Figure 71.

To verify the LDP configuration, perform these verification tasks:

- “Verifying LDP Neighbors” on page 338
- “Verifying LDP Sessions” on page 339
- “Verifying the Presence of LDP-Signaled LSPs” on page 340
- “Verifying Traffic Forwarding over the LDP-Signaled LSP” on page 340

### Verifying LDP Neighbors

**Purpose** Verify that each Services Router shows the appropriate LDP neighbors—for example, that router R5 has both router R6 and router R7 as LDP neighbors.

**Action** From the CLI, enter the show ldp neighbor command.

**Sample Output**      `user@r5> show ldp neighbor`

| Address   | Interface  | Label space ID | Hold time |
|-----------|------------|----------------|-----------|
| 10.0.8.5  | fe-0/0/0.0 | 10.0.9.6:0     | 14        |
| 10.0.8.10 | fe-0/0/1.0 | 10.0.9.7:0     | 11        |

**What It Means**      The output shows the IP addresses of the neighboring interfaces along with the interface through which the neighbor adjacency is established. Verify the following information:

- Each interface on which LDP is enabled is listed.
- Each neighboring LDP interface address is listed with the appropriate corresponding LDP interface.
- Under Label space ID, the appropriate loopback address for each neighbor appears.

## Verifying LDP Sessions

**Purpose**      Verify that a TCP-based LDP session has been established between all LDP neighbors. Also, verify that the modified keepalive value is active.

**Action**      From the CLI, enter the `show ldp session detail` command.

**Sample Output**      `user@r5> show ldp session detail`

```
Address: 10.0.9.7, State: Operational, Connection: Open, Hold time: 28
 Session ID: 10.0.3.5:0--10.0.9.7:0
 Next keepalive in 3 seconds
 Passive, Maximum PDU: 4096, Hold time: 30, Neighbor count: 1
 Keepalive interval: 5, Connect retry interval: 1
 Local - Restart: disabled, Helper mode: enabled
 Remote - Restart: disabled, Helper mode: disabled
 Local maximum recovery time: 240000 msec
 Next-hop addresses received:
 10.0.8.10
 10.0.2.17
```

- What It Means** The output shows the detailed information, including session IDs, keepalive interval, and next-hop addresses, for each established LDP session. Verify the following information:
- Each LDP neighbor address has an entry, listed by loopback address.
  - The state for each session is **Operational**, and the connection for each session is **Open**. A state of **Nonexistent** or a connection of **Closed** indicates a problem with one of the following:
    - LDP configuration
    - Passage of traffic between the two Services Routers
    - Physical link between the two routers
  - For Keepalive interval, the appropriate value, 5, appears.

### Verifying the Presence of LDP-Signaled LSPs

- Purpose** Verify that each Services Router's **inet.3** routing table has an LSP for the loopback address on each of the other routers.
- Action** From the CLI, enter the **show route table inet.3** command.
- Sample Output**
- ```
user@r5> show route table inet.3

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.9.6/32          *[LDP/9/0] 00:05:29, metric 1
                    > to 10.0.8.5 via fe-0/0/0.0
10.0.9.7/32          *[LDP/9/0] 00:05:37, metric 1
                    > to 10.0.8.10 via fe-0/0/1.0
```
- What It Means** The output shows the LDP routes that exist in the **inet.3** routing table. Verify that an LDP-signaled LSP is associated with the loopback addresses of the other routers in the MPLS network.

Verifying Traffic Forwarding over the LDP-Signaled LSP

- Purpose** Verify that traffic between hosts **C1** and **C2** is forwarded over the LDP-signaled LSP between Services Router **R6** and Services Router **R7**. Because traffic uses any configured gateway address by default, you must explicitly specify that the gateway address is to be bypassed.
- Action** If host **C1** is a Juniper Networks router, from the CLI enter the **traceroute 220.220.0.0 source 200.200.0.1 bypass-routing gateway 172.16.0.1** command.

Sample Output	<pre> user@c1> traceroute 220.220.0.0 source 200.200.0.1 bypass-routing gateway 172.16.0.1 traceroute to 220.220.0.1 (172.16.0.1) from 200.200.0.1, 30 hops max, 40 byte packets 1 172.16.0.1 (172.16.0.1) 0.661 ms 0.538 ms 0.449 ms 2 10.0.8.9 (10.0.8.9) 0.511 ms 0.479 ms 0.468 ms MPLS Label=100004 CoS=0 TTL=1 S=1 3 10.0.8.5 (10.0.8.5) 0.476 ms 0.512 ms 0.441 ms 4 220.220.0.1 (220.220.0.1) 0.436 ms 0.420 ms 0.416 ms </pre>
What It Means	<p>The output shows the route that traffic travels between C1 and C2, without using the default gateway. Verify that traffic sent from C1 to C2 travels through router R7. The 10.0.8.9 address is the interface address for router R5.</p>

Verifying an RSVP-Signaled LSP

Suppose that RSVP is configured to establish an LSP as shown in Figure 72.

To verify the RSVP configuration, perform these verification tasks:

- “Verifying RSVP Neighbors” on page 341
- “Verifying RSVP Sessions” on page 341
- “Verifying the Presence of RSVP-Signaled LSPs” on page 342

Verifying RSVP Neighbors

Purpose	Verify that each Services Router shows the appropriate RSVP neighbors—for example, that router R1 lists both router R3 and router R2 as RSVP neighbors.
Action	From the CLI, enter the <code>show rsvp neighbor</code> command.
Sample Output	<pre> user@r1> show rsvp neighbor RSVP neighbor: 2 learned Address Idle Up/Dn LastChange HelloInt HelloTx/Rx 10.0.6.2 0 3/2 13:01 3 366/349 10.0.3.3 0 1/0 22:49 3 448/448 </pre>
What It Means	The output shows the IP addresses of the neighboring routers. Verify that each neighboring RSVP router loopback address is listed.

Verifying RSVP Sessions

Purpose	Verify that an RSVP session has been established between all RSVP neighbors. Also, verify that the bandwidth reservation value is active.
----------------	---

Action From the CLI, enter the show rsvp session detail command.

Sample Output

```
user@r1> show rsvp session detail

Ingress RSVP: 1 sessions

10.0.9.7
  From: 10.0.6.1, LSPstate: Up, ActiveRoute: 0
  LSPname: r1-r7, LSPpath: Primary
  Bidirectional, Upstream label in: -, Upstream label out: -
  Suggested label received: -, Suggested label sent: -
  Recovery label received: -, Recovery label sent: 100000
  Resv style: 1 FF, Label in: -, Label out: 100000
  Time left: -, Since: Thu Jan 26 17:57:45 2002
  Tspec: rate 10Mbps size 10Mbps peak Infbps m 20 M 1500
  Port number: sender 3 receiver 17 protocol 0
  PATH rcvfrom: localclient
  PATH sentto: 10.0.4.13 (fe-0/0/1.0) 1467 pkts
  RESV rcvfrom: 10.0.4.13 (fe-0/0/1.0) 1467 pkts
  Record route: <self> 10.0.4.13 10.0.2.1 10.0.8.10
```

What It Means The output shows the detailed information, including session IDs, bandwidth reservation, and next-hop addresses, for each established RSVP session. Verify the following information:

- Each RSVP neighbor address has an entry for each neighbor, listed by loopback address.
- The state for each LSP session is Up.
- Under Tspec, the appropriate bandwidth value, 10Mbps, appears.

Verifying the Presence of RSVP-Signaled LSPs

Purpose Verify that the inet.3 routing table of the entry (ingress) Services Router, R1, has a configured LSP to the loopback address of router R7.

Action From the CLI, enter the show route table inet.3 command.

Sample Output

```
user@r1> show route table inet.3

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.9.7/32          *[RSVP/7] 00:05:29, metric 10
                    > to 10.0.4.17 via fe-0/0/0.0, label-switched-path r1-r7
```

What It Means The output shows the RSVP routes that exist in the inet.3 routing table. Verify that an RSVP-signaled LSP is associated with the loopback address of the exit (egress) router R7, in the MPLS network.

Chapter 14

Configuring Virtual Private Networks

You can configure a Services Router to participate in several types of virtual private networks (VPNs). A VPN allows remote sites and users to use a public communication infrastructure to create secure access to an organization's network. VPNs are a cost-effective alternative to expensive dedicated lines.

There are many ways to set up a VPN and direct traffic through it. This chapter describes the most common tasks involved in setting up a basic Layer 2 VPN, Layer 2 circuit, or Layer 3 VPN configuration. For more information about VPNs, including other configurations and advanced or less common tasks, see the *JUNOS VPNs Configuration Guide*.

You can use either the J-Web configuration editor or the CLI configuration editor to configure VPNs.

This chapter contains the following topics:

- VPN Configuration Overview on page 343
- Before You Begin on page 346
- Configuring VPNs with a Configuration Editor on page 346
- Verifying a VPN Configuration on page 364

VPN Configuration Overview

To configure VPN functionality on a Services Router, you must enable support on the provider edge (PE) Services Router as well as configure the Services Router to distribute routing information to other Services Routers in the VPN. The sample configurations in this chapter describe setting up a basic Multiprotocol Label Switching (MPLS) Layer 2 VPN, Layer 3 VPN, and Layer 2 circuit.

This section contains the following topics:


- Sample VPN Topology on page 344
- Basic Layer 2 VPN Configuration on page 344
- Basic Layer 2 Circuit Configuration on page 345

- Basic Layer 3 VPN Configuration on page 345

Sample VPN Topology

Figure 73 shows the overview of a basic VPN topology for the sample configurations in this chapter.

Figure 73: Basic VPN Topology



Basic Layer 2 VPN Configuration

Implementing a Layer 2 VPN on the Services Router is similar to implementing a VPN using a Layer 2 technology such as Asynchronous Transfer Mode (ATM) or Frame Relay. However, for a Layer 2 VPN on the Services Router, traffic is forwarded to the router in a Layer 2 format. Traffic is then carried by Multiprotocol Label Switching (MPLS) over the service provider's network, and then converted back to Layer 2 format at the receiving end.

On a Layer 2 VPN, routing occurs on the customer's Services Routers, typically on the customer edge (CE) router. The CE Services Router connected to a service provider on a Layer 2 VPN must select the appropriate circuit on which to send traffic. The provider edge (PE) Services Router receiving the traffic sends it across the service provider's network to the PE Services Router connected to the receiving site. PE Services Routers are not required to learn the customer's routes or routing topology, but they must identify the tunnel through which to send the data.

In this sample Layer 2 VPN configuration, the PE routers use the same autonomous system (AS). Within the AS, routing information is communicated through an interior gateway protocol (IGP). Outside the AS, routing information is shared with other ASs through Border Gateway Protocol (BGP). Each AS has a single routing policy and uses a group of one or more IP prefixes. The PE routers must use the same signaling protocols to communicate.

Each routing instance that you configure on a PE router must have a unique route distinguisher associated with it. VPN routing instances need a route distinguisher to help BGP identify overlapping network layer reachability information (NLRIs) messages from different VPNs.

Basic Layer 2 Circuit Configuration

A Layer 2 circuit is a point-to-point Layer 2 connection that transports traffic by means of Multiprotocol Label Switching (MPLS) or another tunneling technology on the service provider network. The Layer 2 circuit creates a virtual connection to direct traffic between two CE Services Routers across a service provider network. The main difference between a Layer 2 VPN and a Layer 2 circuit is the method of setting up the virtual connection. As with a leased line, a Layer 2 circuit forwards all packets received from the local interface to the remote interface.

On the interface communicating with the other PE router, you must specify MPLS and IPv4, and include the IP address. For the loopback interface, you must specify `inet`, and include the IP address. For IPv4, you must designate the loopback interface as primary so it can receive control packets. Because it is always operational, the loopback interface is best able to perform the control function.

On the PE router interface facing the CE router, you must specify a circuit cross-connect (CCC) encapsulation type. The type of encapsulation depends on the interface type. For example, an Ethernet interface uses `ethernet-ccc`. The encapsulation type determines how the packet is constructed for that interface.

On the CE router interface that faces the PE router, you must specify `inet` (for IPv4), and include the IP address. You also specify a routing protocol such as Open Shortest Path First (OSPF) which specifies the area and IP address of the Services Router interface.

With this information, the Services Routers can send and receive packets across the circuit.

Basic Layer 3 VPN Configuration

A Layer 3 VPN operates at the Layer 3 level of the OSI model, the Network layer. In this configuration, the service provider network must learn the IP addresses of devices sending traffic across the VPN. The Layer 3 VPN requires more processing power on the PE Services Routers, because it has larger routing tables for managing network traffic on the customer sites.

A Layer 3 VPN is a set of sites that share common routing information, and connectivity of the sites is controlled by a collection of policies. The sites making up a Layer 3 VPN are connected over a service provider's existing public Internet backbone.

An interface on each CE Services Router communicates with an interface on a PE Services Router through the external Border Gateway Protocol (EBGP).

On the provider Services Router, you configure two interfaces: one to communicate with each PE Services Router. The interfaces communicate with the PE Services Routers by using IPv4 and MPLS. The provider router is in the same AS as the PE routers, which is typically the case for Layer 3 VPNs.

The provider router uses OSPF and Label Distribution Protocol (LDP) to communicate with the PE Services Routers. For OSPF, the provider Services

Router interfaces that communicate with the PE routers are specified, as well as the loopback interface. For the PE routers, the loopback interface is in passive mode, meaning it does not send OSPF packets to perform the control function. In this example, the provider router and PE routers are in the same backbone area. For the LDP configuration, the provider router interfaces that communicate with the PE routers are specified.

Before You Begin

Before you begin configuring VPNs, perform the following tasks:

- Determine which Services Routers are participating in the VPN configuration. This chapter describes configuring an interface for basic VPN connectivity. To configure an interface, see “Configuring Network Interfaces” on page 103.
- Determine the protocols to use in the VPN configuration. These protocols include
 - MPLS—See “Multiprotocol Label Switching Overview” on page 315 and the *JUNOS Routing Protocols Configuration Guide*.
 - BGP, EBGp, and internal BGP (IBGP)—See “Configuring BGP Sessions” on page 295 and the *JUNOS Routing Protocols Configuration Guide*.
 - LDP and Resource Reservation Protocol (RSVP)—See “Configuring Signaling Protocols for Traffic Engineering” on page 331 and the *JUNOS MPLS Applications Configuration Guide*.
 - OSPF—See “Configuring an OSPF Network” on page 265 and the *JUNOS Routing Protocols Configuration Guide*.

Configuring VPNs with a Configuration Editor

To configure a basic Layer 3 VPN, Layer 2 VPN, or Layer 2 circuit, perform the following tasks. Use Table 86 to help you select the tasks for your VPN type. For information about using the J-Web and CLI configuration editors, see “Using J-Web Configuration Tools” on page 3.

- Configuring Interfaces Participating in a VPN on page 347
- Configuring Protocols Used by a VPN on page 349
- Configuring a VPN Routing Instance on page 357
- Configuring a VPN Routing Policy on page 359

Table 86: VPN Configuration Task Summary

Section	Layer 3 VPN	Layer 2 VPN	Layer 2 Circuit
“Configuring Interfaces Participating in a VPN” on page 347	All Services Routers	All Services Routers	All Services Routers
“Configuring Protocols Used by a VPN” on page 349	All Services Routers	All Services Routers	All Services Routers
“Configuring a VPN Routing Instance” on page 357	PE Services Routers	PE Services Routers	N/A
“Configuring a VPN Routing Policy” on page 359	CE Services Routers (PE Services Routers if you are not using a route target)	PE Services Routers if you are not using a route target	N/A

Configuring Interfaces Participating in a VPN

Configuring the Services Router interfaces that participate in the VPN is similar to configuring them for other uses, with a few requirements for VPN.

Before following the procedures in this section, make sure you have initially configured the interface as described in “Configuring Network Interfaces” on page 103.

To configure an interface for a VPN:

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 87 for each interface involved in the VPN, except Layer 3 loopback interfaces, which do not require other configuration.
3. Go on to “Configuring Protocols Used by a VPN” on page 349.

Table 87: Configuring an Interface for a VPN

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure IPv4. (interfaces on all Services Routers)	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Interfaces. 2. In the Interface name column, select the interface. 3. For Layer 2 VPNs on the interface facing a CE router, select an encapsulation type, such as ethernet-ccc from the Encapsulation list. For Fast Ethernet interfaces, you also must select Vlan tagging from the Vlan tag mode list. 4. In the Interface unit number column, select the logical interface. 5. In the Family group, select Inet and click Edit. 6. Next to Address, click Add new entry 7. In the Source box, type the IPv4 address—for example, 10.49.102.1/30. For a loopback address on a Layer 2 configuration, select Primary. 8. Click OK to return to the Unit page. 	<ul style="list-style-type: none"> ■ For all interfaces except loopback, and a Layer 2 VPN interface facing a CE router: From the top of the configuration hierarchy, enter <code>edit interfaces interface-name unit logical_interface family inet address ipv4_address</code> ■ For a loopback address on a Layer 2 configuration: From the top of the configuration hierarchy, enter <code>edit interfaces loO unit logical_interface family inet address ipv4_address primary</code> ■ For a Layer 2 VPN interface facing a CE router: From the top of the configuration hierarchy, enter <code>set interfaces interface-name vlan-tagging encapsulation vlan-ccc unit logical_interface encapsulation vlan-ccc vlan-id id-number</code>
Configure the MPLS address family. (for interfaces on a PE or provider Services Router that communicate with a PE or provider Services Router only, and not for loopback addresses)	On the Unit page, select Mpls in the Family group.	At the [edit interfaces <i>interface</i>] level, enter <code>set unit logical_interface family mpls</code>
For Layer 2 VPNs and circuits, configure encapsulation. If multiple logical units are configured, the encapsulation type is needed at the interface level only. It is always required at the unit level. (for interfaces on a PE Services Router that communicate with a CE Services Router)	<ol style="list-style-type: none"> 1. On the Unit page, select an encapsulation type from the Encapsulation list. 2. Click OK. 3. On the Interface page, select an encapsulation type from the Encapsulation list. 4. Click OK until you see the Configuration Interfaces page displaying all interfaces on the router. 	<ol style="list-style-type: none"> 1. At the [edit interfaces <i>interface</i>] level, enter <code>set encapsulation encapsulation_type</code> 2. Enter <code>set unit logical_interface encapsulation encapsulation_type</code>

Configuring Protocols Used by a VPN

The Services Routers in a VPN use a variety of protocols to communicate between PE and provider Services Routers. Use Table 88 to help you select the tasks for your VPN type. For more information about configuring routing protocols, see the *JUNOS Routing Protocols Configuration Guide* and the *JUNOS MPLS Applications Configuration Guide*.

This section contains the following topics:

- “Configuring MPLS for VPNs” on page 349
- “Configuring a BGP Session” on page 351
- “Configuring Routing Options for VPNs” on page 352
- “Configuring an IGP and a Signaling Protocol” on page 353
- “Configuring LDP for Signaling” on page 353
- “Configuring RSVP for Signaling” on page 355
- “Configuring a Layer 2 Circuit” on page 356

Table 88: VPN Protocol Configuration Task Summary

Section	Layer 3 VPN	Layer 2 VPN	Layer 2 Circuit
“Configuring MPLS for VPNs” on page 349	N/A unless you are using RSVP	PE and provider Services Routers	PE Services Routers
“Configuring a BGP Session” on page 351	PE Services Routers	PE Services Routers	PE Services Routers
“Configuring Routing Options for VPNs” on page 352	All Services Routers	All Services Routers	All Services Routers
“Configuring an IGP and a Signaling Protocol” on page 353— <i>one</i> of the following tasks: <ul style="list-style-type: none"> ■ “Configuring LDP for Signaling” on page 353 ■ “Configuring RSVP for Signaling” on page 355 	PE and provider Services Routers	PE Services Routers	PE Services Routers
“Configuring a Layer 2 Circuit” on page 356	N/A	N/A	PE Services Routers

Configuring MPLS for VPNs

For Layer 2 VPN and Layer 2 circuit interfaces that communicate with other PE Services Routers and provider Services Routers, you must advertise the

interface using MPLS. Unless you are using RSVP, this section does not apply to Layer 3 VPNs because MPLS is configured on the interface.

For more information about configuring MPLS, see “Multiprotocol Label Switching Overview” on page 315 and the *JUNOS MPLS Applications Configuration Guide*.

To configure MPLS for VPNs:

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 89 on each PE Services Router and provider Services Router interface that communicates with another PE Services Router.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 364.
5. Go on to “Configuring a BGP Session” on page 351.

Table 89: Configuring MPLS for VPNs

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the top of the configuration hierarchy and specify the interfaces used for communication between PE routers and between PE routers and provider routers. (PE and provider Services Routers)	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Protocols > Mpls > Interface. 2. In the Interface name box, type <i>interface-name</i>. 3. Click OK. 	<p>From the top of the configuration hierarchy, enter the following command for each interface you want to enable:</p> <pre>edit protocols mpls interface <i>interface-name</i></pre>
For RSVP only, configure an MPLS label-switched path (LSP) to the destination point on the PE router for LSP. During configuration, you specify the IP address of the LSP destination point, which is an address on the remote PE router. The path name is defined on the source Services Router only and is unique between two routers. (PE Services Router interface communicating with another PE Services Router)	<ol style="list-style-type: none"> 1. In the MPLS page, click Add New Entry in the Label switched path group. 2. Type a path name in the Path name box and an IP address in the To box. 3. Click OK. 4. Next to Interface, click Add New Entry. 5. Type <i>interface-name</i> in the Interface name box. 6. Click OK. 7. Repeat Steps 4 through 6 for each interface. 	<ol style="list-style-type: none"> 1. From the top of the configuration hierarchy, enter <pre>edit protocols mpls label-switched-path <i>path-name</i></pre> 2. Enter <pre>set to <i>ip-address</i></pre> 3. Enter <i>up</i>. 4. Enter <pre>interface <i>interface-name</i></pre>

Configuring a BGP Session

You must configure an internal BGP (IBGP) session between PE Services Routers so the Services Routers can exchange information about routes originating and terminating in the VPN. The PE routers use this information to determine which labels to use for traffic destined for remote sites. The IBGP session for the VPN runs through the loopback address. This section is valid for Layer 2 VPNs and Layer 3 VPNs, but not Layer 2 circuits.

For the Layer 3 example, you also configure an EBGp session.

For more information about configuring IBGP sessions, see “Configuring BGP Sessions” on page 295 and the *JUNOS Routing Protocols Configuration Guide*.

To configure an IBGP session:

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 90 on each PE router.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, “Verifying a VPN Configuration” on page 364.
5. Go on to “Configuring Routing Options for VPNs” on page 352.

Table 90: Configuring an IGBP Session

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the top of the configuration hierarchy and configure the IGBP session. (PE Services Router)	1. In the configuration editor hierarchy, select Protocols > Bgp .	1. From the top of the configuration hierarchy, enter edit protocols bgp group <i>group-name</i>
	2. Next to Group, click Add New Entry .	
	3. Type a name in the Group name box.	2. Enter
	4. From the Type list, select Internal .	set type internal
	5. In the Local address box, type the local loopback IP address.	3. Enter set local-address loopback-interface-ip-address
	6. In the Family group, select L2vpn for a Layer 2 VPN or Inet vpn for a Layer 3 VPN.	4. Enter
	7. Select Unicast .	set family <i>family-type</i> unicast
	8. Click OK .	Replace <i>family-type</i> with l2vpn for a Layer 2 VPN or inet-vpn for a Layer 3 VPN.
	9. In the Neighbor group, click Add new entry .	5. Enter up.
	10. In the Address box, type the loopback IP address of the neighboring PE router.	6. Enter the loopback address of the neighboring PE router:
	11. Click OK until you return to the BGP page.	set neighbor <i>ip-address</i>

Configuring Routing Options for VPNs

The only required routing option for VPNs is the autonomous system (AS) number. You must specify it on each router involved in the VPN.

To configure routing options for a VPN:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration task described in Table 91.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 364.
5. Go on to “Configuring an IGP and a Signaling Protocol” on page 353.

Table 91: Configuring Routing Options for a VPN

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the AS number.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, click Routing Options. 2. In the AS number box, type the AS number. 3. Click OK. 	<p>From the top of the configuration hierarchy, enter</p> <pre>set routing-options autonomous-system as-number</pre>

Configuring an IGP and a Signaling Protocol

The PE Services Routers and provider Services Routers must be able to exchange routing information. To enable this exchange, you must configure either an IGP such as OSPF or static routes on these routers. You must configure the IGP at the [edit protocols] level, not within the routing instance at the [edit routing-instances] level.

You can use LDP or RSVP between PE routers and between PE routers and provider routers, but not for interfaces between PE routers and CE routers. LDP routes traffic using IGP metrics. RSVP has traffic engineering that lets you override IGP metrics as needed. For more information about these protocols, see “Signaling Protocols Overview” on page 322.

Each PE Services Router’s loopback address must appear as a separate route. Do not configure any summarization of the PE Services Router’s loopback addresses at the area boundary.

For more information about configuring IGPs and static routes, see “Configuring an OSPF Network” on page 265, “Configuring Static Routes” on page 237, and the *JUNOS Routing Protocols Configuration Guide*.

Configure the appropriate signaling protocol for your VPN:

- “Configuring LDP for Signaling” on page 353
- “Configuring RSVP for Signaling” on page 355

Configuring LDP for Signaling

You must configure LDP and OSPF on PE and provider routers. For more information about configuring OSPF see “Configuring an OSPF Network” on page 265.

To configure LDP and OSPF:

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 92 on PE and provider router interfaces that communicate with a PE router or provider router.

For the protocols to work properly, you also must configure the MPLS address family for each interface that uses LDP or RSVP, as described previously in “Configuring Interfaces Participating in a VPN” on page 347.

3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 364.
5. Go on to “Configuring a VPN Routing Instance” on page 357.

Table 92: Configuring LDP and OSPF for Signaling

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Navigate to the top of the configuration hierarchy and specify the LDP protocol. Enable local interfaces that communicate with a PE router or provider router, and the loopback interface of the PE router.</p> <p>(PE and provider Services Routers)</p>	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Protocols > Ldp > Interface. 2. In the Interface name column, type <i>interface-name</i>. 3. Click OK. 4. Repeat Steps 2 and 3 for each interface you want to enable. 	<p>From the top of the configuration hierarchy, enter the following command for each interface you want to enable:</p> <pre>edit protocols ldp interface <i>interface-name</i></pre>
<p>Configure OSPF for each interface that uses LDP.</p> <p>For OSPF, you must configure at least one area on at least one of the router's interfaces. An AS can be divided into multiple areas. This example uses the backbone area 0.0.0.0.</p> <p>(PE and provider Services Routers)</p>	<p>For OSPF:</p> <ol style="list-style-type: none"> 1. In the configuration editor hierarchy, click Protocols > Ospf. 2. For Layer 2 VPN or circuit, select Traffic engineering. 3. Next to Area group, click Add new entry and add the area. 4. Next to Area group, select the area (0.0.0.0). 5. Next to Interface group, select Add new entry. 6. In the Interface name box, type <i>interface-name</i>. 7. Click OK. 8. Repeat Steps 5 through 7 to enable additional interfaces. 9. Click OK twice to return to the Protocols page. 	<p>For OSPF:</p> <ol style="list-style-type: none"> 1. From the top of the configuration hierarchy, enter the following command for each interface you want to enable: <pre>edit protocols ospf area 0.0.0.0 interface <i>interface-name</i></pre> 2. For Layer 2 VPN or circuit, move up to the [edit protocols ospf] level and enter <pre>set traffic-engineering</pre>

Configuring RSVP for Signaling

You must enable RSVP for all connections that participate in the label-switched path (LSP) on PE and provider Services Routers. In addition, you must configure OSPF on various interfaces.

For more information about configuring OSPF see “Configuring an OSPF Network” on page 265.

To configure RSVP and OSPF:

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 93 on each PE router and provider router, as specified.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 364.
5. Go on to “Configuring a VPN Routing Instance” on page 357.

Table 93: Configuring RSVP and OSPF for Signaling

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the top of the configuration hierarchy and configure OSPF with traffic engineering support. (PE Services Router)	For OSPF, follow these steps: <ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Protocols > Ospf. 2. Select Traffic engineering, and then click Configure. 3. Select Shortcuts. 4. Click OK until you return to the Protocols page. 	For OSPF, from the top of the configuration hierarchy, enter the following command for each interface you want to enable: edit protocols ospf traffic-engineering shortcuts
Enable RSVP on interfaces that participate in the LSP. (PE Services Router) Enable interfaces on the source and destination points. (provider Services Router) Enable interfaces that connect the LSP between the PE Services Routers.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Protocols > Rsvp. 2. In the Interface group, click Add New Entry. 3. Type an interface name. 4. Click OK. 5. Repeat Steps 2 through 4 for each interface you want to enable. 6. Click OK. 	From the top of the configuration hierarchy, enter the following command for each interface you want to enable: edit protocols rsvp interface <i>interface-name</i>

Configuring a Layer 2 Circuit

Each Layer 2 circuit is represented by the logical interface connecting the local PE Services Router to the local CE Services Router. All Layer 2 circuits using a particular remote PE Services Router neighbor is identified by its IP address and is usually the endpoint destination for the LSP tunnel transporting the Layer 2 circuit.

You configure a virtual circuit ID on each interface. Each virtual circuit ID uniquely identifies the Layer 2 circuit among all the Layer 2 circuits to a specific neighbor. The key to identifying a particular Layer 2 circuit on a PE router is the neighbor address and the virtual circuit ID. Based on the virtual circuit ID and the neighbor relationship, an LDP label is bound to an LDP circuit. LDP uses the binding for sending traffic on that Layer 2 circuit to the remote CE router.

To configure a Layer 2 circuit:

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 94 on each PE router and provider router, as specified.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 364.

Table 94: Configuring a Layer 2 Circuit

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the top of the configuration hierarchy and enable a Layer 2 circuit on the appropriate interface. (PE Services Router)	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Protocols > L2circuit. 2. Next to Neighbor, click Add new entry. 3. In the Neighbor box, enter the loopback address of the local router. 4. Next to Interface, click Add new entry. 5. In the Interface box, type the interface name of the remote PE router. 6. In the Virtual circuit id box, type an ID number. 7. Click OK until you return to the Protocols page. 	<ol style="list-style-type: none"> 1. From the top of the configuration hierarchy, enter <code>edit protocols l2circuit neighbor</code> <code>interface-name interface interface-name</code> For neighbor, specify the local loopback address, and for interface, specify the interface name of the remote PE router. 2. Enter <code>set virtual-circuit-id id-number</code>

Configuring a VPN Routing Instance

You must configure a routing instance for each VPN on each PE Services Router participating in the VPN. The routing instance has the same name on each PE router. VPN routing instances need a route distinguisher to help BGP distinguish between potentially identical network layer reachability information (NLRI) messages received from different VPNs. This section does not apply to Layer 2 circuit configurations.

Each routing instance that you configure on a PE router must have a unique route distinguisher. There are two possible formats:

- *as-number:number*, where *as-number* is an autonomous system (AS) number (a 2-byte value) in the range 1 through 65,535, and *number* is any 4-byte value. We recommend that you use an Internet Assigned Numbers Authority (IANA)-assigned, nonprivate AS number, preferably the ISP or the customer AS number.
- *ip-address:number*, where *ip-address* is an IP address (a 4-byte value) and *number* is any 2-byte value. The IP address can be any globally unique unicast address. We recommend that you use the address that you configure in the *router-id* statement, which is a public IP address in your assigned prefix range.

The route target defines which route is part of a VPN. A unique route target helps distinguish between different VPN services on the same router. Each VPN also has a policy that defines how routes are imported into the VPN routing and forwarding (VRF) table on the router. A Layer 2 VPN is configured with import and export policies. A Layer 3 VPN uses a unique route target to distinguish between VPN routes.

To configure a VPN routing instance:

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 95 on each PE router.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 364.
5. Go on to “Configuring a VPN Routing Policy” on page 359.

Table 95: Configuring a VPN Routing Instance

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the top of the configuration hierarchy and create the routing instance. (PE Services Router)	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Routing instances > Mpls. 2. In the Instance group, click Add New Entry. 3. Type a name in the Instance name box. 	<p>From the top of the configuration hierarchy, enter</p> <p>edit routing-instances <i>routing-instance-name</i></p>
Specify a text description for the routing instance. This text appears in the output of the show route instance detail command. (PE Services Router)	In the Description box, type a description.	<p>Enter</p> <p>set description " <i>text</i> "</p>
Specify the instance type, either l2vpn for Layer 2 VPNs or vrf for Layer 3 VPNs. (PE Services Router)	From the Instance type list, select an instance type.	<p>Enter</p> <p>set instance-type <i>instance-type</i></p>
Specify the interface of the remote PE Services Router. (PE Services Router)	<ol style="list-style-type: none"> 1. Next to Interface group, click Add New Entry. 2. In the Interface name box, enter <i>interface-name</i> . 3. Click OK. 	<p>Enter</p> <p>set interface <i>interface-name</i></p>
Specify the route distinguisher. (PE Services Router)	In the Rd type box, enter a route distinguisher in the format <i>as-number : number</i> or <i>ip-address : number</i> .	<p>Enter one of the following commands:</p> <ul style="list-style-type: none"> ■ set route-distinguisher <i>as-number : number</i> ■ set route-distinguisher <i>ip-address : number</i>

Table 95: Configuring a VPN Routing Instance (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Specify the policy for the Layer 2 VRF table.	For the sample Layer 2 VPN configuration, which uses import and export policies:	For the sample Layer 2 VPN configuration, which uses import and export policies, enter
For the Layer 2 VPN example, the routing policies are defined in “Configuring a Routing Policy for Layer 2 VPNs” on page 360. (PE Services Router)	<ol style="list-style-type: none"> Next to Vrf export group, select Add new entry. In the Value box, type the export routing policy name. Click OK. Next to Vrf import group, click Add new entry. In the Value box, type the import routing policy name. Click OK. 	<pre>set vrf-import <i>import-policy-name</i> vrf-export export-policy-name</pre>
Specify the policy for the Layer 3 VRF table.	For the sample Layer 3 VPN configuration, which uses a route target:	For the sample Layer 3 VPN configuration, which uses a route target, enter
For the Layer 3 VPN example, the routing policy is defined in “Configuring a Routing Policy for Layer 3 VPNs” on page 363. (PE Services Router)	<ol style="list-style-type: none"> In the Vrf target box, click Configure. In the Community box, type the community (target: <i>community-id</i>, where <i>community-id</i> is <i>as-number: number</i> or <i>ip-address: number</i>). Click OK. 	<pre>set vrf-target target: <i>community-id</i></pre> <p>Replace <i>community-id</i> with either of the following:</p> <ul style="list-style-type: none"> ■ <i>as-number: number</i> ■ <i>ip-address: number</i>

Configuring a VPN Routing Policy

Layer 2 and Layer 3 VPNs require a routing policy that describes which packets are sent and received across the VPN. Layer 2 circuits do not use a policy, and therefore, Layer 2 circuits send and receive all packets. For Layer 2 VPNs, the routing policy resides on the PE Services Routers. For the Layer 3 VPN example, the routing policy resides on the CE Services Routers.

This section contains the following topics. For more information about configuring routing policies, see “Configuring Routing Policies” on page 459 and the *JUNOS Routing Protocols Configuration Guide*.

- “Configuring a Routing Policy for Layer 2 VPNs” on page 360
- “Configuring a Routing Policy for Layer 3 VPNs” on page 363

Configuring a Routing Policy for Layer 2 VPNs

If the routing instance uses a policy for accepting and rejecting packets instead of a route target, you must specify the import and export routing policies and the community on each PE Services Router.

To configure a Layer 2 VPN routing policy on a PE 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 96 and Table 97 on each PE router.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 364.

Table 96: Configuring an Import Routing Policy for Layer 2 VPNs

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the top of the configuration hierarchy and configure the import routing policy. (PE Services Router)	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Policy options > Policy statement. 2. In the Policy name box, type the policy name—for example, <code>import_vpn</code>. 	<p>From the top of the configuration hierarchy, enter</p> <pre>edit policy-options policy-statement import-policy-name</pre>

Table 96: Configuring an Import Routing Policy for Layer 2 VPNs (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define the term for accepting packets. (PE Services Router)	<ol style="list-style-type: none"> Next to Term group, click Add new entry. In the Term name box, type a term name—for example, 10. Next to From, click Configure. Click Add new entry. Click Protocol and select bgp from the Value menu. Click OK. Next to Community, click Add new entry. Type the <i>community-name</i> value in the Community Name box. Click OK. Next to Then, click Configure. From the Accept reject list, select accept. Click OK until you are at the Policy statement page. 	<ol style="list-style-type: none"> Enter set term <i>term-name-accept</i> from protocol bgp community <i>community-name</i> Enter set term <i>term-name-accept</i> then accept
Define the term for rejecting packets. (PE Services Router)	<ol style="list-style-type: none"> Next to the Term group, click Add new entry. In the Term name box, type a term name—for example, 20. Next to Then, click Configure. From the Accept list, select reject. Click OK until you return to the Policy options page. 	Enter set term <i>term-name-reject</i> then reject

After configuring an import routing policy for a Layer 2 VPN, configure an export routing policy for the Layer 2 VPN. The export routing policy defines how routes are exported from the PE Services Router routing table. An export policy is applied to routes sent to other PE Services Routers in the VPN. The export policy must also evaluate all routes received over the routing protocol session with the CE Services Router. The export policy must also contain a second term for rejecting all other routes.

Table 97: Configuring an Export Routing Policy for Layer 2 VPNs

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the export routing policy. (PE Services Router)	<ol style="list-style-type: none"> Next to the Policy statement group, click Add new entry. In the Policy name box, type the policy name—for example, <code>export_vpn</code>. 	<p>From the top of the configuration hierarchy, enter</p> <pre>edit policy-options policy-statement export-policy-name</pre>
Define the term for accepting packets. (PE Services Router)	<ol style="list-style-type: none"> Next to the Term group, click Add new entry. In the Term name box, type a term name—for example, <code>10</code>. Next to From, click Configure. Next to Community, click Add new entry. Type the <code>community-name</code> value in the Community Name box. Click OK. Next to Then, click Configure. From the Accept reject list, select accept. Click OK twice until you are at the Policy statement page. 	<ol style="list-style-type: none"> Enter <pre>set term term-name-accept from community add community-name</pre> Enter <pre>set term term-name-accept then accept</pre>
Define the term for rejecting packets. (PE Services Router)	<ol style="list-style-type: none"> Next to the Term group, click Add new entry. In the Term name box, type a term name—for example, <code>20</code>. Next to Then, click Configure. From the Accept reject list, select reject. Click OK until you return to the Policy options page. 	<ol style="list-style-type: none"> Enter <pre>set term term-name-reject from community add community-name</pre> Enter <pre>set term term-name-reject then reject</pre>
Define the community. (PE Services Router)	<ol style="list-style-type: none"> In the Community group, click Add new entry. In the Community name box, type a community name—for example, <code>VPN</code>. In the Members group, click Add new entry. In the Value box, type <code>target: community-id</code>, where <code>community-id</code> is <code>as-number : number</code> or <code>ip-address : number</code>. Click OK until you return to the Policy options page. 	<p>Type the following commands:</p> <pre>community community-name target: as-number or ip-address : number</pre>

Configuring a Routing Policy for Layer 3 VPNs

To configure a Layer 3 VPN routing policy on a CE 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 98 on each CE Services Router.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 364.

Table 98: Configuring a Routing Policy for Layer 3 VPNs

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the top of the configuration hierarchy and configure the routing policy for the loopback interface. (CE Services Router)	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Policy options > Policy statement. 2. In the Policy name box, type the policy name—for example, <code>loopback</code>. 	<p>From the top of the configuration hierarchy, enter</p> <pre>edit policy-options policy-statement policy-name</pre>

Table 98: Configuring a Routing Policy for Layer 3 VPNs (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define the term for accepting packets. (CE Services Router)	<ol style="list-style-type: none"> 1. In the Term group, click Add new entry. 2. In the Term name box, type a term name—for example, 1. 3. Next to From, click Configure. 4. Click protocol, then Add new entry. 5. Select direct from the Value menu, and click OK. 6. 7. Next to Route Filter, click Add new entry. 8. Type <i>local-loopback-address/netmask</i> in the Address box. 9. Select exact from the Modifier list. 10. Click OK twice. 11. Next to Then, click Configure. 12. From the Accept reject list, select accept. 13. Click OK until you are at the Policy statement page. 	<ol style="list-style-type: none"> 1. Enter <code>set term term-name-accept</code> <code>from protocol direct route-filter</code> <code>local-loopback-address/netmask exact</code> 2. Enter <code>set term term-name-accept</code> then <code>accept</code>
Define the term for rejecting packets. (CE Services Router)	<ol style="list-style-type: none"> 1. Next to the Term group, click Add new entry. 2. In the Term name box, type a term name—for example, 2. 3. Next to Then, click Configure. 4. From the Accept reject list, select reject. 5. Click OK until you return to the Policy options page. 	<ol style="list-style-type: none"> Enter <code>set term term-name-reject</code> then <code>reject</code>

Verifying a VPN Configuration

To verify the connectivity of Layer 2 VPNs, Layer 3 VPNs, and Layer 2 circuits, use the `ping mpls` command. This command helps to verify that a VPN or circuit has been enabled. This command tests the integrity of the VPN or Layer 2 circuit connection between the PE Services Routers. It does not test the connection between a PE and a CE Services Router.

This section contains the following topics:

- Pinging a Layer 2 VPN on page 365

- Pinging a Layer 3 VPN on page 365
- Pinging a Layer 2 Circuit on page 365

Pinging a Layer 2 VPN

To ping a Layer 2 VPN, use one of the following commands:

- `ping mpls l2vpn interface interface-name`

Ping an interface configured for the Layer 2 VPN on the PE router.

- `ping mpls l2vpn instance l2vpn-instance-name local-site-id local-site-id-number remote-site-id remote-site-id-number`

Ping a combination of the Layer 2 VPN routing instance name, the local site identifier, and the remote site identifier to test the integrity of the Layer 2 VPN connection (specified by identifiers) between the two PE Services routers.

Pinging a Layer 3 VPN

To ping a Layer 3 VPN, use the following command:

```
ping mpls l3vpn l3vpn-name prefix prefix <count count>
```

Ping a combination of a IPv4 destination prefix and a Layer 3 VPN name on the destination PE Services Router to test the integrity of the VPN connection between the source and destination Services Routers. The destination prefix corresponds to a prefix in the Layer 3 VPN. However, ping tests only whether the prefix is present in a PE VRF table.

Pinging a Layer 2 Circuit

To ping a Layer 2 circuit, use one of the following commands:

- `ping mpls l2circuit interface interface-name`

Ping an interface configured for the Layer 2 circuit on the PE Services Router.

- `ping mpls l2circuit virtual-circuit <prefix> <virtual-circuit-id>`

Ping a combination of the IPv4 prefix and the virtual circuit ID on the destination PE Services Router to test the integrity of the Layer 2 circuit between the source and destination Services Routers.

Chapter 15

Configuring CLNS VPNs

Connectionless Network Service (CLNS) is a Layer 3 protocol similar to IPv4 for linking hosts (end systems) with routers (intermediate systems) in an Open Systems Interconnection (OSI) network. CLNS and its related OSI protocols, Intermediate System-to-Intermediate System (IS-IS) and End System-to-Intermediate System (ES-IS), are International Organization for Standardization (ISO) standards.

You can configure Services Routers as provider edge (PE) routers within a CLNS network. CLNS networks can be connected over an IP MPLS network core using BGP and MPLS Layer 3 virtual private networks (VPNs). For more information, see RFC 2547, *BGP/MPLS VPNs*.

You can use either the J-Web configuration editor or CLI configuration editor to configure CLNS.

This chapter contains the following topics. For more information about CLNS, IS-IS, and ES-IS, see the *JUNOS Routing Protocols Configuration Guide*.

- CLNS Terms on page 367
- CLNS Overview on page 368
- Before You Begin on page 369
- Configuring CLNS with a Configuration Editor on page 369
- Verifying CLNS VPN Configuration on page 376

CLNS Terms

Before configuring CLNS, become familiar with the terms defined in Table 99.

Table 99: CLNS Terms

Term	Definition
CLNS island	Typically one IS-IS level 1 area that is part of a single IGP routing domain. An island can contain more than one area. CLNS islands can be connected by virtual private networks (VPNs).
Connectionless Network Service (CLNS)	Layer 3 protocol similar to IPv4 for linking hosts (end systems) with routers (intermediate systems) in an Open Systems Interconnection (OSI) network, by using network service access points (NSAPs) instead of prefix addresses to specify hosts and routers.

Table 99: CLNS Terms (continued)

Term	Definition
customer edge (CE) router	Router or switch in the customer's network that is connected to a service provider's provider edge (PE) router and participates in a Layer 3 VPN.
end system	A host in an Open Systems Interconnection (OSI) network.
End System-to-Intermediate System (ES-IS)	Protocol that enables end systems (hosts) and intermediate systems (routers) to discover each other, by a method similar to Address Resolution Protocol (ARP) discovery in an IPv4 network.
intermediate system	A router in an Open Systems Interconnection (OSI) network.
International Organization for Standardization (ISO)	Worldwide federation of standards bodies that promotes international standardization and published international agreements as International Standards.
network layer reachability information (NLRI)	Information about routes exchanged in update messages by Border Gateway Protocol (BGP) systems, to enable routers to determine the relationships among all known BGP autonomous systems.
network services access point (NSAP)	International Standards Organization (ISO) addressing method for identifying hosts (end systems) and routers (intermediate systems) at the data-link layer (Layer 3) in an Open Systems Interconnection (OSI) network. An NSAP is from 8 to 20 bytes long and consists of an area address, a system ID, and a network selector (NSEL) byte.
Open Systems Interconnection (OSI)	Standard reference model for representing the way messages are transmitted between two points on a network.
provider edge (PE) router	Services Router in the service provider network that is connected to a customer edge (CE) device and participates in a virtual private network (VPN).
virtual private network (VPN)	Private data network that uses a public TCP/IP network, typically the Internet, while maintaining privacy with a tunneling protocol, encryption, and security procedures.

CLNS Overview

CLNS uses network service access points (NSAPs), similar to IP addresses found in IPv4, to identify end systems (hosts) and intermediate systems (routers). ES-IS enables the hosts and routers to discover each other. IS-IS is the interior gateway protocol (IGP) that carries ISO CLNS routes through a network.

Depending on your network topology, one or more of the following components are needed to route within a CLNS environment:

- ES-IS—Provides the basic interaction between CLNS hosts (end systems) and routers (intermediate systems). Using ES-IS, hosts advertise their ISO NSAP addresses and subnetwork point-of-attachment (SNPA) addresses to other routers and hosts attached to the subnetwork. The resolution of Layer 3 ISO NSAPs to Layer 2 SNPAs by ES-IS is equivalent to ARP within an IPv4 network.

If a CLNS island does not contain any end systems, you do not need to configure ES-IS on a Services Router.

- IS-IS extensions—Provide the basic IGP support for collecting intradomain routing information for CLNS destinations within a CLNS network. Routers learning host addresses through ES-IS can advertise them to other routers (intermediate systems) using IS-IS.
- Static routes—You can configure static routes to exchange CLNS routes within a CLNS island. You can use static routing with or without IS-IS.
- Border Gateway Protocol (BGP) extensions—BGP extensions allow BGP to carry CLNS VPN network layer reachability information (NLRI) between PE routers. Each CLNS route is encapsulated into a CLNS VPN NLRI and propagated between remote sites in a VPN.

For more information about CLNS, see the ISO 8473 standards. For more information about IS-IS, see the ISO 10589 standard. For more information about ES-IS, see the ISO 9542 standard.

Before You Begin

Before you begin configuring CLNS, complete the following tasks:

- Configure IS-IS. See the *JUNOS Routing Protocols Configuration Guide*.
- Configure the network interfaces. See “Configuring Network Interfaces” on page 103.
- If applicable, configure BGP and VPNs. See “Configuring BGP Sessions” on page 295 and “Configuring Virtual Private Networks” on page 343.

Configuring CLNS with a Configuration Editor

To configure CLNS on a Services Router, you must perform the first task and then one or more of the following tasks (depending on your network):

- Configuring a VPN Routing Instance (Required) on page 370
- Configuring ES-IS on page 371
- Configuring IS-IS for CLNS on page 372
- Configuring CLNS Static Routes on page 374
- Configuring BGP for CLNS on page 375



NOTE: Many of the configuration statements used in this section can be included at different hierarchy levels in the configuration. For more information, see the *JUNOS Routing Protocols Configuration Guide*.

Configuring a VPN Routing Instance (Required)

You typically configure ES-IS, IS-IS, and CLNS static routes using a VPN routing instance. For more information about routing instances, see “Configuring a VPN Routing Instance” on page 357.

To configure a VPN routing instance:

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 100.
3. Go on to one of the following tasks:
 - “Configuring IS-IS for CLNS” on page 372
 - “Configuring CLNS Static Routes” on page 374
 - “Configuring BGP for CLNS” on page 375
 - “Verifying CLNS VPN Configuration” on page 376

Table 100: Configuring a VPN Routing Instance for CLNS

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the top of the configuration hierarchy and create the routing instance aaaa .	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Routing instances. 2. Next to Instance, click Add new entry. 3. In the Instance name box, type aaaa. 4. Click OK. 	From the top of the configuration hierarchy, enter edit routing-instances aaaa
Specify the instance type vrf for Layer 3 VPNs.	In the Instance type list, select vrf .	Enter set instance-type vrf

Table 100: Configuring a VPN Routing Instance for CLNS (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Specify the interfaces that belong to the routing instance aaaa —for example, lo0.1 , e1-2/0/0.0 , and t1-3/0/0.0 .	<ol style="list-style-type: none"> 1. Next to Interface, click Add New Entry. 2. In the Interface name box, type lo0.1. 3. Click OK. 4. Next to Interface, click Add New Entry. 5. In the Interface name box, type e1-2/0/0.0. 6. Click OK. 7. Next to Interface, click Add New Entry. 8. In the Interface name box, type t1-3/0/0.0. 9. Click OK. 	<p>Enter</p> <ol style="list-style-type: none"> 1. <code>set interface lo0.1</code> 2. <code>set interface e1-2/0/0.0</code> 3. <code>set interface t1-3/0/0.0</code>
Specify the route distinguisher—for example, 10.255.245.1:1 .	In the Rd type box, type 10.255.245.1:1 .	<p>Enter</p> <p><code>set route-distinguisher 10.255.245.1:1</code></p>
Specify the policy for the Layer 3 VRF table—for example, target:11111:1 .	<ol style="list-style-type: none"> 1. Next to Vrf target, click Configure. 2. In the Community box, type target:11111:1. 3. Click OK. 	<p>Enter</p> <p><code>set vrf-target target:11111:1</code></p>

Configuring ES-IS

If a Services Router is a PE router within a CLNS island that contains any end systems, you must configure ES-IS on the Services Router.

To configure ES-IS 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 101.
3. If you are finished configuring the router, commit the configuration.
4. If applicable, go on to one of the following tasks:
 - “Configuring IS-IS for CLNS” on page 372

- “Configuring CLNS Static Routes” on page 374
- “Configuring BGP for CLNS” on page 375
- “Verifying CLNS VPN Configuration” on page 376

Table 101: Configuring ES-IS

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Routing instances level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Routing instances. 2. Under Instance name, click aaaa. 	<p>From the top of the configuration hierarchy, enter</p> <pre>edit routing-instances aaaa</pre>
Enable ES-IS on all interfaces.	<ol style="list-style-type: none"> 1. Next to Protocols, click Configure. 2. Next to Esis, click Configure. 3. Next to Interface, click Add new entry. 4. In the Interface name box, type all. 5. Click OK until you return to the Protocols statement page. 	<pre>Enter set protocols esis interface all</pre>

Configuring IS-IS for CLNS

You can configure IS-IS to exchange CLNS routes within a CLNS island. To export BGP routes into IS-IS, you must configure and apply an export policy. For more information about policies, see “Configuring Routing Policies” on page 459.

If you have a pure CLNS island—an island that does not contain any IP devices—you must disable IPv4 and IPv6 routing.

To configure IS-IS for CLNS:

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 102.
3. If you are finished configuring the router, commit the configuration.
4. If applicable, go on to one of the following tasks:
 - “Configuring CLNS Static Routes” on page 374
 - “Configuring BGP for CLNS” on page 375
 - “Verifying CLNS VPN Configuration” on page 376

Table 102: Configuring IS-IS to Exchange CLNS Routes

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Routing instances level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Routing instances. 2. Under Instance name, click aaaa. 	<p>From the top of the configuration hierarchy, enter</p> <pre>edit routing-instances aaaa</pre>
Enable CLNS routing.	<ol style="list-style-type: none"> 1. Next to Protocols, click Configure. 2. Next to Isis, click Configure. 3. Next to CLNS routing, select the Yes box. 	<p>Enter</p> <pre>set protocols isis clns-routing</pre>
Enable IS-IS on all interfaces.	<ol style="list-style-type: none"> 1. Next to Interface, click Add new entry. 2. In the Interface name box, type all. 3. Click OK. 	<p>Enter</p> <pre>set protocols isis interface all</pre>
(Optional) To configure a pure CLNS network, disable IPv4 and IPv6 routing.	<ol style="list-style-type: none"> 1. Next to No ipv4 routing, select the Yes box. 2. Next to No ipv6 routing, select the Yes box. 3. Click OK. 	<p>Enter</p> <pre>set protocols isis no-ipv4-routing no-ipv6-routing</pre>
Define the BGP export policy name—for example, dist-bgp —and the family and protocol.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Policy options. 2. Next to Policy statement, click Add new entry. 3. In the Policy name box, type dist-bgp. 4. Next to From, click Configure. 5. In the Family list, select iso. 6. Next to Protocol, click Add new entry. 7. In the Value list, select bgp. 8. Click OK until you return to the Policy statement page. 	<p>From the top of the configuration hierarchy, enter</p> <pre>set policy-options policy-statement dist-bgp from family iso protocol bgp</pre>

Table 102: Configuring IS-IS to Exchange CLNS Routes (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define the action for the export policy.	<ol style="list-style-type: none"> 1. Next to Then, click Configure. 2. In the Accept reject list, select accept. 3. Click OK until you return to the Configuration page. 	<p>From the top of the configuration hierarchy, enter</p> <pre>set policy-options policy-statement dist-bgp then accept</pre>
Apply the export policy to IS-IS.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Routing instances. 2. Next to aaaa, click Protocols. 3. Next to Isis, click Edit. 4. Next to Export, click Add new entry. 5. In the Value box, type dist-bgp. 6. Click OK until you return to the Instance page. 	<p>From the top of the configuration hierarchy, enter</p> <pre>set routing-instances aaaa protocols isis export dist-bgp</pre>

Configuring CLNS Static Routes

If some devices in your network do not support IS-IS, you must configure CLNS static routes. You might also consider using static routes if your network is simple.

This procedure, as well as the configuration provided in “Verifying CLNS VPN Configuration” on page 376, uses the following ISO NET address and NSAP prefix:

- 47.0005.80ff.f800.0000.aaaa.1000.1921.6800.4196.00
- 47.0005.80ff.f800.0000.bbbb.1022/104

To configure CLNS static routes:

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.
3. If you are finished configuring the router, commit the configuration.
4. If applicable, go on to one of the following tasks:
 - “Configuring BGP for CLNS” on page 375
 - “Verifying CLNS VPN Configuration” on page 376

Table 103: Configuring Static CLNS Routes

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Routing instances level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Routing instances. 2. Under Instance name, click aaaa. 	<p>From the top of the configuration hierarchy, enter</p> <pre>edit routing-instances aaaa</pre>
Configure the next-hop ISO NET address for an NSAP prefix.	<ol style="list-style-type: none"> 1. Next to Routing options, click Configure. 2. Next to Rib, click Add new entry. 3. In the Rib name box, type aaaa.iso.0. 4. Next to Static, click Configure. 5. Next to Iso route, click Add new entry. 6. In the Destination box, type 47.0005.80ff.f800.0000.bbbb.1022/104. 7. From the Next hop list, select Next hop. 8. Next to Next hop, click Add new entry. 9. In the Value box, type 47.0005.80ff.f800.0000.aaaa.1000.1921.6800.4196.00. 10. Click OK. 	<p>Enter</p> <pre>set routing-options iso-route 47.0005.80ff.f800.0000.bbbb.1022/104 next-hop 47.0005.80ff.f800.0000.aaaa.1000.1921.6800.4196.00</pre>

Configuring BGP for CLNS

To configure BGP to carry CLNS VPN NLRI:

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 104.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying CLNS VPN Configuration” on page 376.

Table 104: Configuring BGP to Carry CLNS VPN NLRI Messages

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Bgp level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Bgp .	From the top of the configuration hierarchy, enter
Define a BGP group name—for example, pedge-pedge.	<ol style="list-style-type: none"> Next to Group, click Add new entry. In the Group name box, type pedge-pedge. 	<pre>set protocols bgp group pedge-pedge neighbor 10.255.245.215 family iso-vpn unicast</pre>
Define a BGP peer neighbor address for the group—for example, 10.255.245.215.	<ol style="list-style-type: none"> Next to Neighbor, click Add new entry. In the Address box, type 10.255.245.215. 	
Define the family.	<ol style="list-style-type: none"> Under Family, next to Iso vpn, click Configure. Next to Unicast, select the Yes box. Click OK. 	

Verifying CLNS VPN Configuration

Verify that the Services Router is configured correctly for CLNS VPNs.

Displaying CLNS VPN Configuration

- Purpose** Verify the configuration of CLNS VPNs.
- Action** From the J-Web interface, select **Configuration > View and Edit > View Configuration Text**. Alternatively, from configuration mode in the CLI, enter the show command.

Sample Output

```
[edit]
user@host# show
interfaces {
  e1-2/0/0.0 {
    unit 0 {
      family inet {
        address 192.168.37.51/31;
      }
      family iso;
      family mpls;
    }
  }
  t1-3/0/0.0 {
    unit 0 {
      family inet {
        address 192.168.37.24/32;
```

```

    }
    family iso;
    family mpls;
  }
}
lo0 {
  unit 0 {
    family inet {
      address 127.0.0.1/32;
      address 10.255.245.215/32;
    }
    family iso {
      address 47.0005.80ff.f800.0000.0108.0001.1921.6800.4215.00;
    }
  }
  unit 1 {
    family iso {
      address 47.0005.80ff.f800.0000.0108.aaa2.1921.6800.4215.00;
    }
  }
}
routing-options {
  autonomous-system 230;
}
protocols {
  bgp {
    group pedge-pegde {
      type internal;
      local-address 10.255.245.215;
      neighbor 10.255.245.212 {
        family iso-vpn {
          unicast;
        }
      }
    }
  }
}
policy-options {
  policy-statement dist-bgp {
    from {
      protocol bgp;
      family iso;
    }
    then accept;
  }
}
routing-instances {
  aaaa {
    instance-type vrf;
    interface lo0.1;
    interface e1-2/0/0.0;
    interface t1-3/0/0.0;
    route-distinguisher 10.255.245.1:1;
  }
}

```

```

vrf-target target:11111:1;
routing-options {
  rib aaaa.iso.0 {
    static {
      iso-route 47.0005.80ff.f800.0000.bbbb.1022/104
      next-hop 47.0005.80ff.f800.0000.aaaa.1000.1921.6800.4196.00;
    }
  }
}
protocols {
  esis {
    interface all;
  }
  isis {
    export dist-bgp;
    no-ipv4-routing;
    no-ip64-routing;
    clns-routing;
    interface all;
  }
}
}

```

What It Means Verify that the output shows the intended configuration of CLNS VPNs. For more information about the format of a configuration file, see “Viewing the Configuration Text” on page 12.

Chapter 16

Configuring IPSec for Secure Packet Exchange

An IP Security (IPSec) tunnel allows access to a private network through a secure tunnel. This feature is particularly useful when a private network is divided among multiple sites, and transit between the sites must occur on a public network. To ensure secure transport of packets across the public network to the multiple sites, individual tunnels are configured. Network Address Translation (NAT) enables packets outbound through a tunnel to be filtered by source address.



NOTE: You must have a license to configure an IPSec tunnel. For license details, see the *J-series Services Router Administration Guide*.

This chapter contains the following topics. For more information about IPSec and NAT, see the *JUNOS System Basics Configuration Guide* and the *JUNOS Services Interfaces Configuration Guide*.

- IPSec Tunnel Overview on page 379
- Before You Begin on page 380
- Configuring an IPSec Tunnel with Quick Configuration on page 380
- Configuring an IPSec Tunnel with a Configuration Editor on page 382
- Verifying the IPSec Tunnel Configuration on page 391

IPSec Tunnel Overview

Each IPSec tunnel is defined by a local tunnel endpoint and a remote tunnel endpoint. Packets with a destination address matching the private network prefix are encrypted and encapsulated in a tunnel packet that is routable through the outside network. The source address of the tunnel packet is the local gateway, and the destination address is the remote gateway. Once the encapsulation packet reaches the other side, the remote end determines how to route the packet.

Security Associations

An IPsec security association (SA) is a set of rules used by IPsec tunnel gateways by which traffic is transported. IPsec security associations are established either manually, through configuration statements, or by Internet Key Exchange (IKE). In the case of manually configured security associations, the connection is established when both ends of the tunnel are configured, and the connections last until one of the endpoints is taken offline. For IKE security associations, connections are established only when traffic is sent through the tunnel, and they dissolve after a preset amount of time or traffic.

Translating Outgoing Traffic

Outgoing (egress) traffic across the tunnel must be marked with the outbound tunnel endpoint address so that it can be filtered by the stateful firewall filter on the opposite side of the tunnel. Packet tagging is performed by Network Address Translation (NAT). The source address for outbound packets is translated to the local gateway address so that, to the remote gateway, all packets appear to originate from the local endpoint. Address translation enables the remote gateway to filter packets based on source address to determine which packets are to be transported through the tunnel.

Before You Begin

Before you begin configuring an IPsec tunnel, you must have completed these tasks:

- Establish basic connectivity. See the *J-series Services Router Getting Started Guide*.
- Configure network interfaces. See “Configuring Network Interfaces” on page 103.
- Configure one or more routing protocols. See “Configuring Static Routes” on page 237, “Configuring a RIP Network” on page 249, “Configuring an OSPF Network” on page 265, or “Configuring BGP Sessions” on page 295.

Configuring an IPsec Tunnel with Quick Configuration

J-Web Quick Configuration allows you to create IPsec tunnels. Figure 74 shows the Quick Configuration page for IPsec tunnels.

Figure 74: Quick Configuration Page for IPSec Tunnels

Juniper NETWORKS **ROUTER - J4300** Logged in as: **regress**
[Help](#) [About](#) [Logout](#)

Monitor Configuration Diagnose Manage

Quick Configuration

Set Up
 SSL
 Interfaces
 Users
 SNMP
 Routing
 Firewall/NAT
IPSec Tunnels
 Realtime Performance Monitoring

► **View and Edit**
 ► **History**
 ► **Rescue**

Configuration > Quick Configuration > IPSec Tunnels

Quick Configuration

IPSec Tunnels Add an IPSec Tunnel

Tunnel Information

* **Local Tunnel Endpoint** ?

* **Remote Tunnel Endpoint** ?

* **IKE Secret Key** ?

* **Verify IKE Secret Key**

Private Prefix List ?

Private Prefix List
<input type="text"/>

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. [Trademark Notice](#).

To configure an IPSec tunnel with Quick Configuration:

1. In the J-Web user interface, select **Configuration > IPSec Tunnels**.
2. Enter information into the Quick Configuration page for IPSec Tunnels, as described in Table 105.
3. From the IPSec Tunnels Quick Configuration page, click one of the following buttons:
 - To apply the configuration and return to the Quick Configuration IPSec Tunnels page, click **OK**.
 - To cancel your entries and return to the Quick Configuration for IPSec Tunnels page, click **Cancel**.

4. To check the configuration, see “Verifying the IPSec Tunnel Configuration” on page 391.

Table 105: IPSec Tunnels Quick Configuration Summary

Field	Function	Your Action
Tunnel Information		
Local Tunnel Endpoint (required)	Externally routable IP address that is the local endpoint of the IPSec tunnel	Type the IPSec tunnel's local endpoint 32-bit IP address, in dotted decimal notation.
Remote Tunnel Endpoint (required)	Externally routable IP address that is the peer endpoint of the IPSec tunnel	Type the IPSec tunnel's peer endpoint 32-bit IP address, in dotted decimal notation.
IKE Secret Key (required)	Internet Key Exchange key that is preshared to ensure authentication across the IPSec tunnel	Type the IKE key to be used for authentication across the IPSec tunnel. Characters are disguised as you type.
Verify IKE Secret Key (required)	Internet Key Exchange key that is preshared to ensure authentication across the IPSec tunnel	Verify the IKE key by retyping the key to be used for authentication across the IPSec tunnel. Characters are disguised as you type.
Private Prefix List	List of addresses or address prefixes for which the IPSec tunnel is used. Packets whose destination address matches any of the addresses or prefixes in this list are transported through the IPSec tunnel to the remote tunnel endpoint.	<ol style="list-style-type: none"> 1. In the text box at the bottom of the list, enter an IP address or address prefix, in dotted decimal notation. 2. Click Add.

Configuring an IPSec Tunnel with a Configuration Editor

To configure a Services Router to transport traffic across a secure IPSec tunnel, you must define the tunnel and configure its components. To configure an IPSec tunnel, perform the following tasks:

- Configuring IPSec Services Interfaces on page 383
- Configuring IPSec Service Sets on page 384
- Configuring an IPSec Stateful Firewall Filter Rule on page 387
- Configuring a NAT Pool on page 389

Configuring IPSec Services Interfaces

To configure an IPSec tunnel, you must configure the following services interfaces:

- *Inside services interface* —Logical interface used to apply the service sets that define the behavior of the IPSec tunnel for outbound traffic (traffic whose next hop is inside the IPSec tunnel).
- *Outside services interface* —Logical interface used to apply the service sets that define the behavior of the IPSec tunnel for inbound traffic (traffic whose next hop is outside the IPSec tunnel).

For the services to be applied, you must first define the logical interfaces to be used.

To configure IPSec inside services interfaces and outside services 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 106.
3. Go on to “Configuring IPSec Service Sets” on page 384.

Table 106: Configuring IPSec Interfaces

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	In the configuration editor hierarchy, select Interfaces .	From the top of the configuration hierarchy, enter edit interfaces

Table 106: Configuring IPSec Interfaces (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the inside services interface for the IPSec tunnel.	1. Next to Interface, click Add new entry .	1. Configure the services interface as an inside-service interface:
On the J-series Services Router, the services interface is always sp-0/0/0.unit . The logical interface must have a unit number other than 0. By default, J-Web Quick Configuration uses the unit number 1001 for inside-service logical interfaces.	2. In the Interface name box, type sp-0/0/0 , and click OK .	set sp-0/0/0 unit 1001 service-domain inside
	3. In the Interface box, click sp-0/0/0 .	2. Configure the services interface as an inet interface:
	4. In the Unit box, click Add new entry .	set sp-0/0/0 unit 1001 family inet
	5. In the Interface unit number box, type 1001 .	
	6. In the Service domain box, select inside from the list.	
	7. In the Family box, click inet .	
	8. Select the Primary box, and click OK .	
Configure the outside services interface for the IPSec tunnel.	1. Next to Interface, click Add new entry .	1. Configure the services interface as an outside-service interface:
On the J-series Services Router, the services interface is always sp-0/0/0.unit . The logical interface must have a unit number other than 0. By default, J-Web Quick Configuration uses the unit number 2001 for outside-service logical interfaces.	2. In the Interface name box, type sp-0/0/0 , and click OK .	set sp-0/0/0 unit 2001 service-domain outside
	3. Next to Interface, click sp-0/0/0 .	2. Configure the services interface as an inet interface:
	4. Next to Unit, click Add new entry .	set sp-0/0/0 unit 2001 family inet
	5. In the Interface unit number box, type 2001 .	
	6. In the Service domain box, select outside from the list.	
	7. In the Family box, click inet .	
	8. Select the Primary check box, and click OK .	

Configuring IPSec Service Sets

The next-hop service set defines which services interface to use for all inside-service next hops and all outside-service next hops (traffic inside the network and outside the network). The unit numbers used to define the next-hop interfaces must match exactly the unit numbers used in the interfaces configuration.

When you configure an IPSec service set, you must also configure the local gateway. You then configure an IPSec rule to set the remote gateway on all traffic, configure a security association (SA) with a static IKE key, and configure another rule to act

on input traffic. This configuration allows you to set the remote gateway address and perform IKE validation on all incoming traffic through the IPSec tunnel.

Finally, you apply the entire service set.

To configure IPSec service sets:

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.
3. Go on to “Configuring an IPSec Stateful Firewall Filter Rule” on page 387.

Table 107: Configuring IPSec Service Sets

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the next-hop service set for the IPSec tunnel.	1. From the top of the configuration hierarchy, click Services .	1. From the top of the configuration hierarchy, enter
Use any unique string for the service set name.	2. Next to Service sets, click Add new entry .	<code>edit services</code>
You must include an interface name and unit number for the inside-service interface and the outside-service interface. By default, J-Web Quick Configuration uses the following values:	3. In the Service set name box, type the name of the service set.	2. Set the inside-service interface:
<ul style="list-style-type: none"> ■ For the inside-service interface—<code>sp-0/0/0.1001</code> ■ For the outside-service interface—<code>sp-0/0/0.2002</code> 	4. In the Service type choice box, select Next hop service from the list.	<code>set service-set service-set-name next-hop-service inside-service-interface sp-0/0/0.1001</code>
	5. In the Nested configuration box, click Next hop service .	3. Set the outside-service interface:
	6. In the Inside service interface box, type <code>sp-0/0/0.1001</code> .	<code>set service-set service-set-name next-hop-service outside-service-interface sp-0/0/0.2002</code>
	7. Click OK .	
	8. In the Nested configuration box, click Next hop service .	
	9. In the Outside service interface box, type <code>sp-0/0/0.2002</code> .	
	10. Click OK .	
Configure the IP address of the local gateway for the IPSec service set to the local tunnel endpoint—for example, <code>1.1.1.1</code> .	1. Next to Ipsec vpn options, click Configure .	Set the local gateway address for the service set:
	2. In the Local gateway box, type <code>1.1.1.1</code> .	<code>set service-set service-set-name ipsec-vpn-options local-gateway 1.1.1.1</code>

Table 107: Configuring IPSec Service Sets (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure IPSec rules to set the IP address of the remote gateway—for example, 2.2.2.2—on all traffic.	1. From the top of the configuration hierarchy, click Services > Ipsec-vpn .	1. From the top of the configuration hierarchy, enter edit services ipsec-vpn
Use any unique string for the rule name.	2. Next to Rule, click Add new entry .	2. Configure a rule with a term that sets the remote gateway to 2.2.2.2:
Because the rule applies to all traffic, you must only configure the action (or then statement) for the term. Use any unique string for the term name.	3. In the Rule name box, type the name of the rule. 4. Next to the term, click Add new entry . 5. In the Term name box, type the name of the term. 6. To configure an action, click Then . 7. In the Remote gateway box, type 2.2.2.2. 8. Click OK .	set rule <i>rule-name</i> term <i>term-name</i> then remote-gateway 2.2.2.2
Configure a security association with a static IKE key.	1. From the top of the configuration hierarchy, select Services > Ipsec-vpn > Ike .	1. From the top of the configuration hierarchy, enter edit services ipsec-vpn ike
The IKE key is a preshared key and must be configured exactly the same way at both the local and remote endpoints of the IPSec tunnel.	2. Next to Policy, click Add new entry .	2. Configure the IKE pre-shared key in ASCII text format:
The IKE key is configured as ike policy and then applied with the dynamic statement. Use any unique string for the IKE policy name.	3. In the Name box, type the name of the IKE policy. 4. Click Pre-shared key . 5. In the Key choice box, select Ascii text from the list. 6. In the Ascii text box, type the IKE key in plain text. 7. Click OK . 8. Navigate to the IPSec rule configured previously. From the top of the configuration hierarchy, click Services > Ipsec-vp > rule-name > term term-name > then . 9. Click Dynamic . 10. In the Ike-policy box, type the name of the IKE policy you configured. 11. Click OK .	set policy <i>policy-name</i> pre-shared-key ascii-text <i>ike-key</i> 3. Navigate to the IPSec rule configured previously. From the top of the configuration hierarchy, enter edit services ipsec-vpn <i>rule-name</i> term <i>term-name</i> then. 4. Configure a dynamic security association that applies the IKE policy: set dynamic ike-policy <i>policy-name</i>

Table 107: Configuring IPsec Service Sets (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the IPsec rule so that it acts on input traffic.	<ol style="list-style-type: none"> From the top of the configuration hierarchy, click Services > Ipsec-vpn > Rule > rule-name. In the Match direction box, select Input from the list. Click OK. 	<ol style="list-style-type: none"> From the top of the configuration hierarchy, enter <code>edit services ipsec-vpn rule rule-name</code> Set the match direction for the rule: <code>set match-direction input</code>
Apply the IPsec rule to all traffic through the previously configured service set.	<ol style="list-style-type: none"> From the top of the configuration hierarchy, click Services > Service-set > service-set-name. In the Ipsec vpn rules choice box, select Ipsec vpn rules from the list. Next to Ipsec vpn rules, click Add new entry. In the Rule name box, type the name of the previously configured IPsec rule. Click OK. 	<ol style="list-style-type: none"> From the top of the configuration hierarchy, enter <code>edit services service-set service-set-name</code> Apply the IPsec rule previously configured: <code>set ipsec-vpn-rules rule-name</code>

Configuring an IPsec Stateful Firewall Filter Rule

If you have configured a stateful firewall filter that designates the interface through which an IPsec tunnel is configured as an *untrusted* interface, you must create a new stateful firewall filter rule that allows IPsec traffic to pass.

For more information about firewall filters, see “Configuring Firewall Filters and NAT” on page 475.

To configure an IPsec stateful firewall filter:

- Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- Perform the configuration tasks described in Table 108.
- Go on to “Configuring a NAT Pool” on page 389.

Table 108: Configuring an IPSec Stateful Firewall Filter Rule

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Create the stateful firewall rule and apply it to inbound traffic.</p> <p>Use any unique string for the rule name.</p>	<ol style="list-style-type: none"> From the top of the configuration hierarchy, click Services > Stateful firewall. Next to the rule, click Add new entry. In the Rule name box, type the name of the rule. From the Match direction list, select Input. 	<ol style="list-style-type: none"> From the top of the configuration hierarchy, enter <code>edit services stateful-firewall</code> Create the firewall rule and apply it to input traffic: <code>set rule rule-name match-direction input</code>
<p>Create the firewall term to match only desired traffic.</p> <p>Use any unique string for the term name.</p>	<ol style="list-style-type: none"> Next to Term, click Add new entry. In the Term name box, type the name of the term. Click From. Next to Destination address, click Add new entry. From the address list, select Enter specific value. In the Address box, type the IP address of the local tunnel endpoint, in dotted decimal notation, and click OK. Next to Source address, click Add new entry. From the address list, select Enter specific value. In the Address box, type the IP address of the remote tunnel endpoint, in dotted decimal notation, and click OK. Next to Applications, click Add new entry. In the Application name box, type <code>junos-ipsec-esp</code>, and click OK. Next to Applications, click Add new entry. In the Application name box, type <code>junos-ike</code>, and click OK. 	<ol style="list-style-type: none"> Create the firewall term and match all packets with a destination address that matches the local tunnel endpoint: <code>set term term-name from destination-address local-tunnel-end-point-address</code> Match all packets with a source address that matches the remote tunnel endpoint: <code>set term term-name from source-address remote-tunnel-end-point-address</code> Match all packets using IPSec as an application protocol: <code>set term term-name from applications junos-ipsec-esp</code> Match all packets using IKE as an application protocol: <code>set term term-name from applications junos-ike</code>

Table 108: Configuring an IPSec Stateful Firewall Filter Rule (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the firewall term to accept only desired traffic.	<ol style="list-style-type: none"> 1. Click OK to return to the Term name page, and click Then. 2. From the Designation list, select Accept, then select the Yes box. 3. Click OK. 	<p>Set the match action to accept:</p> <p>set term <i>term-name</i> then accept</p>
Create the firewall term to reject all other traffic. Use any unique string for the term name.	<ol style="list-style-type: none"> 1. From the top of the configuration hierarchy, click Services > Stateful firewall > Rule > <i>rule-name</i> 2. Next to Term, click Add new entry. 3. In the Term name box, type the name of the term. 4. Click Then. 5. From the Designation list, select Discard. 6. Click OK. 	<ol style="list-style-type: none"> 1. From the top of the configuration hierarchy, enter edit services stateful-firewall rule <i>rule-name</i> 2. Configure a term to discard all traffic: set term <i>term-name</i> then discard

Configuring a NAT Pool

To hide internal IP addresses from the rest of the Internet, you configure the local tunnel endpoint as the only address in a Network Address Translation (NAT) pool, to ensure that it is the address used for address translation.

To configure a NAT pool for IPSec:

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.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying the IPSec Tunnel Configuration” on page 391.

Table 109: Configuring a NAT Pool for IPSec

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the NAT pool from which the addresses for Network Address Translation are taken.	1. From the top of the configuration hierarchy, click Services > Nat .	1. From the top of the configuration hierarchy, enter
Name the NAT pool with any unique string of fewer than 64 characters.	2. Next to Pool, click Add new entry .	<code>edit services nat</code>
Provide the IP address of the local tunnel endpoint—for example, 1.1.1.1.	3. In the Pool name box, type the name of the NAT pool.	2. Add the local tunnel endpoint to the NAT address pool:
	4. From the the Address choice list, select Address .	<code>set pool pool-name address 1.1.1.1</code>
	5. In the Address box, type 1.1.1.1.	
Configure the router so that all outgoing traffic is matched against the IP address of the local tunnel endpoint.	1. From the top of the configuration hierarchy, click Services > Nat .	1. From the top of the configuration hierarchy, enter
Use any unique string for the NAT rule name and for the name of the term in the rule.	2. Next to Rule, click Add new entry .	<code>edit services nat</code>
The source address must be the IP address of the local tunnel endpoint—for example, 1.1.1.1.	3. In the Rule name box, type the name of the rule.	2. Configure a NAT rule and apply it to all output traffic:
	4. From the Match direction list, select Output .	<code>set rule rule-name match-direction output</code>
	5. Next to Term, click Add new entry .	3. Configure the rule to match traffic with a source address that is the same as the local tunnel endpoint:
	6. In the Term name box, type the name of the term.	<code>set rule rule-name term term-name from source-address 1.1.1.1</code>
	7. Click From .	
	8. Next to Source address, click Add new entry .	
	9. From the address list, select Enter specific value .	
	10. In the Address box, type 1.1.1.1.	
	11. Click OK .	
Configure the router so that the source address for traffic through the local endpoint is translated to the local endpoint address.	1. From the top of the configuration hierarchy, click Services > Nat > Rule > rule-name Term > term-name .	1. From the top of the configuration hierarchy, enter
	2. Click Then .	<code>edit services nat rule rule-name term term-name</code>
	3. Click Translated .	2. Configure the source pool:
	4. In the Source pool box, type the name of the NAT pool in which the local tunnel endpoint is configured.	<code>set then translated source-pool pool-name</code>
	5. From the Source list, select Static .	3. Configure the type of translation:
	6. Click OK .	<code>set then translated translation-type source static</code>

Verifying the IPsec Tunnel Configuration

To verify the IPsec tunnel configuration, perform the following task.

Verifying IPsec Tunnel Statistics

Purpose Verify that traffic is being sent through the configured IPsec tunnel.

Action From the CLI, enter the `show services ipsec-vpn ipsec statistics` command.

Sample Output

```
user@host> show services ipsec-vpn ipsec statistics

PIC: sp-0/0/0, Service set: service-set-1

Local gateway: 1.1.1.1, Remote gateway: 2.2.2.2, Tunnel index: 1
ESP Statistics:
  Encrypted bytes:                0
  Decrypted bytes:                0
  Encrypted packets:             0
  Decrypted packets:             0
AH Statistics:
  Input bytes:                   0
  Output bytes:                  0
  Input packets:                 0
  Output packets:                0
Errors:
  AH authentication failures: 0, Replay errors: 0
  ESP authentication failures: 0, Decryption errors: 0
  Bad headers: 0 Bad trailers: 0
```

What It Means The output shows the statistics for the particular service set that defines the IPsec tunnel, including the local and remote gateway addresses, the number of packets that have been encrypted and transported, and the number of errors and failures. Verify the following information:

- The local and remote tunnel endpoints are configured correctly.
- The number of Authentication Header (AH) and Encapsulation Security Payload (ESP) errors is zero. If these numbers are nonzero, the Services Router might be having a problem either transmitting or receiving encrypted packets through the IPsec tunnel.

For more information about `show services ipsec-vpn ipsec statistics`, see the *JUNOS System Basics and Services Command Reference*.

Part 5

Managing Multicast Transmissions

- Multicast Overview on page 395
- Configuring a Multicast Network on page 405

Chapter 17

Multicast Overview

Multicast traffic lies between the extremes of unicast (one source, one destination) and broadcast (one source, all destinations). Multicast is a “one source, many destinations” method of traffic distribution, meaning that the destinations needing to receive the information from a particular source receive the traffic stream.

IP network destinations (clients) do not often communicate directly with sources (servers), so the routers between source and destination must be able to determine the topology of the network from the unicast or multicast perspective to avoid routing traffic haphazardly. The multicast router must find multicast sources on the network, send out copies of packets on several interfaces, prevent routing loops, connect interested destinations with the proper source, and keep the flow of unwanted packets to a minimum. Standard multicast routing protocols provide most of these capabilities.

This chapter contains the following topics. For more information about multicast, see the *JUNOS Multicast Protocols Configuration Guide*. For configuration instructions, see “Configuring a Multicast Network” on page 405.

- Multicast Terms on page 395
- Multicast Architecture on page 398
- Dense and Sparse Routing Modes on page 400
- Strategies for Preventing Routing Loops on page 400
- Multicast Protocol Building Blocks on page 401

Multicast Terms

To understand multicast routing, you must be familiar with the terms defined in Table 110. See Figure 75 for a general view of some of the elements commonly used in an IP multicast network architecture.

Table 110: Multicast Terms

Term	Definition
administrative scoping	Multicast routing strategy that limits the routers and interfaces used to forward a multicast packet by reserving a range of multicast addresses.
Auto-RP	Cisco multicast routing protocol that allows sparse-mode routing protocols to find rendezvous points (RPs) within a routing domain.
bootstrap router (BSR)	Multicast mechanism that allows routers running PIM sparse mode to find rendezvous points (RPs) within a routing domain.
branch	Part of a multicast network that is formed when a leaf subnetwork is joined to the multicast distribution tree. Branches with no interested receivers are pruned from the tree so that multicast packets are no longer replicated on the branch.
broadcast routing protocol	Protocol that distributes traffic from a particular source to all destinations.
dense mode	Multicast routing mode appropriate for LANs with many interested receivers.
Distance Vector Multicast Routing Protocol (DVMRP)	Distributed multicast routing protocol that dynamically generates IP multicast distribution trees using reverse-path multicasting (RPM) to forward multicast traffic to downstream interfaces.
distribution tree	Path linking multicast receivers (listeners) to sources. The root of the tree is at the source, and the branches connect subnetworks of interested receivers (leaves). Multicast packets are replicated only where a distribution tree branches. To shorten paths to a source at the edge of a network, sparse mode multicast protocols can use a <i>shared</i> distribution tree located more centrally in the network backbone.
downstream interface	Interface on a multicast router that is leading toward the receivers. You can configure all the logical interfaces except one as downstream interfaces.
group address	Multicast destination address. A multicast network uses the Class D IP address of a logical group of multicast receivers to identify a destination. IP multicast packets have a multicast group address as the destination address and a unicast source address.
Internet Group Management Protocol (IGMP)	Multicast routing protocol that runs between receiver hosts and routers to determine whether group members are present. Services Routers support IGMPv1, IGMPv2, and IGMPv3.
leaf	IP subnetwork that is connected to a multicast router and that includes at least one host interested in receiving IP multicast packets. The router must send a copy of its multicast packets out on each interface with a leaf, and its action is unaffected by the number of leaves on the interface.
listener	Another name for a receiver in a multicast network.

Table 110: Multicast Terms (continued)

Term	Definition
multicast routing protocol	Protocol that distributes traffic from a particular source to only the destinations needing to receive it. Typical multicast routing protocols are the Distance Vector Multicast Routing Protocol (DVMRP) and Protocol Independent Multicast (PIM).
Multicast Source Discovery Protocol (MSDP)	Multicast routing protocol that connects multicast routing domains and allows them to find rendezvous points (RPs).
Pragmatic General Multicast (PGM)	Special protocol layer for multicast traffic that can be used between the IP layer and the multicast application to add reliability to multicast traffic.
Protocol Independent Multicast (PIM) protocol	Protocol-independent multicast routing protocol that can be used in either sparse or dense mode. In sparse mode, PIM routes to multicast groups that might span WANs and interdomain Internets. In dense mode, PIM is a flood-and-prune protocol.
pruning	Removing from a multicast distribution tree branches that no longer include subnetworks with interested hosts. Pruning ensures that packets are replicated only as needed.
reverse-path forwarding (RPF)	Multicast routing strategy that allows a router to receive packets through an interface if it is the same interface a unicast packet uses as the shortest path back to the source.
rendezvous point (RP)	Core router operating as the root of a shared distribution tree in a multicast network.
Session Announcement Protocol (SAP)	Multicast routing protocol used with other multicast protocols—typically Session Description Protocol (SDP)—to handle session conference announcements.
Session Description Protocol (SDP)	Session directory protocol that advertise multimedia conference sessions and communicates setup information to participants who want to join the session.
shortest-path tree (SPT)	Multicast routing strategy for sparse mode multicast protocols. SPT uses a shared distribution tree rooted in the network backbone to shorten paths to sources at the edge of a network.
source-specific multicast (SSM)	Service that allows a client to receive multicast traffic directly from the source, without the help of a rendezvous point (RP).
sparse mode	Multicast routing mode appropriate for WANs with few interested receivers.
unicast routing protocol	Protocol that distributes traffic from one source to one destination.
upstream interface	Interface on a multicast router that is leading toward the source. To minimize bandwidth use, configure only one upstream interface on a router receiving multicast packets.

Multicast Architecture

Multicast-capable routers replicate packets on the multicast network, which has exactly the same topology as the unicast network it is based on. Multicast routers use a multicast routing protocol to build a distribution tree that connects receivers (also called listeners) to sources.

Upstream and Downstream Interfaces


A single upstream interface on the router leads toward the source to receive multicast packets. The downstream interfaces on the router lead toward the receivers to transmit packets. A router can have as many downstream interfaces as it has logical interfaces, minus 1. To prevent looping, the router's upstream interface must never receive copies of its own downstream multicast packets.

Subnetwork Leaves and Branches

On a multicast router, each subnetwork of hosts that includes at least one interested receiver is a leaf on the multicast distribution tree (see Figure 75). The router must send out a copy of the IP multicast packet on each interface with a leaf. When a new leaf subnetwork joins the tree, a new branch is built so that the router can send out replicated packets on the interface. The number of leaves on an interface does not affect the router. The action is the same for one leaf or a hundred.

A branch that no longer has leaves is pruned from the distribution tree. No multicast packets are sent out on a router interface leading to an IP subnetwork with no interested hosts. Because packets are replicated only where the distribution tree branches, no link ever carries a duplicate flow of packets.

In IP multicast networks, traffic is delivered to multicast groups based on an IP multicast group address instead of a unicast destination address. The groups determine the location of the leaves, and the leaves determine the branches on the multicast network.

Figure 75: Multicast Elements in an IP Network

Multicast IP Address Ranges

Multicast uses the Class D IP address range (224.0.0.0 through 239.255.255.255). Multicast addresses usually have a prefix length of /32, although other prefix lengths are allowed. Multicast addresses represent logical groupings of receivers and not physical collections of devices, and can appear only as the destination in an IP packet, never as the source address.

Notation for Multicast Forwarding States

The multicast forwarding state in a router is usually represented by one of the following notations:

- (S,G) notation—S refers to the unicast IP address of the source for the multicast traffic and G refers to the particular multicast group IP address for which S is the source. All multicast packets sent from this source have S as the source address and G as the destination address.
- (*, G) notation—The asterisk (*) is a wildcard for the address of any multicast application source sending to group G. For example, if two sources are originating exactly the same content for multicast group 224.1.1.2, a router can use (*, 224.1.1.2) to represent the state of a router forwarding traffic from both sources to the group.

Dense and Sparse Routing Modes

To keep packet replication to a minimum, multicast routing protocols use the two primary modes shown in Table 111.



CAUTION: A common multicast guideline is *not to run dense mode on a WAN under any circumstances*.

Table 111: Primary Multicast Routing Modes

Multicast Mode	Description	Appropriate Network for Use
Dense mode	Network is flooded with traffic on all possible branches, then pruned back as branches explicitly (by message) or implicitly (time-out silence) eliminate themselves.	LANs—Networks in which all possible subnets are likely to have at least one receiver.
Sparse mode	Network establishes and sends packets only on branches that have at least one leaf indicating (by message) a need for the traffic.	WANs—Network in which very few of the possible receivers require packets from this source.

Strategies for Preventing Routing Loops

Routing loops are disastrous in multicast networks because of the risk of repeatedly replicated packets, which can overwhelm a network. One of the complexities of modern multicast routing protocols is the need to avoid routing loops, packet by packet, much more rigorously than in unicast routing protocols. Three multicast strategies—reverse-path forwarding (RPF), shortest-path tree (SPT), and administrative scoping—help prevent routing loops by defining routing paths in different ways.

Reverse-Path Forwarding for Loop Prevention

The router's multicast forwarding state runs more logically based on the reverse path, from the receiver back to the root of the distribution tree. In reverse-path forwarding (RPF), every multicast packet received must pass an RPF check before it can be replicated or forwarded on any interface. When it receives a multicast packet on an interface, the router verifies that the *source* address in the multicast IP packet is the *destination* address for a unicast IP packet back to the source.

If the outgoing interface found in the unicast routing table is the same interface that the multicast packet was received on, the packet passes the RPF check. Multicast packets that fail the RPF check are dropped, because the incoming interface is not on the shortest path back to the source. Routers can build and maintain separate tables for RPF purposes.

Shortest-Path Tree for Loop Prevention

The distribution tree used for multicast is rooted at the source and is the shortest-path tree (SPT), but this path can be long if the source is at the periphery of the network. Providing a *shared tree* on the backbone as the distribution tree locates the multicast source more centrally in the network. Shared distribution trees with roots in the core network are created and maintained by a multicast router operating as a rendezvous point (RP), a feature of sparse mode multicast protocols.

Administrative Scoping for Loop Prevention

Scoping limits the routers and interfaces that can forward a multicast packet. Multicast scoping is *administrative* in the sense that a range of multicast addresses is reserved for scoping purposes, as described in RFC 2365, *Administratively Scoped IP Multicast*. Routers at the boundary must filter multicast packets and ensure that packets do not stray beyond the established limit.

Multicast Protocol Building Blocks

Multicast is not a single protocol, but a collection of protocols working together to form trees, prune branches, locate sources and groups, and prevent routing loops:

- Distance Vector Multicast Routing Protocol (DVMRP) and Protocol Independent Multicast (PIM) operate between routers. PIM can operate in dense mode and sparse mode.
- Three versions of the Internet Group Management Protocol (IGMP) run between receiver hosts and routers.
- Several other routing mechanisms and protocols enhance multicast networks by providing useful functions not included in other protocols. These include the bootstrap router (BSR) mechanism, Auto-RP protocol, Multicast Source Discovery Protocol (MSDP), Session Announcement Protocol (SAP) and Session Discovery Protocol (SDP), and Pragmatic General Multicast (PGM) protocol.

Table 112 lists and summarizes these protocols.

Table 112: Multicast Protocol Building Blocks

Multicast Protocol	Description	Uses
DVMRP	Dense-mode-only protocol that uses the flood-and-prune or implicit join method to deliver traffic everywhere and then determine where the uninterested receivers are. DVMRP uses source-based distribution trees in the form (S,G) and builds its own multicast routing tables for RPF checks.	Not appropriate for large-scale Internet use.

Table 112: Multicast Protocol Building Blocks (continued)

Multicast Protocol	Description	Uses
PIM dense mode	<p>Sends an <i>implicit</i> join message, so routers use the flood-and-prune method to deliver traffic everywhere and then determine where the uninterested receivers are.</p> <p>PIM dense mode uses source-based distribution trees in the form (S,G), and also supports sparse-dense mode, with mixed sparse and dense groups. Both PIM modes use unicast routing information for RPF checks.</p>	Most promising multicast protocol in use for LANs.
PIM sparse mode	<p>Sends an <i>explicit</i> join message, so routers determine where the interested receivers are and send join messages upstream to their neighbors, building trees from receivers to a rendezvous point (RP) router, which is the initial source of multicast group traffic.</p> <p>PIM sparse mode builds distribution trees in the form (*,G), but migrates to an (S,G) source-based tree if that path is shorter than the path through the RP router for a particular multicast group's traffic. Both PIM modes use unicast routing information for RPF checks.</p>	Most promising multicast protocol in use for WANs.
PIM source-specific multicast (SSM)	Enhancement to PIM sparse mode that allows a client to receive multicast traffic directly from the source, without the help of a rendezvous point (RP).	Used with IGMPv3 to create a shortest-path tree between receiver and source.
IGMPv1	The original protocol defined in RFC 1112, <i>Host Extensions for IP Multicasting</i> . IGMPv1 sends an explicit join message to the router, but uses a time-out to determine when hosts leave a group.	
IGMPv2	Defined in RFC 2236, <i>Internet Group Management Protocol, Version 2</i> . Among other features, IGMPv2 adds an explicit leave message to the join message.	Used by default.
IGMPv3	Defined in RFC 3376, <i>Internet Group Management Protocol, Version 3</i> . Among other features, IGMPv3 optimizes support for a single source of content for a multicast group, or <i>source-specific multicast (SSM)</i> .	Used with PIM SSM to create a shortest-path tree between receiver and source.
BSR Auto-RP	Allow sparse-mode routing protocols to find rendezvous points (RPs) within the routing domain (autonomous system, or AS). RP addresses can also be statically configured.	

Table 112: Multicast Protocol Building Blocks (continued)

Multicast Protocol	Description	Uses
MSDP	Allows groups located in one multicast routing domain to find rendezvous points (RPs) in other routing domains. MSDP is not used on an RP if all receivers and sources are located in the same routing domain.	Typically runs on the same router as PIM sparse mode rendezvous point (RP). Not appropriate if all receivers and sources are located in the same routing domain.
SAP and SDP	Display multicast session names and correlate the names with multicast traffic. SDP is a session directory protocol that advertises multimedia conference sessions and communicates setup information to participants who want to join the session. A client commonly uses SDP to announce a conference session by periodically multicasting an announcement packet to a well-known multicast address and port using SAP.	
PGM	Special protocol layer for multicast traffic that can be used between the IP layer and the multicast application to add reliability to multicast traffic. PGM allows a receiver to detect missing information in all cases and request replacement information if the receiver application requires it.	

Chapter 18

Configuring a Multicast Network

You configure a router network to support multicast applications with a related family of protocols. To use multicast, you must understand the basic components of a multicast network and their relationships, and then configure the J-series Services Router to act as a node in the network.



NOTE: The J-series Services Router supports both Protocol Independent Multicast (PIM) version 1 and PIM version 2. In this chapter, the term *PIM* refers to both versions of the protocol.

You use either the J-Web configuration editor or CLI configuration editor to configure multicast protocols. The J-Web interface does not include Quick Configuration pages for multicast protocols.

This chapter contains the following topics. For more information about multicast, see the *JUNOS Multicast Protocols Configuration Guide*.

- Before You Begin on page 406
- Configuring a Multicast Network with a Configuration Editor on page 406
- Verifying a Multicast Configuration on page 411

Before You Begin

Before you begin configuring a multicast network, complete the following tasks:

- If you do not already have a basic understanding of multicast, read “Multicast Overview” on page 395.
- Determine whether the Services Router is directly attached to any multicast sources. Receivers must be able to locate these sources.
- Determine whether the Services Router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
- Determine whether to use the sparse, dense, or sparse-dense mode of multicast operation. Each mode has different configuration considerations.
- Determine the address of the rendezvous point (RP) if sparse or sparse-dense mode is used.
- Determine whether to locate the RP with the static configuration, bootstrap router (BSR), or Auto-RP method.
- Determine whether to configure multicast to use its own reverse-path forwarding (RPF) routing table when configuring PIM in sparse, dense, or sparse-dense modes.

Configuring a Multicast Network with a Configuration Editor

To configure the Services Router as a node in a multicast network, you must perform the following tasks marked *(Required)*. For information about using the J-Web and CLI configuration editors, see “Using J-Web Configuration Tools” on page 3.

- Configuring SAP and SDP (Optional) on page 406
- Configuring IGMP (Required) on page 407
- Configuring the PIM Static RP (Optional) on page 408
- Configuring a PIM RPF Routing Table (Optional) on page 410

Configuring SAP and SDP (Optional)

Multicast session announcements are handled by two protocols, the Session Announcement Protocol (SAP) and Session Description Protocol (SDP). These two protocols display multicast session names and correlate the names with multicast traffic. Enabling SDP and SAP allows the router to receive announcements about multimedia and other multicast sessions from sources. Enabling SAP automatically enables SDP.

For more information on SAP and SDP, see the *JUNOS Multicast Protocols Configuration Guide*.

The Services Router listens for session announcements on one or more addresses and ports. By default, the router listens to address and port 224.2.127.254:9875.

To configure SAP and SDP 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 113.
3. Go on to “Configuring IGMP (Required)” on page 407.

Table 113: Configuring SAP and SDP

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Listen level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Protocols > Sap. 2. Click Add new entry next to Listen. 	<p>From the top of the configuration hierarchy, enter</p> <pre>edit protocols sap</pre>
(Optional) Enter one or more addresses and ports for the Services Router to listen to session announcements on. By default, the Services Router listens to address and port 224.2.127.254:9875.	<ol style="list-style-type: none"> 1. In the Address box, type the multicast address the Services Router can listen to session announcements on, in dotted decimal notation. 2. In the Port box, type the port number in decimal notation. 3. Click OK. 	<ol style="list-style-type: none"> 1. Set the address value to the IP address that the Services Router can listen to session announcements on, in dotted decimal notation. For example: <pre>set listen 224.2.127.254</pre> 2. Set the port value to the number of the port that the Services Router can listen to session announcements on, in decimal notation. For example: <pre>set listen 224.2.127.254 port 9875</pre>

Configuring IGMP (Required)

The Internet Group Management Protocol (IGMP) manages the membership of hosts and routers in multicast groups. IGMP is an integral part of IP and must be enabled on all routers and hosts that need to receive IP multicasts. IGMP is automatically enabled on all broadcast interfaces when you configure PIM or DVMRP.

For more information on IGMP, see *JUNOS Multicast Protocols Configuration Guide*.

By default, the Services Router runs IGMPv2. However, you might still want to set the IGMP version explicitly on an interface, or all interfaces. Routers running different versions of IGMP negotiate the lowest common version of IGMP supported by hosts on their subnet. One host running IGMPv1 forces the Services Router to use that version and lose features important to other hosts.

To explicitly configure the IGMP version, 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.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To configure PIM sparse mode, see “Configuring the PIM Static RP (Optional)” on page 408.
 - To check the configuration, see “Verifying a Multicast Configuration” on page 411.

Table 114: Explicitly Configuring the IGMP version

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 configuration editor hierarchy, select Protocols > Igmp. 2. Click Add new entry next to Interface. 	<p>From the top of the configuration hierarchy, enter</p> <p>edit protocols igmp</p>
Set the IGMP version. By default, the Services Router uses IGMPv2, but this version can be changed through negotiation with hosts unless explicitly configured.	<ol style="list-style-type: none"> 1. In the Interface name box, type the name of the interface, or all. 2. In the Version box, type the version number: 1, 2, or 3. 3. Click OK. 	<ol style="list-style-type: none"> 1. Set the interface value to the interface name, or all. For example: set igmp interface all 2. Set the version value to 1, 2, or 3. For example: set igmp interface all version 2

Configuring the PIM Static RP (Optional)

Protocol Independent Multicast (PIM) sparse mode is the most common multicast protocol used on the Internet. PIM sparse mode is the default mode whenever PIM is configured on any interface of the Services Router. However, because PIM must not be configured on the network management interface of the Services Router, you must disable it on that interface.

Each any-source multicast (ASM) group has a shared tree through which receivers learn about new multicast sources and new receivers learn about all multicast sources. The rendezvous point (RP) router is the root of this shared tree and receives the multicast traffic from the source. To receive

multicast traffic from the groups served by the RP, the Services Router must determine the IP address of the RP for the source.

One common way for the Services Router to locate RPs is by static configuration of the IP address of the RP. For information about alternate methods of locating RPs, see the *JUNOS Multicast Protocols Configuration Guide*.

To configure PIM sparse mode, disable PIM on `fe-0/0/0`, and configure the IP address of the RP 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 115.
3. Go on to “Configuring a PIM RPF Routing Table (Optional)” on page 410.

Table 115: Configuring PIM Sparse Mode and the RP

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 configuration editor hierarchy, select Protocols > Pim. 2. Click Add new entry next to Interface. 	From the top of the configuration hierarchy, enter <code>edit protocols pim</code>
Enable PIM on all network interfaces.	In the Interface name box, type <code>all</code> .	Set the <code>interface</code> value to <code>all</code> . For example: <code>set pim interface all</code>
Apply your configuration changes.	Click OK to apply your entries to the configuration.	Changes in the CLI are applied automatically when you execute the <code>set</code> command.
Remain at the Interface level in the configuration hierarchy.	Click Add new entry next to Interface.	Remain at the <code>edit protocols pim interface</code> configuration hierarchy level.
Disable PIM on the network management interface.	<ol style="list-style-type: none"> 1. In the Interface name box, type <code>fe-0/0/0</code>. 2. Select the check box next to Disable. 	Disable the <code>fe-0/0/0</code> interface: <code>set pim interface fe-0/0/0 unit 0 disable</code>
Apply your configuration changes.	Click OK to apply your entries to the configuration.	Changes in the CLI are applied automatically when you execute the <code>set</code> command.

Table 115: Configuring PIM Sparse Mode and the RP (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Rp level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Pim > Rp .	From the top of the configuration hierarchy, enter edit protocols pim rp
Configure the IP address of the RP—for example, 192.168.14.27.	<ol style="list-style-type: none"> 1. Click Configure next to Static. 2. Click Add new entry next to Address. 3. In the Addr box, type 192.168.14.27. 4. Click OK. 	Set the address value to the IP address of the RP: set static address 192.168.14.27

Configuring a PIM RPF Routing Table (Optional)

By default, PIM uses inet.0 as its reverse-path forwarding (RPF) routing table group. PIM uses an RPF routing table group to resolve its RPF neighbor for a particular multicast source address and for the RP address. PIM can optionally use inet.2 as its RPF routing table group. The inet.2 routing table is organized more efficiently for RPF checks.

Once configured, the RPF routing table must be applied to PIM as a routing table group.

To configure and apply a PIM RPF routing table, 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 116.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying a Multicast Configuration” on page 411.

Table 116: Configuring a PIM RPF Routing Table

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Routing options level in the configuration hierarchy.	In the configuration editor hierarchy, select Routing options .	From the top of the configuration hierarchy, enter edit routing-options

Table 116: Configuring a PIM RPF Routing Table (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a new group for the RPF routing table.	Next to Rib groups, click Add new entry .	Enter edit rib-groups
Configure a name for the new RPF routing table group—for example, multicast-rpf-rib —and use inet.2 for its export routing table.	1. In the Ribgroup name box, type multicast-rpf-rib . 2. In the Export rib box, type inet.2 .	Enter set multicast-rpf-rib export-rib inet.2
Configure the new RPF routing table group to use inet.2 for its import routing table.	1. Click Add new entry next to Import rib. 2. In the Value box, type inet.2 . 3. Click OK three times.	Enter set multicast-rpf-rib import-rib inet.2
Navigate to the Rib group level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Pim > Rib group .	From the top of the configuration hierarchy, enter edit protocols pim
Apply the new RPF routing table to PIM.	1. In the Inet box, type the name of the RPF routing table group— multicast-rpf-rib . 2. Click OK three times.	Enter set rib-group multicast-rpf-rib
Create a routing table group for the interface routes.	1. Navigate to the Routing options level in the configuration hierarchy. 2. Next to Rib groups, click Add new entry .	From the top of the configuration hierarchy, enter edit routing-options rib-groups.
Configure a name for the RPF routing table group—for example, if-rib —and use inet.2 and inet.0 for its import routing tables.	1. In the Ribgroup name box, type if-rib . 2. Click Add new entry next to Import rib. 3. In the Value box, type inet.2 inet.0 . 4. Click OK twice.	Enter set if-rib import-rib inet.2 set if-rib import-rib inet.0
Add the new interface routing table group to the interface routes.	1. On the Routing options page, select Interface routes > Rib group . 2. In the Inet box, type if-rib . 3. Click OK .	From the top of the configuration hierarchy, enter edit routing-options interface-routes set rib-group inet if-rib

Verifying a Multicast Configuration

To verify a multicast configuration, perform these tasks:

- Verifying SAP and SDP Addresses and Ports on page 412

- Verifying the IGMP Version on page 412
- Verifying the PIM Mode and Interface Configuration on page 413
- Verifying the PIM RP Configuration on page 413
- Verifying the RPF Routing Table Configuration on page 414

Verifying SAP and SDP Addresses and Ports

Purpose Verify that SAP and SDP are configured to listen on the correct group addresses and ports.

Action From the CLI, enter the `show sap listen` command.

Sample Output `user@host> show sap listen`

```
Group Address    Port
224.2.127.254    9875
```

What It Means The output shows a list of the group addresses and ports that SAP and SDP listen on. Verify the following information:

- Each group address configured, especially the default 224.2.127.254, is listed.
- Each port configured, especially the default 9875, is listed.

For more information about `show sap listen`, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying the IGMP Version

Purpose Verify that IGMP version 2 is configured on all applicable interfaces.

Action From the CLI, enter the `show igmp interface` command.

Sample Output `user@host> show igmp interface`

```
Interface: fe-0/0/0.0
  Querier: 192.168.4.36
  State:           Up Timeout:      197 Version:  2 Groups:      0

Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0
```

What It Means The output shows a list of the Services Router interfaces that are configured for IGMP. Verify the following information:

- Each interface on which IGMP is enabled is listed.
- Next to Version, the number 2 appears.

For more information about `show igmp interface`, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying the PIM Mode and Interface Configuration

Purpose Verify that PIM sparse mode is configured on all applicable interfaces.

Action From the CLI, enter the `show pim interfaces` command.

Sample Output

```
user@host> show pim interfaces

Instance: PIM.master
Name                               Stat Mode      IP V State Count DR address
lo0.0                             Up   Sparse      4 2 DR        0 127.0.0.1
pim.32769                          Up   Sparse      4 2 P2P        0
```

What It Means The output shows a list of the Services Router interfaces that are configured for PIM. Verify the following information:

- Each interface on which PIM is enabled is listed.
- The network management interface, `fe-0/0/0`, is *not* listed.
- Under Mode, the word Sparse appears.

For more information about `show pim interfaces`, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying the PIM RP Configuration

Purpose Verify that the PIM RP is statically configured with the correct IP address.

Action From the CLI, enter the `show pim rps` command.

Sample Output

```
user@host> show pim rps

Instance: PIM.master
Address family INET
RP address      Type      Holdtime Timeout Active groups Group prefixes
192.168.14.27   static    0         None      2 224.0.0.0/4
```

What It Means The output shows a list of the RP addresses that are configured for PIM. At least one RP must be configured. Verify the following information:

- The configured RP is listed with the proper IP address.
- Under Type, the word `static` appears.

Verifying the RPF Routing Table Configuration

Purpose Verify that the PIM RPF routing table is configured correctly.

Action From the CLI, enter the `show multicast rpf` command.

Sample Output

```
user@host> show multicast rpf

Multicast RPF table: inet.0 , 2 entries...
```

What It Means The output shows the multicast RPF table that is configured for PIM. If no multicast RPF routing table is configured, RPF checks use `inet.0`. Verify the following information:

- The configured multicast RPF routing table is `inet.0`.
- The `inet.0` table contains entries.

For more information about `show multicast rpf`, see the *JUNOS Routing Protocols and Policies Command Reference*.

Part 6

Configuring DLSw Services

- Configuring Data Link Switching on page 417

Chapter 19

Configuring Data Link Switching

Data link switching (DLSw) was developed in the early 1990s as a method to transport IBM System Network Architecture (SNA) over a WAN network. To route traffic over a WAN link or the Internet, DLSw encapsulates the SNA network traffic in IP. The Services Router supports DLSw as part of an SNA implementation.



NOTE: You must have a license to configure DLSw. For license details, see the *J-series Services Router Administration Guide*.

You use either the J-Web configuration editor or CLI configuration editor to configure DLSw. For more information on DLSw, see the *JUNOS Services Interfaces Configuration Guide*.

This chapter contains the following topics.

- DLSw Terms on page 417
- Data Link Switching (DLSw) Overview on page 418
- Before You Begin on page 420
- Configuring Basic DLSw with a Configuration Editor on page 420
- Configuring Class-of-Service (CoS) for DLSw (Optional) on page 424
- Verifying DLSw Configuration on page 426

DLSw Terms

Before configuring DLSw on a Services Router, become familiar with the terms defined in Table 117.

Table 117: DLSw Terms


Term	Definition
DLSw circuit	Path formed by establishing a data link control (DLC) connection between each locally configured SNA end system and a local router configured for DLSw. A DLSw circuit is identified by the circuit ID, which includes the SNA end system MAC address, local service access point (LSAP), destination MAC address, and destination service access point (DSAP). Multiple DLSw circuits can operate over the same DLSw connection.
DLSw connection	Set of TCP connections between two DLSw peers that is established after the initial handshake and successful capabilities exchange.
destination service access point (DSAP)	Service access point (SAP) that identifies the destination for which an logical link control protocol data unit (LPDU) is intended.
I-frames	Information frame used to transfer sequentially numbered logical link control protocol data units (LPDUs) between link stations.
Logical Link Control (LLC)	Data-link layer protocol used on a LAN. LLC1 provides connectionless data transfer, and LLC2 provides connection-oriented data transfer.
LLC protocol data unit (LPDU)	Logical link control (LLC) frame on a DLSw network.
service access point (SAP)	OSI term for the component of a network address that identifies the individual application sending or receiving a packet on a host.
source service access point (SSAP)	Service access point (SAP) that identifies the origin of an LPDU on a DLSw network.
Switch-to-Switch Protocol (SSP)	Protocol implemented between two DLSw routers that establishes connections, locates resources, forwards data, and handles error recovery and flow control.

Data Link Switching (DLSw) Overview

Data link switching (DLSw) was developed in the 1990s as a method to transport IBM Systems Network Architecture (SNA) traffic over an IP WAN network. Switch-to-Switch Protocol (SSP) is used to forward network traffic between routers configured for DLSw (DLSw peers). Then, to route traffic over a WAN link or the Internet, DLSw encapsulates the SNA network traffic into IP packets.

DLSw was developed as a forwarding mechanism for IBM Systems Network Architecture (SNA) protocol. Although DLSw does not provide full routing capabilities, it provides switching at the data link layer and encapsulation in TCP/IP for transport over the Internet.

Because DLSw provides support for SNA, a connection-oriented protocol, the Services Router supports Logical Link Control (LLC) type 2 as part of the DLSw implementation. Figure 76 shows a possible DLSw network.

Figure 76: Sample DLSw Network

Switch-to-Switch Protocol for DLSw

Switch-to-Switch Protocol (SSP) is used between DLSw peers to establish connections, locate resources, forward data, and handle error recovery as well as flow control. Generally, SSP does not provide full routing between peers, because routing is typically handled by common routing protocols such as OSPF or BGP. Instead, packets are switched at the SNA data link layer and encapsulated in TCP/IP for transport over IP-based networks. TCP is used as reliable transport method between DLSw peers.

DLSw Operational Stages

There are several operational stages that take place in DLSw connections. First, two DLSw peers establish a TCP connection with each other. After the connection is established, each peer router exchanges supported capabilities with the other router. The TCP connection ensures reliable and guaranteed delivery of IP traffic, and also ensures the integrity and delivery of traffic encapsulated in the IP protocol. After capability information is exchanged, the DLSw peers establish circuits between SNA end systems and begin transmitting information frames (I-frames) over the network.

DLSw Capabilities Exchange

DLSw capabilities exchange is based on a switch-to-switch protocol message describing the capabilities of the sending data-link switch. Sent just after the DLSw peers establish a connection, a capabilities exchange control message communicates the following operational parameters between the two peers:

- DLSw version number
- Initial pacing window size (receive window size)
- List of supported link SAPs (LSAPs)
- Number of supported TCP sessions
- Lists of media access control (MAC) addresses

DLSw Circuits Establishment

Establishing DLSw circuits is a process in which local and remote DLSw peers locate each other and set up data link control (DLC) connections between the remote router and a local router. The specific details of establishing circuits are determined by the traffic type, but the process is the same for all types of traffic.

The first step in the process enables the SNA devices on a LAN to find other SNA devices by sending out an explorer frame with the MAC address of the target SNA device. When a DLSw peer receives the explorer frame, it sends a canureach message frame to each of its DLSw peer connections. The canureach message frame queries the DLSw peers to determine if one of the peers can locate the target SNA device. If one of the DLSw peers can reach the target SNA device, it returns an icanreach message frame to the originating DLSw peer to indicate that it can provide a path to the SNA device in question.

After canureach and icanreach message frames are exchanged, the two DLSw peers establish a circuit consisting of a DLC connection between each router and the local SNA end system and a TCP connection between the two DLSw peers. The resulting circuit is uniquely identified by source and destination circuit IDs. Each SNA DLSw circuit ID includes the following information:

- MAC address of the SNA end system
- Link service access point (LSAP)
- DLC port ID

Circuit priority is negotiated when the circuit is set up on the network.

Before You Begin

Before you begin configuring DLSw, complete the following tasks:

- Establish basic connectivity. See the *J-series Services Router Getting Started Guide*.
- Configure network interfaces. See “Configuring Network Interfaces” on page 103.
- If you do not already have an understanding of DLSw, read “Data Link Switching (DLSw) Overview” on page 418.

Configuring Basic DLSw with a Configuration Editor

To configure DLSw on a Services Router, perform the following tasks marked *(Required)*:

- Configuring LLC2 Properties on an Ethernet Interface (Required) on page 421
- Configuring DLSw on the Local Services Router (Required) on page 422

- Configuring DLSw on the Remote Services Router (Required) on page 423



NOTE: To configure advanced properties for DSLw, see the *JUNOS Services Interfaces Configuration Guide*.

Configuring LLC2 Properties on an Ethernet Interface (Required)

Before configuring DLSw on the Services Router, you must configure the LLC2 properties on the Ethernet interfaces of the router. The Logical Link Control (LLC) layer is one of two sublayers into which the OSI data link layer is subdivided for data link protocols used on the LAN. LLC2 is implemented anytime SNA is running on a LAN or virtual LAN.



NOTE: LLC2 properties must be configured on the local Services Router and the remote Services Router.

To configure LLC2 properties:

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 118.
3. Go on to one of the following required configurations:
 - To configure DLSw on the local Services Router, go on to “Configuring DLSw on the Local Services Router (Required)” on page 422.
 - To configure DLSw on the remote Services Router, go on to “Configuring DLSw on the Remote Services Router (Required)” on page 423.
4. To verify the basic DLSw properties, see “Verifying DLSw Configuration” on page 426.

Table 118: Configuring LLC2 Properties on a Fast Ethernet Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy and select a Fast Ethernet interface—for example, fe-3/0/1 .	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Interfaces > Edit. 2. Click fe-3/0/1. 	<p>From the top of the configuration hierarchy, enter</p> <p>edit interfaces fe-3/0/1</p>
Configure LLC2 properties on the fe-3/0/1 interface.	<ol style="list-style-type: none"> 1. Under Unit and Interface unit number, click 0. 2. Under Family, select Llc2. 3. Click OK until you return to the Configuration page. 	<ol style="list-style-type: none"> 1. Enter edit unit 0 2. Enter set family llc2

Configuring DLSw on the Local Services Router (Required)

To configure DLSw on the local Services Router, you do the following:

- Define a local peer.
- Define a remote peer.
- Finally, define connection behavior.

The example in this section shows how to configure DLSw on the local and remote Services Routers with IP addresses listed in Table 119. The remote Services Router initiates the peer connection.

Table 119: Sample DLSw Peer Router Values

Option	Value
remote-peer	217.110.111.134
local-peer	110.0.10.1

In this example, the local router is configured with **remote-peer** settings because the local router is initiating the connection for SNA traffic over the WAN interface. The remote router is accepting DLSw connections from any DLSw peers.

To configure basic DLSw on the local 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 120.
3. Go on to “Configuring DLSw on the Remote Services Router (Required)” on page 423.

Table 120: Configuring DLSw on the Local Router

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Dls w level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Protocols > Dlsw. 2. Click Configure. 	<p>From the top of the configuration hierarchy, enter</p> <pre>edit protocols dls</pre>
Configure the local router properties.	In the Local peer box, type 110.0.10.1 .	<p>Enter</p> <pre>set local-peer 110.0.10.1</pre>
<p>Configure the remote peer settings.</p> <p>Because the remote router is initiating the peer connection, configure the remote-peer setting.</p>	<ol style="list-style-type: none"> 1. Next to Remote peer, click Configure. 2. Click Add new entry. 3. In the Peer ip box, type 217.110.111.134. 4. Click OK until you return to the Protocols page. 	<p>Enter</p> <pre>set remote-peer 217.110.111.134</pre>

Configuring DLSw on the Remote Services Router (Required)

To configure DLSw on the remote Services Router, you do the following:

- Define a local peer.
- Define a remote peer.
- Finally, define the connection behavior.

To configure DLSw on a remote 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 121.
3. If you are finished configuring the router, commit the configuration.
4. To verify the DLSw configuration, see “Verifying DLSw Configuration” on page 426.

Table 121: Configuring DLSw on the Remote Router

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Dls w level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Protocols > Dlsw. 2. Click Configure. 	From the top of the configuration hierarchy, enter <code>edit protocols dls</code> w
Configure the local router properties. promiscuous —Allows all incoming peer connections.	<ol style="list-style-type: none"> 1. In the Local peer box, type 217.110.111.134. 2. Next to Promiscuous, select Yes. 3. Click OK. 	<ol style="list-style-type: none"> 1. Enter <code>set local-peer 217.110.111.134</code> 2. Enter <code>set promiscuous</code>

Configuring Class-of-Service (CoS) for DLSw (Optional)

The J-series Services Router CoS features provide differentiated services when best-effort traffic delivery is not enough. You can use CoS to classify DLSw packets. The packets are sent to a logical tunnel interface on the router, where they are classified and queued based on the configured type-of-service (ToS) value.

For information about CoS, see “Class-of-Service Overview” on page 449 or the *JUNOS Class of Service Configuration Guide*.

To configure CoS for DLSw on the Services Router, you do the following:

- Configure the logical tunnel `lt-0/0/0` interface.
- Configure the CoS classifier on the `lt-0/0/0` interface.
- Configure the DLSw type-of-service (ToS) precedence on the `lt-0/0/0` interface.

To configure CoS classification for DLSw on a 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 122.
3. If you are finished configuring the router, commit the configuration.

Table 122: Configuring DLSw on the Remote Router

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	In the configuration editor hierarchy, select Interfaces > Edit .	From the top of the configuration hierarchy, enter edit interfaces It-0/0/0
Configure the first logical unit on the It-0/0/0 interface.	<ol style="list-style-type: none"> Next to Interface, click Add new entry. In the Interface name box, type It-0/0/0. Click OK. Next to It-0/0/0, click Edit. Next to Unit, click Add new entry. In the Interface unit number box, type 0. In the Dlci box, type 10. From the Encapsulation list, select frame-relay. In the Peer unit box, type 1. Under Family, select Inet. Click OK. 	<ol style="list-style-type: none"> Enter set unit 0 Enter set dlci 10 Enter set encapsulation frame-relay Enter set peer-unit 1 Enter set family inet
Configure the second logical unit on the It-0/0/0 interface.	<ol style="list-style-type: none"> Next to Unit, click Add new entry. In the Interface unit number box, type 1. In the Dlci box, type 10. From the Encapsulation list, select frame-relay. In the Peer unit box, type 0. Under Family, select Inet. Click OK until you return to the Configuration page. 	<ol style="list-style-type: none"> Enter set unit 1 Enter set dlci 10 Enter set encapsulation frame-relay Enter set peer-unit 0 Enter set family inet

Table 122: Configuring DLSw on the Remote Router (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the default CoS classifier on the It-0/0/0 interface.	<ol style="list-style-type: none"> Next to Class of service, click Edit. Next to Interfaces, click Add new entry. In the Interface name box, type It-0/0/0. Next to Unit, click Add new entry. In the Unit number box, type 1. Next to Classifiers, click Configure. Under Dscp, in the Classifier name box, type default. Click OK until you return to the Configuration page. 	<p>Enter</p> <p>edit class-of-service interfaces It-0/0/0 unit 1</p> <p>Enter</p> <p>set classifiers dscp default</p>
Configure the type-of-service precedence value for DLSw packets—for example, 192.	<ol style="list-style-type: none"> Click Protocols. Next to DlsW, click Edit. Next to Sna cos, click Configure. In the Destination sna interface box, type It-0/0/0.0. In the Type of service box, type 192. Click OK. 	<ol style="list-style-type: none"> Enter edit protocols dlsW sna-cos Enter set destination-sna-interface It-0/0/0.0 type-of-service 192

Verifying DLSw Configuration

To verify DLSw configuration, perform these tasks:

- Displaying LLC2 Properties on a Fast Ethernet Interface on page 427
- Displaying DLSw Capabilities on page 427
- Displaying DLSw Circuit State on page 427
- Displaying Details of a DLSw Circuit State on page 428
- Displaying DLSw Peers on page 428
- Displaying Details of DLSw Peers on page 428
- Displaying DLSw Reachability Information on page 429

Displaying LLC2 Properties on a Fast Ethernet Interface

Purpose	Verify the configuration of LLC2 properties on a Fast Ethernet interface.
Action	From the CLI, enter the show interfaces fe-3/0/0 command.
Sample Output	<pre> user@host# show interfaces fe-3/0/0 fe-3/0/0 { unit 0 { family inet{ address 172.5.20.1/24; } family llc2} } } </pre>
What It Means	Verify that the output shows the intended LLC2 configuration. For more information about the format of a configuration file, see “Viewing the Configuration Text” on page 12.

Displaying DLSw Capabilities

Purpose	Verify DLSw capabilities of remote DLSw peers.
Action	From the CLI, enter the show dlsw capabilities command.
Sample Output	<pre> user@host> show dlsw capabilities Peer: 50.50.50.50 Vendor ID :000585 Version number :0200 Initial pacing window size :32 Version String Juniper Networks, Inc. j2300 internet router JUNOS Software Release 7.4I0 [builder] Build date: 2005-07-15 07:13:17 UTC Copyright (c) 1996-2005 Juniper Networks, Inc. Compiled Wed 26-Jan-05 02:49 by pwade </pre>
What It Means	Verify that the output displays the capabilities of remote DLSw peers.

Displaying DLSw Circuit State

Purpose	Display DLSw circuits currently established after configuration in “Configuring Basic DLSw with a Configuration Editor” on page 420.
Action	From the CLI, enter the show dlsw circuits command.
Sample Output	<pre> user@host> show dlsw circuits Local Remote Address LSAP Address DSAP State Peer 00:40:cd:92:4b:7b 04 00:40:cd:92:4b:51 04 CONNECTED 50.50.50.50 </pre>

Displaying Details of a DLSw Circuit State

Purpose	Display the details of DLSw circuits currently established after configuration in “Configuring Basic DLSw with a Configuration Editor” on page 420.
Action	From the CLI, enter the <code>show dlsw circuits detail</code> command.
Sample Output	<pre> user@host> show dlsw circuits detail Circuit id:2240022c00 local addr:00:40:cd:92:4b:7b lsap:04 remote addr:00:40:cd:92:4b:51 dsap:04 remote peer address: 50.50.50.50 circuit state CONNECTED created time 200238 max btu size 1033 circuit priority 2 </pre>
What It Means	In addition to the local and remote MAC addresses, the created time of the circuit as well as the priority and maximum basic transmission unit (BTU) size are displayed.

Displaying DLSw Peers

Purpose	Display information about the DLSw peers on the network.
Action	From the CLI, enter the <code>show dlsw peers</code> command.
Sample Output	<pre> user@host> show dlsw peers Peer State Circuits 50.50.50.50 Connected 1 10.10.10.10 Connected 1 </pre>
What It Means	The output displays the number of active or inactive DLSw peers.

Displaying Details of DLSw Peers

Purpose	Display detailed information about DLSw peers on a network.
Action	From the CLI, enter the <code>show dlsw peers detail</code> command.
Sample Output	<pre> user@host> show dlsw peers detail Peer: 50.50.50.50 State: Connected, Circuits: 7, Local address:10.10.10 Created time: 21977, Connected time: 4059 Receive initial pacing: 20, No circuits timeout: 300 Type-of-service value: 192 Statistics: Packets received: :0 Packets sent: :752 Bytes received :0 Bytes sent :10 CANUREACH_ex received :0 CANUREACH_ex sent :6 ICANREACH_ex received :6 ICANREACH_ex sent :0 </pre>

- What It Means** The output displays the DLSw peer state and the following statistics:
- Packets received—Number of packets received from DLSw peers
 - Packets sent—Number of packets sent to the DLSw peers
 - Bytes received—Number of bytes received from DLSw peers
 - Bytes sent—Number of bytes sent to the DLSw peers
 - CANUREACH_ex received—Number of exploratory messages received from remote DLSw peers
 - CANUREACH_ex sent—Number of exploratory messages sent to remote DLSw peers
 - ICANREACH_ex received—Number of confirmation messages received from remote DLSw peers
 - ICANREACH_ex sent—Number of confirmation messages sent to remote DLSw peers

Displaying DLSw Reachability Information

Purpose Display information about the MAC cache entries and peer IP addresses currently maintained on the DSLw router.

Action From the CLI, enter the `show dlsw reachability` command.

Sample Output `user@host> show dlsw reachability`

```
MAC index MAC address      Remote DLSw address
  1 22:22:00:00:00:03 3.3.3.1
  2 22:22:00:00:00:04 3.3.3.1
  3 22:22:00:00:00:05 3.3.3.1
  4 22:22:00:00:00:06 3.3.3.1
  5 22:22:00:00:00:07 3.3.3.1
  6 22:22:00:00:00:08 3.3.3.1
  7 22:22:00:00:00:09 3.3.3.1
  8 22:22:00:00:00:0a 3.3.3.1
  9 22:22:00:00:00:0b 3.3.3.1
 10 22:22:00:00:00:0c 3.3.3.1
 11 00:0c:f1:e8:4e:ad 3.3.3.1
 12 44:44:00:00:00:01 3.3.3.1
 13 44:44:00:00:00:02 3.3.3.1
 14 44:44:00:00:00:03 3.3.3.1
 15 44:44:00:00:00:04 3.3.3.1
```


Part 7

Configuring Routing Policy, Firewall Filters, and Class of Service

- Policy, Firewall Filter, and Class-of-Service Overview on page 433
- Configuring Routing Policies on page 459
- Configuring Firewall Filters and NAT on page 475
- Configuring Class of Service with DiffServ on page 529

Chapter 20

Policy, Firewall Filter, and Class-of-Service Overview

Several mechanisms can help you control the way routing information and data packets are handled by a router—routing policy, firewall filters, and class-of-service (CoS) rules. Routing policies control how information is imported to and exported from the routing tables, acting exclusively at the Routing Engine level. Firewall filters examine packets at the entry (ingress) and exit (egress) points of the Services Router, filtering traffic at the router level. CoS rules determine packet scheduling, buffering, and queueing within the router. These three mechanisms are at the core of managing how a router forwards traffic.



NOTE: You must have a license to configure a stateful firewall filter and Network Address Translation (NAT). For license details, see the *J-series Services Router Administration Guide*.

To manage the flow of information into and out of a Services Router, you must understand the fundamentals of routing policies, firewall filters, and CoS rules. To read this chapter, you need a basic understanding of IP routing protocols.

This chapter contains the following topics. For more information see the *JUNOS Policy Framework Configuration Guide* and the *JUNOS Class of Service Configuration Guide*.

- Policy, Firewall Filter, and CoS Terms on page 433
- Routing Policy Overview on page 435
- Firewall Filter Overview on page 440
- Class-of-Service Overview on page 449

Policy, Firewall Filter, and CoS Terms

Before configuring routing policies, firewall filters, or class of service (CoS) with Differentiated Services (DiffServ) on a Services Router, become familiar with the terms defined in Table 123.

Table 123: Policy, Firewall Filter, and CoS Terms

Term	Definition
assured forwarding (AF)	CoS packet forwarding class that provides a group of values you can define and includes four subclasses, AF1, AF2, AF3, and AF4, each with three drop probabilities, low, medium, and high.
behavior aggregate (BA) classifier	Feature that can be used to determine the forwarding treatment for each packet. The BA classifier maps a code point to a loss priority. The loss priority is used later in the work flow to select one of the two drop profiles used by random early detection (RED).
best-effort (BE)	CoS packet forwarding class that provides no service profile. For the BE forwarding class, loss priority is typically not carried in a code point, and random early detection (RED) drop profiles are more aggressive.
class of service (CoS)	Method of classifying traffic on a packet-by-packet basis, using information in the type-of-service (ToS) byte to assign traffic flows to different service levels.
Differentiated Services (DiffServ)	Services based on RFC 2474, <i>Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers</i> . The DiffServ method of CoS uses the type-of-service (ToS) byte to identify different packet flows on a packet-by-packet basis. DiffServ adds a Class Selector code point (CSCP) and a DiffServ code point (DSCP).
DiffServ code point (DSCP)	Values for a 6-bit field defined in IP packet headers that can be used to enforce class-of-service (CoS) distinctions in a Services Router.
drop profile	Drop probabilities for different levels of buffer fullness that are used by random early detection (RED) to determine from which Services Router scheduling queue to drop packets.
expedited forwarding (EF)	CoS packet forwarding class that provides end-to-end service with low loss, low latency, low jitter, and assured bandwidth.
firewall filter	See <i>stateful firewall filter</i> ; <i>stateless firewall filter</i> .
multifield (MF) classifier	Firewall filter that scans through a variety of packet fields to determine the forwarding class and loss priority for a packet and polices traffic to a specific bandwidth and burst size. Typically, a classifier performs matching operations on the selected fields against a configured value.
network address port translation (NAPT)	Method of concealing a set of host ports on a private network behind a pool of public addresses. It can be used as a security measure to protect the host ports from direct targeting in network attacks.
Network Address Translation (NAT)	Method of concealing a set of host addresses on a private network behind a pool of public addresses. It can be used as a security measure to protect the host addresses from direct targeting in network attacks.
network control (NC)	CoS packet forwarding class that is typically high priority because it supports protocol control.
PLP bit	Packet loss priority bit. Used to identify packets that have experienced congestion or are from a transmission that exceeded a service provider's customer service license agreement. A Services Router can use the PLP bit as part of a congestion control strategy. The bit can be configured on an interface or in a filter.
policer	Feature that limits the amount of traffic passing into or out of an interface. It is an essential component of firewall filters that is designed to thwart denial-of-service (DoS) attacks. A policer applies rate limits on bandwidth and burst size for traffic on a particular Service Router interface.
policing	Applying rate and burst size limits to traffic on an interface.
random early detection (RED)	Gradual drop profile for a given class, used for congestion avoidance. RED attempts to anticipate congestion and reacts by dropping a small percentage of packets from the head of a queue to prevent congestion.
rule	Guide that the Services Router follows when applying services. A rule consists of a match direction and one or more terms.

Table 123: Policy, Firewall Filter, and CoS Terms (continued)

Term	Definition
service set	Collection of services. Examples of services include stateful firewall filters and Network Address Translation (NAT).
stateful firewall filter	Type of firewall filter that evaluates the context of connections, permits or denies traffic based on the context, and updates this information dynamically. Context includes IP source and destination addresses, TCP port numbers, TCP sequencing information, and TCP connection flags.
stateless firewall filter	Type of firewall filter that statically evaluates the contents of packets transiting the router, and packets originating from, or destined for, the router. Information about connection states is not maintained.
term	Firewall filters contain one or more terms that specify filter match conditions and actions.
trusted network	Network from which all originating traffic can be trusted—for example, an internal enterprise LAN. Stateful firewall filters allow traffic to flow from trusted to untrusted networks.
untrusted network	Network from which all originating traffic cannot be trusted—for example, a WAN. Unless configured otherwise, stateful firewall filters do not allow traffic to flow from untrusted to trusted networks.

Routing Policy Overview

Routing protocols send information about routes to a router's neighbors. This information is processed and used to create routing tables, which are then distilled into forwarding tables. Routing policies control the flow of information between the routing protocols and the routing tables and between the routing tables and the forwarding tables. Using policies, you can determine which routes are advertised, specify which routes are imported into the routing table, and modify routes to control which routes are added to the forwarding table.

This overview contains the following topics:

- Routing Policy Components on page 435
- Applying Routing Policies on page 440

Routing Policy Components

Routing policies are made up of one or more terms, which contain a set of match conditions and a set of actions. Match conditions are criteria that a route must match before the actions can be applied. If a route matches all criteria, one or more actions are applied to the route. These actions specify whether to accept or reject the route, control how a series of policies are evaluated, and manipulate the characteristics associated with a route.

This section contains the following topics:

- “Routing Policy Terms” on page 436
- “Routing Policy Match Conditions” on page 436
- “Routing Policy Actions” on page 438
- “Default and Final Actions” on page 440

Routing Policy Terms

A term is a named structure in which match conditions and actions are defined. Each routing policy contains one or more terms.

Generally, a Services Router compares a route against the match conditions of each term in a routing policy, starting with the first and moving through the terms in the order in which they are defined, until a match is made and an explicitly configured or default action of `accept` or `reject` is taken. If none of the terms in the policy match the route, the Services Router compares the route against the next policy, and so on, until either an action is taken or the default policy is evaluated.

Routing Policy Match Conditions

A match condition defines the criteria that a route must match for an action to take place. Each term can have one or more match conditions. If a route matches all the match conditions for a particular term, the actions defined for that term are processed.

Each term can consist of two statements, `to` and `from`, that define match conditions:

- In the `from` statement, you define the criteria that an *incoming* route must match. You can specify one or more match conditions. If you specify more than one, all conditions must match the route for a match to occur.
- In the `to` statement, you define the criteria that an *outgoing* route must match. You can specify one or more match conditions. If you specify more than one, all conditions must match the route for a match to occur.

The order of match conditions in a term is not important, because a route must match all match conditions in a term for an action to be taken.

Table 124 summarizes the routing policy match conditions.

Table 124: Summary of Routing Policy Match Conditions

Match Condition	Description
aggregate-contributor	Matches routes that are contributing to a configured aggregate. This match condition can be used to suppress a contributor in an aggregate route.

Table 124: Summary of Routing Policy Match Conditions (continued)

Match Condition	Description
area <i>area-id</i>	Matches a route learned from the specified OSPF area during the exporting of OSPF routes into other protocols.
as-path <i>name</i>	Name of an AS path regular expression. BGP routes whose AS path matches the regular expression are processed.
color <i>preference</i>	Color value. You can specify preference values that are finer-grained than those specified in the <i>preference</i> match conditions. The <i>color</i> value can be a number from 0 through 4,294,967,295 ($2^{32} - 1$). A lower number indicates a more preferred route.
community	Name of one or more communities. If you list more than one name, only one name needs to match for a match to occur. (The matching is effectively a logical OR operation.)
external [type <i>metric-type</i>]	Matches external OSPF routes, including routes exported from one level to another. In this construct <i>type</i> is an optional keyword. The <i>metric-type</i> value can be either 1 or 2. When you do not specify <i>type</i> , this condition matches all external routes.
interface <i>interface-name</i>	Name or IP address of one or more router interfaces. Do not use this qualifier with protocols that are not interface-specific, such as internal BGP (IBGP). Depending on where the policy is applied, this match condition matches routes learned from or advertised through the specified interface.
internal	Matches a routing policy against the internal flag for simplified next-hop self policies.
level <i>level</i>	Matches the IS-IS level. Routes that are from the specified level or are being advertised to the specified level are processed.
local-preference <i>value</i>	BGP local preference attribute. The preference value can be from 0 through 4,294,967,295 ($2^{32} - 1$).
metric <i>metric</i> metric2 <i>metric</i>	Metric value. The <i>metric</i> value corresponds to the multiple exit discriminator (MED), and <i>metric2</i> corresponds to the interior gateway protocol (IGP) metric if the BGP next hop runs back through another route.
neighbor <i>address</i>	Address of one or more neighbors (peers). For BGP export policies, the address can be a directly connected or indirectly connected peer. For all other protocols, the address is the neighbor from which the advertisement is received.
next-hop <i>address</i>	Next-hop address or addresses specified in the routing information for a particular route. For BGP routes, matches are performed against each protocol next hop.
origin <i>value</i>	BGP origin attribute, which is the origin of the AS path information. The value can be one of the following: <ul style="list-style-type: none"> ■ egp—Path information originated from another AS. ■ igp—Path information originated from within the local AS. ■ incomplete—Path information was learned by some other means.
policy [<i>policy-names</i>]	Name of one or more policies to evaluate as a subroutine.
preference <i>preference</i> preference2 <i>preference</i>	Preference value. You can specify a primary preference value (<i>preference</i>) and a secondary preference value (<i>preference2</i>). The preference value can be a number from 0 through 4,294,967,295 ($2^{32} - 1$). A lower number indicates a more preferred route.

Table 124: Summary of Routing Policy Match Conditions (continued)

Match Condition	Description
<code>prefix-list name</code>	Named list of IP addresses configured at the Policy-options level in the configuration hierarchy. This match condition can be used on import policies only.
<code>protocol protocol</code>	Name of the protocol from which the route was learned or to which the route is being advertised. It can be one of the following: aggregate , bgp , direct , dvmrp , isis , local , ospf , pim-dense , pim-sparse , rip , ripng , or static .
<code>route-filter destination-prefix match-type <actions></code>	List of destination prefixes. When specifying a destination prefix, you can specify an exact match with a specific route or a less precise match using match types. You can configure either a common action that applies to the entire list or an action associated with each prefix. Route filters can be used on import policies only.
<code>route-type value</code>	Type of route. The value can be either external or internal .
<code>source-address-filter destination-prefix match-type <actions></code>	List of multicast source addresses. When specifying a source address, you can specify an exact match with a specific route or a less precise match using match types. You can configure either a common action that applies to the entire list or an action associated with each prefix. Source-address filters can be used on import policies only.

Routing Policy Actions

An action defines what the Services Router does with the route when the route matches all the match conditions in the **from** and **to** statements for a particular term. If a term does not have **from** and **to** statements, all routes are considered to match and the actions apply to all routes.

Each term can have one or more of the following types of actions. The actions are configured under the **then** statement.

- Flow control actions, which affect whether to accept or reject the route and whether to evaluate the next term or routing policy
- Actions that manipulate route characteristics
- Trace action, which logs route matches

Table 125 summarizes the routing policy actions.

If you do not specify an action, one of the following results occurs:

- The next term in the routing policy, if one exists, is evaluated.
- If the routing policy has no more terms, the next routing policy, if one exists, is evaluated.
- If there are no more terms or routing policies, the accept or reject action specified by the default policy is executed.

Table 125: Summary of Key Routing Policy Actions

Action	Description
Flow Control Actions	
accept	Accepts the route and propagates it. After a route is accepted, no other terms in the routing policy and no other routing policies are evaluated.
reject	Rejects the route and does not propagate it. After a route is rejected, no other terms in the routing policy and no other routing policies are evaluated.
next term	Skips to and evaluates the next term in the same routing policy. Any accept or reject action specified in the then statement is ignored. Any actions specified in the then statement that manipulate route characteristics are applied to the route.
next policy	Skips to and evaluates the next routing policy. Any accept or reject action specified in the then statement is ignored. Any actions specified in the then statement that manipulate route characteristics are applied to the route.
Route Manipulation Actions	
as-path-prepend <i>as-path</i>	<p>Appends one or more autonomous system (AS) numbers at the beginning of the AS path. If you are specifying more than one AS number, include the numbers in quotation marks.</p> <p>The AS numbers are added after the local AS number has been added to the path. This action adds AS numbers to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the appended AS numbers are placed within a confederation sequence. Otherwise, the appended AS numbers are placed with a nonconfederation sequence.</p>
as-path-expand last-as count <i>n</i>	<p>Extracts the last AS number in the existing AS path and appends that AS number to the beginning of the AS path <i>n</i> times. Replace <i>n</i> with a number from 1 through 32.</p> <p>The AS numbers are added after the local AS number has been added to the path. This action adds AS numbers to AS sequences only, not to AS sets. If the existing AS path begins with a confederation sequence or set, the appended AS numbers are placed within a confederation sequence. Otherwise, the appended AS numbers are placed with a nonconfederation sequence.</p>
class <i>class-name</i>	Applies the specified class-of-service (CoS) parameters to routes installed into the routing table.
color <i>preference</i> color2 <i>preference</i>	Sets the preference value to the specified value. The color and color2 preference values can be a number from 0 through 4,294,967,295 ($2^{32} - 1$). A lower number indicates a more preferred route.
damping <i>name</i>	<p>Applies the specified route-damping parameters to the route. These parameters override BGP's default damping parameters.</p> <p>This action is useful only in import policies.</p>

Table 125: Summary of Key Routing Policy Actions (continued)

Action	Description
local-preference <i>value</i>	Sets the BGP local preference attribute. The preference can be a number from 0 through 4,294,967,295 ($2^{32} - 1$).
metric <i>metric</i>	Sets the metric. You can specify up to four metric values, starting with metric (for the first metric value) and continuing with metric2 , metric3 , and metric4 . For BGP routes, metric corresponds to the MED, and metric2 corresponds to the IGP metric if the BGP next hop loops through another router.
metric2 <i>metric</i>	
metric3 <i>metric</i>	
metric4 <i>metric</i>	
next-hop <i>address</i>	Sets the next hop. If you specify address as self , the next-hop address is replaced by one of the local router's addresses. The advertising protocol determines which address to use.

Default and Final Actions

If none of the terms' match conditions evaluate to true, the final action is executed. The final action is defined in an unnamed term. Additionally, you can define a default action (either **accept** or **reject**) that overrides any action intrinsic to the protocol.

Applying Routing Policies

Once a policy is created, it must be applied before it is active. You apply routing policies using the **import** and **export** statements at the **Protocols > protocol-name** level in the configuration hierarchy.

In the **import** statement, list the name of the routing policy to be evaluated when routes are imported into the routing table from the routing protocol.

In the **export** statement, list the name of the routing policy to be evaluated when routes are being exported from the routing table into a dynamic routing protocol. Only active routes are exported from the routing table.

To specify more than one policy and create a policy chain, list the policies using a space as a separator. If multiple policies are specified, the policies are evaluated in the order in which they are specified. As soon as an **accept** or **reject** action is executed, the policy chain evaluation ends.

Firewall Filter Overview



NOTE: You must have a license to configure a stateful firewall filter and Network Address Translation (NAT). For license details, see the *J-series Services Router Administration Guide*.

This section contains the following topics:

- Stateful and Stateless Firewall Filters on page 441
- Process for Configuring a Stateful Firewall Filter and NAT on page 442
- Summary of Stateful Firewall Filter and NAT Match Conditions and Actions on page 442
- Planning a Stateless Firewall Filter on page 444
- Stateless Firewall Filter Match Conditions, Actions, and Action Modifiers on page 445

Stateful and Stateless Firewall Filters

In a *stateful* firewall filter, all packets flowing from a trusted network to an untrusted network are allowed. Packets flowing from an untrusted network to a trusted network are allowed only if they are responses to a session originated by the trusted network, or if they are explicitly accepted by a term in the stateful firewall filter rule.



CAUTION: If a packet does not match any terms in a firewall filter rule, the packet is discarded. Take care you do not configure a firewall filter that prevents you from accessing the Services Router after you commit the configuration. For example, if you configure a firewall filter that does not match HTTP or HTTPS packets, you cannot access the router with the J-Web interface.

When Network Address Translation (NAT) is enabled, the source address of a packet flowing from a trusted network to an untrusted network is replaced with an address chosen from a specified range, or *pool*, of addresses. In addition, you can configure the Services Router to dynamically translate the source port of the packet—a process called Network Address Port Translation (NAPT).

A *stateless* firewall filter can filter packets transiting the Services Router from a source to a destination, or packets originating from, or destined for, the Routing Engine. Stateless firewall filters applied to the Routing Engine interface protect the processes and resources owned by the Routing Engine.

You can apply a stateless firewall filter to an input or output interface, or to both. Every packet, including fragmented packets, is evaluated against stateless firewall filters.

All stateful and stateless firewall filters contain one or more terms, and each term consists of two components—match conditions and actions. The match conditions define the values or fields that the packet must contain to be considered a match. If a packet is a match, the corresponding action is taken. By default, a packet that does not match a firewall filter is discarded.



NOTE: A firewall filter with a large number of terms can adversely affect both the configuration commit time and the performance of the Routing Engine.

For more information about firewall filters, see “Configuring IPSec for Secure Packet Exchange” on page 379 and the *JUNOS Policy Framework Configuration Guide*. For more information about NAT, see the *JUNOS Services Interfaces Configuration Guide*.

Process for Configuring a Stateful Firewall Filter and NAT

To configure a stateful firewall filter and NAT, perform the following tasks:

- Define the stateful firewall filter output and input rules. You must define an output rule that allows all traffic (application and nonapplication) to flow from the trusted network to the untrusted network.

To define the match condition in the term that allows application traffic to flow from the trusted network to the untrusted network, we recommend you specify the JUNOS default group `junos-algs-outbound` as the application set. To view the configuration of this group, enter the `show groups junos-defaults applications application-set junos-algs-outbound` configuration mode command. For more information about JUNOS default groups, see the *JUNOS System Basics Configuration Guide*.

You also must define an input rule to discard all traffic from the untrusted network that is not a response to a session originated by the trusted network.

- Define the NAT address and port pool.
- Define the NAT output and input rules.
- Define a service set that includes all stateful firewall filter and NAT rules and the service interface. You must specify the service interface as `sp-0/0/0`. This service interface is a virtual interface that must be included at the [edit interfaces] hierarchy level to support stateful firewall filter and NAT services.
- Apply the service set to the interfaces that make up the untrusted network.



NOTE: Do not apply the service set to the `sp-0/0/0` interface.

For more information about match conditions and actions, see “Summary of Stateful Firewall Filter and NAT Match Conditions and Actions” on page 442.

Summary of Stateful Firewall Filter and NAT Match Conditions and Actions

Table 126 lists the match conditions you can specify in stateful firewall filter and NAT terms. Table 127 and Table 128 list actions you can specify in stateful firewall filter and NAT terms.

Table 126: Stateful Firewall Filter and NAT Match Conditions

Match Condition	Description
application-sets [<i>set-names</i>]	List of application set names. Application sets are defined at the [edit applications] hierarchy level.
applications [<i>application-names</i>]	List of applications. Applications are defined at the [edit applications] hierarchy level.
destination-address <i>address</i>	IP destination address field.
source-address <i>address</i>	IP source address field.

For more information about configuring applications and application sets for stateful firewall filters, see the *JUNOS Services Interfaces Configuration Guide*.

Table 127: Stateful Firewall Filter Actions

Actions	Description
accept	Accept the packet and send it to its destination.
allow-ip-options [<i>values</i>]	If the IP Option header of the packet contains a value that matches one of the specified values, accept the packet. If this action is not included, only packets without IP options are accepted. This action can be specified only with the accept action. You can specify the IP option as text or a numeric value: any (0), ip-security (130), ip-stream (8), loose-source-route (3), route-record (7), router-alert (148), strict-source-route (9), and timestamp (4).
discard	Do not accept the packet, and do not process it further.
reject	Do not accept the packet, and send a rejection message. UDP sends an ICMP unreachable code and RCP sends RST. Rejected packets can be logged or sampled.
syslog	Record information in the system logging facility. This action can be used with all options except discard .

Table 128: NAT Actions

Actions	Description
syslog	Record information in the system logging facility.
translated destination-pool <i>nat-pool-name</i>	Translate the destination address using the specified pool.
translated source-pool <i>nat-pool-name</i>	Translate the source address using the specified pool.

Table 128: NAT Actions (continued)

Actions	Description
translation-type (destination type source type)	<p>Translate the destination and source port using the specified type:</p> <ul style="list-style-type: none"> ■ destination static—Translate the destination address without port mapping. This type requires the size of the source address space to be the same as the size of the destination address space. You must specify a destination-pool name. The referenced pool must contain exactly one address and no port configuration at the [edit nat pool] hierarchy level. ■ source dynamic—Translate the source address with port mapping by means of NAT. You must specify a source-pool name. The referenced pool must include a port configuration at the [edit nat pool] hierarchy level. ■ source static—Translate the source address without port mapping. This type requires the size of the source address space to be the same as the size of the destination address space. You must specify a source-pool name. The referenced pool must contain exactly one address and no port configuration at the [edit nat pool] hierarchy level.
syslog	Information is recorded in the system logging facility.

Planning a Stateless Firewall Filter

Before creating a stateless firewall filter and applying it to an interface, determine what you want the firewall filter to accomplish and how to use its match conditions and actions to achieve your goal. Also, make sure you understand how packets are matched and the default action of the resulting firewall filter.



CAUTION: If a packet does not match any terms in a stateless firewall filter rule, the packet is discarded. Take care that you do not configure a firewall filter that prevents you from accessing the Services Router after you commit the configuration. For example, if you configure a firewall filter that does not match HTTP or HTTPS packets, you cannot access the router with the J-Web interface.

To configure a stateless firewall filter, determine the following:

- Purpose of the firewall filter—for example, to limit traffic to certain protocols, IP source or destination addresses, or data rates, or to prevent denial-of-service (DoS) attacks.
- Appropriate match conditions. The packet header fields to match—for example, IP header fields (such as source and destination IP addresses,

protocols, and IP options), TCP header fields (such as source and destination ports and flags), and ICMP header fields (such as ICMP packet type and code).

- Action to take if a match occurs—for example, accept, discard, or evaluate the next term.
- (Optional) Action modifiers. Additional actions to take if a packet matches—for example, count, log, rate limit, or police a packet.
- Interface on which the firewall filter is applied. The input or output side, or both sides, of the Routing Engine interface or a non-Routing Engine interface.

For more information about what a stateless firewall filter can include, see “Stateless Firewall Filter Match Conditions, Actions, and Action Modifiers” on page 445. For more information about stateless firewall filters, see the *JUNOS Policy Framework Configuration Guide*.

Stateless Firewall Filter Match Conditions, Actions, and Action Modifiers

Table 129 lists the match conditions you can specify in stateless firewall filter terms. Some of the numeric range and bit-field match conditions allow you to specify a text synonym. For a complete list of the synonyms, do any of the following:

- If you are using the J-Web interface, select the synonym from the appropriate list.
- If you are using the CLI, type a question mark (?) after the from statement.
- See the *JUNOS Policy Framework Configuration Guide*.

To specify a bit-field match condition with values, such as `tcp-flags`, you must enclose the values in quotation marks (“ ”). You can use bit-field logical operators to create expressions that are evaluated for matches. For example, if the following expression is used in a filter term, a match occurs if the packet is the initial packet of a TCP session:

`tcp-flags “syn & !ack”`

Table 130 lists the bit-field logical operators in order of highest to lowest precedence.

You can use text synonyms to specify some common bit-field matches. In the previous example, you can specify `tcp-initial` to specify the same match condition.



NOTE: When the Services Router compares the stateless firewall filter match conditions to a packet, it compares only the header fields specified in the match condition. There is no implied protocol match. For example, if you specify a match of `destination-port ssh`, the Services Router checks for a value of 0x22 in the 2-byte field that is two bytes after the IP packet header. The protocol field of the packet is not checked.

Table 129: Stateless Firewall Filter Match Conditions

Match Condition	Description
Numeric Range Match Conditions	
<i>keyword-except</i>	<p>Negates a match. For example, destination-port-except <i>number</i> .</p> <p>The following keywords accept the -except extension: destination-port, dscp, esp-spi, forwarding-class, fragment-offset, icmp-code, icmp-type, interface-group, ip-options, packet-length, port, precedence, protocol and source-port.</p>
<i>destination-port number</i>	<p>TCP or User Datagram Protocol (UDP) destination port field. You cannot specify both the port and destination-port match conditions in the same term. Normally, you specify this match in conjunction with the protocol tcp or protocol udp match statement to determine which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify a text synonym. For example, you can specify telnet or 23.</p>
<i>esp-spi spi-value</i>	<p>IPSec encapsulating security payload (ESP) security parameter index (SPI) value. Match on this specific SPI value. You can specify the ESP SPI value in either hexadecimal, binary, or decimal form.</p>
<i>forwarding-class class</i>	<p>Forwarding class. Specify assured-forwarding, best-effort, expedited-forwarding, or network-control.</p>
<i>fragment-offset number</i>	<p>Fragment offset field.</p>
<i>icmp-code number</i>	<p>ICMP code field. Normally, you specify this match in conjunction with the protocol icmp match statement to determine which protocol is being used on the port.</p> <p>This value or keyword provides more specific information than icmp-type. Because the value's meaning depends on the associated icmp-type, you must specify icmp-type along with icmp-code.</p> <p>In place of the numeric value, you can specify a text synonym. For example, you can specify ip-header-bad or 0.</p>
<i>icmp-type number</i>	<p>ICMP packet type field. Normally, you specify this match in conjunction with the protocol icmp match statement to determine which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify a text synonym. For example, you can specify time-exceeded or 11.</p>
<i>interface-group group-number</i>	<p>Interface group on which the packet was received. An interface group is a set of one or more logical interfaces. For information about configuration interface groups, see the <i>JUNOS Policy Framework Configuration Guide</i>.</p>
<i>packet-length bytes</i>	<p>Length of the received packet, in bytes. The length refers only to the IP packet, including the packet header, and does not include any Layer 2 encapsulation overhead.</p>
<i>port number</i>	<p>TCP or UDP source or destination port field. You cannot specify both the port match and either the destination-port or source-port match conditions in the same term. Normally, you specify this match in conjunction with the protocol tcp or protocol udp match statement to determine which protocol is being used on the port.</p> <p>In place of the numeric value, you can specify a text synonym. For example, you can specify bgp or 179.</p>
<i>precedence ip-precedence-field</i>	<p>IP precedence field. You can specify precedence in either hexadecimal, binary, or decimal form.</p> <p>In place of the numeric value, you can specify a text synonym. For example, you can specify immediate or 0x40.</p>

Table 129: Stateless Firewall Filter Match Conditions (continued)

Match Condition	Description
protocol <i>number</i>	IP protocol field. In place of the numeric value, you can specify a text synonym. For example, you can specify ospf or 89 .
source-port <i>number</i>	TCP or UDP source port field. You cannot specify the port and source-port match conditions in the same term. Normally, you specify this match in conjunction with the protocol tcp or protocol udp match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify a text synonym. For example, you can specify http or 80 .
Address Match Conditions	
address <i>prefix</i>	IP source or destination address field. You cannot specify both the address and the destination-address or source-address match conditions in the same term.
destination-address <i>prefix</i>	IP destination address field. You cannot specify the destination-address and address match conditions in the same term.
destination-prefix-list <i>prefix-list</i>	IP destination prefix list field. You cannot specify the destination-prefix-list and prefix-list match conditions in the same term.
prefix-list <i>prefix-list</i>	IP source or destination prefix list field. You cannot specify both the prefix-list and the destination-prefix-list or source-prefix-list match conditions in the same term.
source-address <i>prefix</i>	IP source address field. You cannot specify the source-address and address match conditions in the same rule.
source-prefix-list <i>prefix-list</i>	IP source prefix list field. You cannot specify the source-prefix-list and prefix-list match conditions in the same term.
Bit-Field Match Conditions with Values	
fragment-flags <i>number</i>	IP fragmentation flags. In place of the numeric value, you can specify a text synonym. For example, you can specify more-fragments or 0x2000 .
ip-options <i>number</i>	IP options. In place of the numeric value, you can specify a text synonym. For example, you can specify record-route or 7 .
tcp-flags <i>number</i>	TCP flags. Normally, you specify this match in conjunction with the protocol tcp match statement to determine which protocol is being used on the port. In place of the numeric value, you can specify a text synonym. For example, you can specify syn or 0x02 .
Bit-Field Text Synonym Match Conditions	
first-fragment	First fragment of a fragmented packet. This condition does not match unfragmented packets.
is-fragment	This condition matches if the packet is a trailing fragment. It does not match the first fragment of a fragmented packet. To match both first and trailing fragments, you can use two terms, or you can use fragment-offset 0-8191 .

Table 129: Stateless Firewall Filter Match Conditions (continued)

Match Condition	Description
tcp-established	TCP packets other than the first packet of a connection. This match condition is a synonym for "(ack rst)". This condition does not implicitly check that the protocol is TCP. To do so, specify the protocol tcp match condition.
tcp-initial	First TCP packet of a connection. This match condition is a synonym for "(syn & !ack)". This condition does not implicitly check that the protocol is TCP. To do so, specify the protocol tcp match condition.

Table 130: Stateless Firewall Filter Bit-Field Logical Operators

Logical Operator	Description
(...)	Grouping
!	Negation
& or +	Logical AND
or ,	Logical OR

Table 131 lists the actions and action modifiers you can specify in stateless firewall filter terms.

Table 131: Stateless Firewall Filter Actions and Action Modifiers

Action or Action Modifier	Description
accept	Accepts a packet. This is the default if the packet matches. However, we strongly recommend that you always explicitly configure an action in the then statement.
discard	Discards a packet silently, without sending an Internet Control Message Protocol (ICMP) message. Packets are available for logging and sampling before being discarded.
next term	Continues to the next term for evaluation.
reject <message-type>	Discards a packet, sending an ICMP destination unreachable message. Rejected packets are available for logging and sampling. You can specify one of the following message types: administratively-prohibited (default), bad-host-tos , bad-network-tos , host-prohibited , host-unknown , host-unreachable , network-prohibited , network-unknown , network-unreachable , port-unreachable , precedence-cutoff , precedence-violation , protocol-unreachable , source-host-isolated , source-route-failed , or tcp-reset . If you specify tcp-reset , a TCP reset is returned if the packet is a TCP packet. Otherwise, nothing is returned.
routing-instance <i>routing-instance</i>	Routes the packet using the specified routing instance.
Action Modifiers	

Table 131: Stateless Firewall Filter Actions and Action Modifiers (continued)

Action or Action Modifier	Description
count <i>counter-name</i>	Counts the number of packets passing this term. The name can contain letters, numbers, and hyphens (-), and can be up to 24 characters long. A counter name is specific to the filter that uses it, so all interfaces that use the same filter increment the same counter.
forwarding-class <i>class-name</i>	Classifies the packet to the specified forwarding class.
log	Logs the packet's header information in the Routing Engine. You can access this information by entering the show firewall log command at the CLI.
loss-priority <i>priority</i>	Sets the scheduling priority of the packet. The priority can be low or high .
policer <i>policer-name</i>	Applies rate limits to the traffic using the named policer.
sample	Samples the traffic on the interface. Use this modifier only when traffic sampling is enabled. For more information, see the <i>JUNOS Policy Framework Configuration Guide</i> .
syslog	Records information in the system logging facility. This action can be used in conjunction with all options except discard .

Class-of-Service Overview

With the class-of-service (CoS) features on a Services Router, you can assign service levels with different delay, jitter (delay variation), and packet loss characteristics to particular applications served by specific traffic flows. CoS is especially useful for networks supporting time-sensitive video and audio applications. To configure CoS features on a Services Router, see “Configuring Class of Service with DiffServ” on page 529.

This overview contains the following topics. For more information about CoS and DiffServ, see the *JUNOS Class of Service Configuration Guide*.

- Benefits of DiffServ CoS on page 449
- DSCPs and Forwarding Service Classes on page 450
- JUNOS CoS Functions on page 451
- How Forwarding Classes and Schedulers Work on page 453

Benefits of DiffServ CoS

IP routers normally forward packets independently, without controlling throughput or delay. This type of packet forwarding, known as best-effort service, is as good as your network equipment and links allow. Best-effort service is sufficient for many traditional IP data delivery applications, such as e-mail or Web browsing. However, newer IP applications such as real-time video and audio (or voice) require lower delay, jitter, and packet loss than simple best-effort networks can provide.

CoS features allow a Services Router to improve its processing of critical packets while maintaining best-effort traffic flows, even during periods of congestion. Network throughput is determined by a combination of available bandwidth and delay. CoS dedicates a guaranteed minimum bandwidth to a particular service class by reducing forwarding queue delays. (The other two elements of overall network delay, serial transmission delays determined by link speeds and propagation delays determined by media type, are not affected by CoS settings.)

Normally, packets are queued for output in their order of arrival, regardless of service class. Queueing delays increase with network congestion and often result in lost packets when queue buffers overflow. CoS packet classification assigns packets to forwarding queues by service class.

Because CoS must be implemented consistently end-to-end through the network, the CoS features on the Services Router are based on IETF Differentiated Services (DiffServ) standards, to interoperate with other vendors' CoS implementations.

DSCPs and Forwarding Service Classes

DiffServ specifications establish a 6-bit field in the IP packet header to indicate the forwarding service class to apply to the packet. The bit values in the DiffServ field form DiffServ code points (DSCPs) that can be set by the application or by a Services Router on the edge of a DiffServ-enabled network.

Each DiffServ forwarding service class has a well-known name and alias. Although not part of the specifications, the aliases are well known through usage. For example, the alias for DSCP 101110 is widely accepted as **ef** (expedited forwarding).

The 21 well-known DSCPs establish five DiffServ service classes. Table 132 identifies the forwarding service classes and aliases that correspond to the 21 DSCPs.

Table 132: Default Forwarding Service Class-to-DSCP Mapping

DiffServ Service Class Alias	IP DSCP	Forwarding Service Class and Use
ef	101110	<p>Expedited forwarding—The Services Router delivers assured bandwidth, low loss, low delay, and low delay variation (jitter) end-to-end for packets in this service class.</p> <p>Routers accept excess traffic in this class, but in contrast to assured forwarding, out-of-profile expedited-forwarding packets can be forwarded out of sequence or dropped.</p>

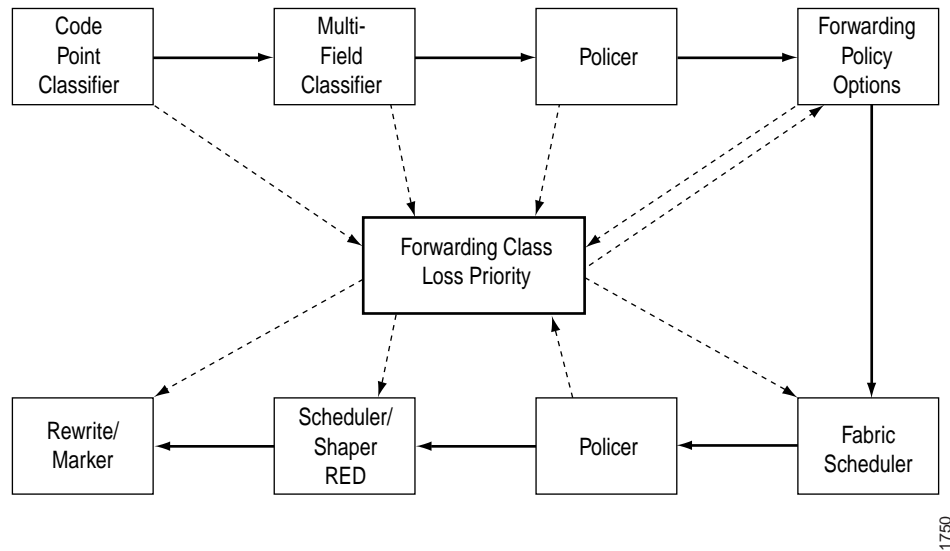
Table 132: Default Forwarding Service Class-to-DSCP Mapping (continued)

DiffServ Service Class Alias	IP DSCP	Forwarding Service Class and Use
af11	001010	Assured forwarding —The Services Router offers a high level of assurance that the packets are delivered as long as the packet flow from the customer stays within a certain service profile that you define.
af12	001100	
af13	001110	
af21	010010	The router accepts excess traffic, but applies a random early discard (RED) drop profile to decide if the excess packets are dropped and not forwarded.
af22	010100	
af23	010110	Three drop probabilities (low, medium, and high) are defined for this service class.
af31	011010	
af32	011100	
af33	011110	
af41	100010	
af42	100100	
af43	100110	
be	000000	Best-effort —The Services Router does not apply any special CoS handling to packets with 000000 in the DiffServ field, a backward compatibility feature. These packets are usually dropped under congested network conditions.
cs1	001000	Conversational services —The Services Router delivers assured (usually low) bandwidth with low delay and jitter for packets in this service class. Packets can be dropped, but are never delivered out of sequence. Packetized voice is a good example of a conversational service.
cs2	010000	
cs3	011000	
cs4	100000	
cs5	101000	
nc1/cs6	110000	Network control —The Services Router delivers packets in this service class with a low priority. (These packets are not delay sensitive.) Typically, these packets represent routing protocol hello or keepalive messages. Because loss of these packets jeopardizes proper network operation, delay is preferable to discard. (See also the conversational services description in this table.)
nc2/cs7	111000	

JUNOS CoS Functions

Although the DiffServ CoS specifications define the position and length of the DSCP in the packet header, the DiffServ implementation is vendor specific. DiffServ CoS functions in JUNOS software are implemented by a series of components that you configure individually or in combination to define particular service offerings.

Figure 77 shows the components of the JUNOS CoS features, illustrating the sequence in which they interact. Table 133 defines the components and explains their use.

Figure 77: Packet Flow Through JUNOS CoS-Configurable Components**Table 133: JUNOS CoS Components**

CoS Component	Use
Classifiers	<p>Associate incoming packets with a forwarding class and packet loss priority (PLP). The following types of classifiers are available:</p> <ul style="list-style-type: none"> ■ Behavior aggregate (BA) or code point traffic classifiers—Allow you to set the forwarding class and PLP based on DSCP. ■ Multifield (MF) traffic classifiers—Allow you to set the forwarding class and PLP based on firewall filter rules. This is usually done at the edge of the network for packets that do not have valid DSCPs in the packet headers.
Forwarding classes	<p>Allow you to set the scheduling and marking of packets as they transit the Services Router. Known as ordered aggregates in the DiffServ architecture, the forwarding class plus the loss priority determine the router's per-hop behavior (PHB in DiffServ) for CoS.</p>
Loss priorities	<p>Allow you to set the priority of dropping a packet before it is sent. Loss priority affects the scheduling of a packet without affecting the packet's relative ordering.</p>
Forwarding policy options	<ul style="list-style-type: none"> ■ Allow you to associate forwarding classes with next hops. ■ Allow you to create classification overrides, which assign forwarding classes to sets of prefixes.

Table 133: JUNOS CoS Components (continued)

CoS Component	Use
Transmission scheduling and rate control	<p>Provide you with a variety of tools to manage traffic flows. The following types are available:</p> <ul style="list-style-type: none"> ■ Schedulers—Allow you to define the priority, bandwidth, delay buffer size, rate control status, and RED drop profiles to be applied to a particular forwarding class for packet transmission. Drop profiles are useful for the assured forwarding service class. ■ Fabric schedulers—For M320 and T-series platforms only, fabric schedulers allow you to identify a packet as high or low priority based on its forwarding class, and to associate schedulers with the fabric priorities. ■ Policers for traffic classes—Allow you to limit traffic of a certain class to a specified bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a different forwarding class or to a different loss priority, or to both. You define policers with filters that can be associated with input or output interfaces. Policers are useful for the expedited forwarding service class.
Rewrite markers	<p>Allow you to redefine the DSCP value of outgoing packets. Rewriting or marking outbound packets is useful when the routing platform is at the border of a network and must alter the code points to meet the policies of the targeted peer.</p>

How Forwarding Classes and Schedulers Work

This section contains the following topics:

- “Default Forwarding Class Queue Assignments” on page 453
- “Default Scheduler Settings” on page 454
- “Default Behavior Aggregate (BA) Classifiers” on page 455
- “DSCP Rewrites” on page 456
- “Sample BA Classification” on page 456

Default Forwarding Class Queue Assignments

J-series routers have eight queues built into the hardware. If a classifier does not assign a packet to any other queue (for example, for other than well-known DSCPs that have not been added to the classifier), the packet is assigned by default to the class associated with queue 0.

Table 134 shows the four default forwarding classes and queues that Juniper Networks classifiers assign to packets based on the DSCP values in arriving packet headers.

Table 134: Default Forwarding Class Queue Assignments

Forwarding Class	Forwarding Queue
best-effort	queue 0
expedited-forwarding	queue 1
assured-forwarding	queue 2
network-control	queue 3

Because the Services Router supports up to eight queues, you can configure two queues for each forwarding class, one with high loss priority and one with low loss priority.

Default Scheduler Settings

Each forwarding class has an associated scheduler priority. Only two forwarding classes, **best-effort** and **network-control** (queue 0 and queue 3), are used in the JUNOS default scheduler configuration.

By default, the **best-effort** forwarding class (queue 0) receives 95 percent of the output link bandwidth and buffer space, and the **network-control** forwarding class (queue 3) receives 5 percent of the output link bandwidth and buffer space. The default drop profile causes the buffer to fill and then discard all packets until it again has space.

The **expedited-forwarding** and **assured-forwarding** classes have no schedulers, because by default no resources are assigned to queue 1 and queue 2. However, you can manually configure resources for **expedited-forwarding** and **assured-forwarding**.

The default scheduler settings are implicit in the configuration, although they do not appear in the output of the **show class-of-service** command.

```
[edit class-of-service]
schedulers {
  network-control {
    transmit-rate percent 5;
    buffer-size percent 5;
    priority low;
    drop-profile-map loss-priority any protocol any;
    drop-profile terminal;
  }
  best-effort {
    transmit-rate percent 95;
    buffer-size percent 95;
    priority low;
    drop-profile-map loss-priority any protocol any;
    drop-profile terminal;
  }
}
drop-profiles {
  terminal {
    fill-level 100 drop-probability 100;
```



```

    }
}

```

Default Behavior Aggregate (BA) Classifiers

Table 135 shows the forwarding class and packet loss priority (PLP) that are assigned by default to each well-known DSCP. Although several DSCPs map to the expedited-forwarding (ef) and assured-forwarding (af) classes, by default no resources are assigned to these forwarding classes. All af classes other than af1x are mapped to best-effort, because RFC 2597, *Assured Forwarding PHB Group*, prohibits a node from aggregating classes. Assignment to best-effort implies that the node does not support that class.

You can modify the default settings through configuration. For instructions, see “Configuring Class of Service with DiffServ” on page 529.

Table 135: Default Behavior Aggregate (BA) Classification

DSCP Alias	Forwarding Class	Packet Loss Priority (PLP)
ef	expedited-forwarding	low
af11	assured-forwarding	low
af12	assured-forwarding	high
af13	assured-forwarding	high
af21	best-effort	low
af22	best-effort	low
af23	best-effort	low
af31	best-effort	low
af32	best-effort	low
af33	best-effort	low
af41	best-effort	low
af42	best-effort	low
af43	best-effort	low
be	best-effort	low
cs1	best-effort	low
cs2	best-effort	low
cs3	best-effort	low
cs4	best-effort	low
cs5	best-effort	low
nc1/cs6	network-control	low

Table 135: Default Behavior Aggregate (BA) Classification (continued)

DSCP Alias	Forwarding Class	Packet Loss Priority (PLP)
nc2/cs7	network-control	low
other	best-effort	low

DSCP Rewrites

Typically, a router rewrites the DSCPs in outgoing packets once, when packets enter the DiffServ portion of the network, either because the packets do not arrive from the customer with the proper DSCP bit set or because the service provider wants to verify that that customer has set the DSCP properly. CoS implementations that accept the DSCP and classify and schedule traffic solely on DSCP value perform behavior aggregate (BA) DiffServ functions and do not usually rewrite the DSCP. DSCP rewrites typically occur in multifield (MF) DiffServ scenarios.

For instructions for configuring rewrite rules, see “Configuring and Applying Rewrite Rules (Required)” on page 537.

Sample BA Classification

Table 136 shows the router forwarding classes associated with each well-known DSCP code point and the resources assigned to their output queues for a sample DiffServ CoS implementation. This example assigns expedited forwarding to queue 1 and a subset of the assured forwarding classes (af1x) to queue 2, and distributes resources among all four forwarding classes.

Other DiffServ-based implementations are possible. For configuration information, see “Configuring Class of Service with DiffServ” on page 529.

Table 136: Sample BA Classification Forwarding Classes and Queues

DSCP Alias	DSCP Bits	Forwarding Class	PLP	Queue
ef	101110	expedited-forwarding	low	1
af11	001010	assured-forwarding	low	2
af12	001100	assured-forwarding	high	2
af13	001110	assured-forwarding	high	2
af21	010010	best-effort	low	0
af22	010100	best-effort	low	0
af23	010110	best-effort	low	0
af31	011010	best-effort	low	0
af32	011100	best-effort	low	0
af33	011110	best-effort	low	0

Table 136: Sample BA Classification Forwarding Classes and Queues (continued)

DSCP Alias	DSCP Bits	Forwarding Class	PLP	Queue
af41	100010	best-effort	low	0
af42	100100	best-effort	low	0
af43	100110	best-effort	low	0
be	000000	best-effort	low	0
cs1	0010000	best-effort	low	0
cs2	010000	best-effort	low	0
cs3	011000	best-effort	low	0
cs4	100000	best-effort	low	0
cs5	101000	best-effort	low	0
nc1/cs6	110000	network-control	low	3
nc2/cs7	111000	network-control	low	3
other	—	best-effort	low	0

Chapter 21

Configuring Routing Policies

Use routing policies as filters to control the information from routing protocols that a Services Router imports into its routing table and the information that the router exports (advertises) to its neighbors. To create a routing policy, you configure criteria against which routes are compared, and the action that is performed if the criteria are met.

You use either the J-Web configuration editor or CLI configuration editor to configure a routing policy.

This chapter contains the following topics. For more information about routing policies, see the *JUNOS Policy Framework Configuration Guide*.

- Before You Begin on page 460
- Configuring a Routing Policy with a Configuration Editor on page 460

Before You Begin

Before you begin configuring a routing policy, complete the following tasks:

- If you do not already have a basic understanding of routing policies, read “Routing Policy Overview” on page 435.
- Determine what you want to accomplish with the policy, and thoroughly understand how to achieve your goal using the various match conditions and actions.
- Make certain that you understand the default policies and actions for the policy you are configuring.
- Configure an interface on the router. See “Configuring Network Interfaces” on page 103.
- Configure an Interior Gateway Protocol (IGP) and Border Gateway Protocol (BGP), if necessary. See “Configuring BGP Sessions” on page 295.
- Configure the router interface to reject or accept routes, if necessary. See “Configuring Firewall Filters and NAT” on page 475.
- Configure static routes, if necessary. See “Configuring Static Routes” on page 237.

Configuring a Routing Policy with a Configuration Editor

A routing policy has a major impact on the flow of routing information or packets within and through the Services Router. The match conditions and actions allow you to configure a customized policy to fit your needs.

To configure a routing policy, you must perform the following tasks marked *(Required)*. Perform additional tasks as needed for your router. For information about using the J-Web and CLI configuration editors, see “Using J-Web Configuration Tools” on page 3.

- Configuring the Policy Name (Required) on page 461
- Configuring a Policy Term (Required) on page 461
- Rejecting Known Invalid Routes (Optional) on page 462
- Injecting OSPF Routes into the BGP Routing Table (Optional) on page 464
- Grouping Source and Destination Prefixes in a Forwarding Class (Optional) on page 466
- Configuring a Policy to Prepend the AS Path (Optional) on page 467
- Configuring Damping Parameters (Optional) on page 470

Configuring the Policy Name (Required)

Each routing policy is identified by a policy name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

Each routing policy name must be unique within a configuration.

To configure the policy name:

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 137.
3. Go on to “Configuring a Policy Term (Required)” on page 461.

Table 137: Configuring the Policy Name

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Policy statement level in the configuration hierarchy.	In the configuration editor hierarchy, select Policy options > Policy statement .	From the top of the configuration hierarchy, enter edit policy-options
Enter the policy name—for example, policy1.	<ol style="list-style-type: none"> 1. In the Policy name box, type policy1. 2. Click OK. 	Type the policy-name value: set policy-statement policy1

Configuring a Policy Term (Required)

Each routing policy term is identified by a term name. The name can contain letters, numbers, and hyphens (-) and can be up to 255 characters long. To include spaces in the name, enclose the entire name in double quotation marks.

To configure a policy term:

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 138.
3. If you are finished configuring the router, commit the configuration.
4. To configure additional routing policy features, go on to one of the following procedures:
 - To remove useless routes, see “Rejecting Known Invalid Routes (Optional)” on page 462.

- To advertise additional routes, see “Injecting OSPF Routes into the BGP Routing Table (Optional)” on page 464.
- To create a forwarding class, see “Grouping Source and Destination Prefixes in a Forwarding Class (Optional)” on page 466.
- To make a route less preferable to BGP, see “Configuring a Policy to Prepend the AS Path (Optional)” on page 467.
- To suppress route information, see “Configuring Damping Parameters (Optional)” on page 470.

Table 138: Configuring a Policy Term

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Policy statement level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Policy options > Policy statement. 2. Under Policy name, click policy1. 	<p>From the top of the configuration hierarchy, enter</p> <pre>edit policy-options policy-statement policy1</pre>
Create and name a policy term—for example, term1 .	<ol style="list-style-type: none"> 1. In the Term box, click Add new entry. 2. In the Term name box, type term1. 3. Click OK. 	<p>Create and name a policy term:</p> <pre>set term term1</pre>

Rejecting Known Invalid Routes (Optional)

You can specify known invalid (“bad”) routes to ignore by specifying matches on destination prefixes. When specifying a destination prefix, you can specify an exact match with a specific route, or a less precise match by using match types. You can configure either a common reject action that applies to the entire list, or an action associated with each prefix. Table 139 lists route list match types.

Table 139: Route List Match Types

Match Type	Match Conditions
exact	The route shares the same most-significant bits (described by <i>prefix-length</i>), and <i>prefix-length</i> is equal to the route’s prefix length.
longer	The route shares the same most-significant bits (described by <i>prefix-length</i>), and <i>prefix-length</i> is greater than the route’s prefix length.
orlonger	The route shares the same most-significant bits (described by <i>prefix-length</i>), and <i>prefix-length</i> is equal to or greater than the route’s prefix length.

Table 139: Route List Match Types (continued)

Match Type	Match Conditions
prefix-length-range <i>prefix-length2</i> - <i>prefix-length3</i>	The route shares the same most-significant bits (described by <i>prefix-length</i>), and the route's prefix length falls between <i>prefix-length2</i> and <i>prefix-length3</i> , inclusive.
through <i>destination-prefix</i>	<p>All the following are true:</p> <ul style="list-style-type: none"> ■ The route shares the same most-significant bits (described by <i>prefix-length</i>) of the first destination prefix. ■ The route shares the same most-significant bits (described by <i>prefix-length</i>) of the second destination prefix for the number of bits in the prefix length. ■ The number of bits in the route's prefix length is less than or equal to the number of bits in the second prefix. <p>You do not use the through match type in most routing policy configurations. For more information, see the <i>JUNOS Policy Framework Configuration Guide</i>.</p>
upto <i>prefix-length2</i>	The route shares the same most-significant bits (described by <i>prefix-length</i>) and the route's prefix length falls between <i>prefix-length</i> and <i>prefix-length2</i> .

For example, you can create a policy named `rejectpolicy1` to reject routes with a mask of /8 and greater (/8, /9, /10, and so on) that have the first 8 bits set to 0, and to accept routes less than 8 bits in length.

To create `rejectpolicy1`:

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 140.
3. If you are finished configuring the router, commit the configuration.
4. To configure additional routing policy features, go on to one of the following procedures:
 - To advertise additional routes, see “Injecting OSPF Routes into the BGP Routing Table (Optional)” on page 464.
 - To create a forwarding class, see “Grouping Source and Destination Prefixes in a Forwarding Class (Optional)” on page 466.
 - To make a route less preferable to BGP, see “Configuring a Policy to Prepend the AS Path (Optional)” on page 467.
 - To suppress route information, see “Configuring Damping Parameters (Optional)” on page 470.

Table 140: Creating a Policy to Reject Known Invalid Routes

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Policy statement level in the configuration hierarchy.	In the configuration editor hierarchy, select Policy options > Policy statement .	From the top of the configuration hierarchy, enter edit policy-options policy-statement
Create a rejection policy and term—for example, rejectpolicy1 and rejectterm1.	<ol style="list-style-type: none"> 1. In the Policy name box, type rejectpolicy1. 2. Next to Term, click Add new entry. 3. In the Term name box, type rejectterm1. 	Enter set rejectpolicy1 term rejectterm1
Specify the routes to accept—for example, routes with a mask of 0/0 up to /7.	<ol style="list-style-type: none"> 1. Next to From, click Configure. 2. Next to Route filter, click Add new entry. 3. In the Address box, type 0/0. 4. From the Modifier list, select Upto. 5. In the Upto box, type /7. 6. From the Accept reject list, select accept. 7. Click OK. 	Accept routes less than 8 bits in length: set from route-filter 0/0 up to /7 accept
Specify the routes to reject—for example, routes with a mask of /8 or greater.	<ol style="list-style-type: none"> 1. Next to Route filter, click Add new entry. 2. In the Address box, type /8. 3. From the Modifier list, select Orlonger. 4. From the Accept reject list, select reject. 5. Click OK. 	<ol style="list-style-type: none"> 1. Specify routes less than 8 bits in length: set from route-filter /8 orlonger 2. Reject these routes: set then reject

Injecting OSPF Routes into the BGP Routing Table (Optional)

You can specify a match condition for policies based on protocols by naming a protocol from which the route is learned or to which the route is being advertised. You can specify one of the following protocols: aggregate, BGP, direct, DVMRP, IS-IS, local, OSPF, PIM-dense, PIM-sparse, RIP, or static

For example, you can inject or redistribute OSPF routes into the BGP routing table by creating a routing policy.

To create a routing policy named injectpolicy1 that redistributes OSPF routes from area 1 only into BGP and does not advertise routes learned by BGP:

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 141.
3. If you are finished configuring the router, commit the configuration.
4. To configure additional routing policy features, go on to one of the following procedures:
 - To create a forwarding class, see “Grouping Source and Destination Prefixes in a Forwarding Class (Optional)” on page 466.
 - To make a route less preferable to BGP, see “Configuring a Policy to Prepend the AS Path (Optional)” on page 467.
 - To suppress route information, see “Configuring Damping Parameters (Optional)” on page 470.

Table 141: Creating a Policy to Inject OSPF Routes into BGP

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Policy statement level in the configuration hierarchy.	In the configuration editor hierarchy, select Policy options > Policy statement .	From the top of the configuration hierarchy, enter edit policy-options policy-statement
Create an injection policy and term—for example, injectpolicy1 and injectterm1.	<ol style="list-style-type: none"> 1. In the Policy name box, type injectpolicy1. 2. Next to Term, click Add new entry. 3. In the Term name box, type injectterm1. 	Enter set injectpolicy1 term injectterm1
Specify the OSPF routes.	<ol style="list-style-type: none"> 1. In the From option, click Configure. 2. In the Protocol box, click Add new entry. 3. In the Value drop box, select ospf. 4. Click OK. 	Specify the OSPF match condition: set from ospf
Specify the routes from a particular OSPF area—for example, area 1.	<ol style="list-style-type: none"> 1. In the Area box, type 1. 2. Click OK. 	Specify Area 1 as a match condition: set from area 1
Specify that the route is to be accepted if the previous conditions are matched. Set the default option to reject other OSPF routes.	<ol style="list-style-type: none"> 1. Next to Then, click Configure. 2. From the Accept reject list, Select accept. 3. From the Default action list, Select reject. 4. Click OK until you return to the Configuration page. 	Specify the action to accept: set then accept

Table 141: Creating a Policy to Inject OSPF Routes into BGP (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Bgp level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Bgp .	From the top of the configuration hierarchy, enter edit protocols bgp
Apply the routing policy injectpolicy1 to BGP.	<ol style="list-style-type: none"> Next to Export, click Add new entry. In the Value option, type injectpolicy1. Click OK. 	Specify the OSPF match condition: set export injectpolicy1

Grouping Source and Destination Prefixes in a Forwarding Class (Optional)

Create a forwarding class called forwarding-class1 that includes packets based on both the destination address and the source address in the packet.

To configure and apply the routing policy policy1, which you configured in Table 137 and Table 138, to group source and destination prefixes in a forwarding class:

- Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- Perform the configuration tasks described in Table 142.
- If you are finished configuring the router, commit the configuration.
- To configure additional routing policy features, go on to one of the following procedures:
 - To make a route less preferable to BGP, see “Configuring a Policy to Prepend the AS Path (Optional)” on page 467.
 - To suppress route information, see “Configuring Damping Parameters (Optional)” on page 470.

Table 142: Creating a Policy to Group Source and Destination Prefixes in a Forwarding Class

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the term1 level in the configuration hierarchy.	In the configuration editor hierarchy, select Policy options > Policy statement policy1 > Term term1 .	From the top of the configuration hierarchy, enter edit policy-options policy-statement policy1 term term1

Table 142: Creating a Policy to Group Source and Destination Prefixes in a Forwarding Class (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Specify the routes to include in the route filter. For example: <ul style="list-style-type: none"> ■ Source routes greater than or equal to 10.210.0.0/16 ■ Destination routes greater than or equal to 10.215.0.0/16 	1. Next to From, click Configure .	Specify the source routes for the route filter: set from route-filter 10.210.0.0/16 orlonger
	2. Next to Route filter, click Add new entry .	
	3. In the Address box, type 10.210.0.0/16.	
	4. From the Modifier list, select Orlonger.	
	5. Click OK to return to the From page.	
	1. Next to Route filter, click Add new entry .	Specify the destination routes for the route filter: set from route-filter 10.215.0.0/16 orlonger
	2. In the Address box, type 10.215.0.0/16.	
	3. From the Modifier list, select Orlonger.	
	4. Click OK until you return to the Term page.	
Group the source and destination prefixes into a forwarding class—for example, forwarding-class1.	1. Next to Then, click Configure .	Specify the forwarding class name:
	2. In the Forwarding class box, type forwarding-class1.	set then forwarding class forwarding-class1
	3. Click OK .	
Navigate to the Forwarding table level in the configuration hierarchy.	In the configuration editor hierarchy, select Routing options > Forwarding table .	From the top of the configuration hierarchy, enter edit routing-options forwarding-table
Apply the policy1 policy to the forwarding table.	1. Next to Export, click Add new entry .	Specify the routing policy to apply:
The routing policy is evaluated when routes are being exported from the routing table into the forwarding table. Only active routes are exported from the routing table.	2. In the Value box, type policy1.	set export policy1
	3. Click OK .	You can refer to the same routing policy one or more times in the same or a different export statement.

Configuring a Policy to Prepend the AS Path (Optional)

You can *prepend* or add one or more autonomous system (AS) numbers at the beginning of an AS path. The AS numbers are added after the local AS number has been added to the path. Prepending an AS path makes a shorter AS path look longer and therefore less preferable to the Border Gateway Protocol (BGP).

For example, from AS 1, there are two equal paths (through AS 2 and AS 3) to reach AS 4. You might want packets from certain sources to use the path through AS 2. Therefore, you must make the path through AS 3 look less preferable so that BGP chooses the path through AS 2. In AS 1, you can prepend multiple AS numbers.

To create a routing policy `prependpolicy1` that prepends multiple AS numbers:

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 143.
3. If you are finished configuring the router, commit the configuration.
4. To suppress route information, see “Configuring Damping Parameters (Optional)” on page 470.

Table 143: Creating a Policy to Prepend AS Numbers

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Policy statement level in the configuration hierarchy.	In the configuration editor hierarchy, select Policy options > Policy statement .	From the top of the configuration hierarchy, enter edit policy-options policy-statement
Create a prepend policy and term—for example, <code>prependpolicy1</code> and <code>prependterm1</code> .	<ol style="list-style-type: none"> 1. In the Policy name box, type <code>prependpolicy1</code>. 2. Next to Term, click Add new entry. 3. In the Term name box, type <code>prependterm1</code>. 	Enter set prependpolicy1 term prependterm1

Table 143: Creating a Policy to Prepend AS Numbers (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Specify the routes to prepend AS numbers to. For example: <ul style="list-style-type: none"> ■ Routes greater than or equal to 172.16.0.0/12 ■ Routes greater than or equal to 192.168.0.0/16 ■ Routes greater than or equal to 10.0.0.0/8 	1. Next to From, click Configure . 2. Next to Route filter, click Add new entry . 3. In the Value box, type 172.16.0.0/12. 4. From the Modifier list, select Orlonger . 5. Click OK .	Specify the first routes to prepend: set from route-filter 172.16.0.0/12 orlonger
	1. Next to From, click Configure . 2. Next to Route filter, click Add new entry . 3. In the Value box, type 192.168.0.0/16. 4. From the Modifier list, select Orlonger . 5. Click OK .	Specify the next routes to prepend: set from route-filter 192.168.0.0/16 orlonger
	1. Next to From, click Configure . 2. Next to Route filter, click Add new entry . 3. In the Value box, type 10.0.0.0/8. 4. From the Modifier list, select Orlonger . 5. Click OK until you return to the Term page.	Specify the last routes to prepend: set from route-filter 10.0.0.0/8 orlonger
Specify the AS numbers to prepend. Separate each AS number with a space—for example, 1 1 1 1.	1. Next to Then, click Configure . 2. In the AS path prepend box, type 1 1 1 1. 3. Click OK .	Specify the AS numbers to prepend, and enclose them inside double quotation marks: set then as-path-prepend "1 1 1 1"
Navigate to the Bgp level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Bgp .	From the top of the configuration hierarchy, enter edit protocols bgp
Apply the prependpolicy1 policy as an import policy for all BGP routes.	1. Next to Import, click Add new entry .	Apply the policy: set import prependpolicy1
The routing policy is evaluated when routes are being imported to the routing table.	2. In the Value box, type prependpolicy1 . 3. Click OK .	You can refer to the same routing policy one or more times in the same or a different import statement.

Configuring Damping Parameters (Optional)

Flap damping reduces the number of update messages by marking routes as ineligible for selection as the active or preferable route. Marking routes in this way leads to some delay, or *suppression*, in the propagation of route information, but the result is increased network stability. You typically apply flap damping to external BGP (EBGP) routes (routes in different ASs). You can also apply flap damping within a confederation, between confederation member ASs. Because routing consistency within an AS is important, do not apply flap damping to internal BGP (IBGP) routes. (If you do, it is ignored.)

You can specify one or more of the damping parameters described in Table 144. If you do not specify a damping parameter, the default value of the parameter is used.

Table 144: Damping Parameters

Damping Parameter	Description	Default Value	Possible Values
half-life <i>minutes</i>	Decay half-life—Number of minutes after which an arbitrary value is halved if a route stays stable.	15 (minutes)	1 through 4
max-suppress <i>minutes</i>	Maximum hold-down time for a route, in minutes.	60 (minutes)	1 through 720
reuse	Reuse threshold—Arbitrary value below which a suppressed route can be used again.	750	1 through 20000
suppress	Cutoff (suppression) threshold—Arbitrary value above which a route can no longer be used or included in advertisements.	3000	1 through 20000

To change the default BGP flap damping values, you define actions by creating a named set of damping parameters and including it in a routing policy with the damping action. For the damping routing policy to work, you also must enable BGP route flap damping.

To configure damping with a policy named `dampenpolicy1`, 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 145.
3. If you are finished configuring the router, commit the configuration.

Table 145: Creating a Policy to Accept and Apply Damping on Routes

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Policy statement level in the configuration hierarchy.	In the configuration editor hierarchy, select Policy options > Policy statement .	From the top of the configuration hierarchy, enter edit policy-options policy-statement
Create a damping policy and term—for example, dampenpolicy1 and dampenterm1.	<ol style="list-style-type: none"> 1. In the Policy name box, type dampenpolicy1. 2. Next to Term, click Add new entry. 3. In the Term name box, type dampenterm1. 	Enter set dampenpolicy1 term dampenterm1
Specify the routes to dampen and associate each group of routes with a group name. For example: <ul style="list-style-type: none"> ■ group1—Routes greater than or equal to 172.16.0.0/12 ■ group2—Routes greater than or equal to 192.168.0.0/16 ■ group3—Routes greater than or equal to 10.0.0.0/8 	<ol style="list-style-type: none"> 1. Next to From, click Configure. 2. Next to Route filter, click Add new entry. 3. In the Address box, type 172.16.0.0/12. 4. In the Damping box, type group1. 5. From the Modifier list, select Orlonger. 6. Click OK. 	Specify the first routes to dampen: set from route-filter 172.16.0.0/12 orlonger damping group 1
	<ol style="list-style-type: none"> 1. Next to Route filter, click Add new entry. 2. In the Address box, type 192.168.0.0/16. 3. In the Damping box, type group2. 4. From the Modifier list, select Orlonger. 5. Click OK. 	Specify the next routes to dampen: set from route-filter 192.168.0.0/16 orlonger
	<ol style="list-style-type: none"> 1. Next to Route filter, click Add new entry. 2. In the Address box, type 10.0.0.0/8. 3. In the Damping box, type group3. 4. From the Modifier list, select Orlonger. 5. Click OK until you return to the Policy options page. 	Specify the last routes to dampen: set from route-filter 10.0.0.0/8 orlonger

Table 145: Creating a Policy to Accept and Apply Damping on Routes (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Create three damping parameter groups with different damping actions. For example:</p> <ul style="list-style-type: none"> ■ group1—Increases the half-life to 30 minutes. All other parameters are left at their default values. ■ group2—Increases the half-life to 40 minutes, decreases the maximum hold-down time for a route to 45 minutes, increases the reuse value to 1000, and reduces the cutoff (suppression) threshold to 400. ■ group3—Disables route damping. 	<p>For <i>each</i> damping group:</p> <ol style="list-style-type: none"> 1. Next to Damping, click Add new entry. 2. In the Damping object name box, type the name of a damping group—for example, group1. 3. In the Half life box, type the half-life duration, in minutes: <ul style="list-style-type: none"> ■ For group1—30 ■ For group2—40 4. In the Max suppress box, type the maximum hold-down time, in minutes: <ul style="list-style-type: none"> ■ For group1—60 (the default) ■ For group2—45 5. In the Reuse box, type the reuse threshold, for this damping group: <ul style="list-style-type: none"> ■ For group1—750 (the default) ■ For group2—1000 6. In the Suppress box, type the cutoff threshold, for this damping group: <ul style="list-style-type: none"> ■ For group1—3000 (the default) ■ For group2—400 7. To disable damping for the group3 damping group, select the Disable check box. 8. Click OK when you finish configuring each group. 	<p>Create and configure the damping parameter groups:</p> <pre>edit damping group1 half-life 30 max-suppress 60 reuse 750 suppress 3000 edit damping group2 half-life 40 max-suppress 45 reuse 1000 suppress 400 edit damping group3 disable</pre>
Navigate to the Bgp level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Bgp .	From the top of the configuration hierarchy, enter <pre>edit protocols bgp</pre>
Enable damping.	<ol style="list-style-type: none"> 1. Select the Damping check box. 2. Click OK. 	<p>Enable damping:</p> <pre>set damping</pre>

Table 145: Creating a Policy to Accept and Apply Damping on Routes (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Neighbor level in the configuration hierarchy, for the BGP neighbor to which you want to apply the damping policy—for example, the neighbor at IP address 172.16.15.14 .	In the configuration editor hierarchy, select Protocols > Bgp > Group GroupA > Neighbor 172.16.15.14 .	From the top of the configuration hierarchy, enter edit protocols bgp group groupA neighbor 172.16.15.14
Apply the policy as an import policy for the BGP neighbor.	1. Next to Import, click Add new entry .	Apply the policy: set import dampenpolicy1
The routing policy is evaluated when routes are imported to the routing table.	2. In the Value box, type the name of the policy. 3. Click OK .	You can refer to the same routing policy one or more times in the same or a different import statement.

Chapter 22

Configuring Firewall Filters and NAT

A *stateful* firewall filter inspects traffic flowing between a trusted network and an untrusted network. Contrasted with a *stateless* firewall filter that inspects packets in isolation, a stateful firewall filter provides an extra layer of security by using state information derived from past communications and other applications to make dynamic control decisions.

The Services Router uses the stateful firewall filter as a basis for performing Network Address Translation (NAT).



NOTE: You must have a license to configure a stateful firewall filter and NAT. For license details, see the *J-series Services Router Administration Guide*.

You can use either J-Web Quick Configuration or a configuration editor to configure stateful firewall filters and NAT and to configure a stateless firewall filter.

This chapter contains the following topics. For more information about firewall filters, see the *JUNOS Policy Framework Configuration Guide*. For more information about NAT, see the *JUNOS Services Interfaces Configuration Guide*.

- Before You Begin on page 476
- Configuring a Stateful Firewall Filter with Quick Configuration on page 476
- Configuring a Stateful Firewall Filter with a Configuration Editor on page 480
- Configuring a Stateless Firewall Filter with Quick Configuration on page 486
- Configuring a Stateless Firewall Filter with a Configuration Editor on page 501
- Verifying Firewall Filter Configuration on page 517

Before You Begin

Before you begin configuring firewall filters, complete the following tasks:

- If you do not already have an understanding of firewall filters, read “Firewall Filter Overview” on page 440.
- Before you begin configuring stateful firewall filters and NAT, you must configure the interfaces on which to apply these services. To configure an interface, see “Configuring Network Interfaces” on page 103.

Unlike a stateful firewall filter, you can configure a stateless firewall filter before configuring the interfaces on which they are applied.

Configuring a Stateful Firewall Filter with Quick Configuration

You can use the Firewall/NAT Quick Configuration pages to configure a stateful firewall filter and NAT. These Quick Configuration pages allow you to designate the interfaces that make up the untrusted network. In addition, you can designate the applications that are allowed to operate from the untrusted network to the trusted network.

Figure 78 and Figure 79 show the Firewall/NAT Quick Configuration main and application pages.

Figure 78: Firewall/NAT Quick Configuration Main Page

Juniper NETWORKS **ROUTER - J6300** Logged in as: regress [Help](#) [About](#) [Logout](#)

Monitor Configuration Diagnose Manage

Configuration > Quick Configuration > Firewall/NAT

Quick Configuration

Firewall/NAT

Stateful Firewall

Stateful firewall inspects traffic flowing between a trusted network and an untrusted network. All packets flowing from a trusted network to an untrusted network are allowed. Packets flowing from an untrusted network to a trusted network are allowed only if they are responses to a session originated by the trusted network.

Enable Stateful Firewall ☐

Trusted Interfaces

Select the interfaces to be part of a trusted network. Stateful firewall is applied to the untrusted interfaces.

Untrusted Interfaces	Trusted Interfaces
	e1-1/0/0.0 fe-0/0/0.0 t3-4/0/0.0

Network Address Translation (NAT)

When NAT is enabled, the source address of a packet flowing from a trusted network to an untrusted network is replaced with an address chosen from the specified range. The source port of the packet is also replaced with a dynamically chosen port.

Enable NAT ☐

• **Low Address in Address Range**

High Address in Address Range

Outside Applications Allowed

The following applications are allowed to operate from the untrusted network to the trusted network.

No applications are allowed from the untrusted network onto the trusted network.

Figure 79: Firewall/NAT Quick Configuration Application Page

Juniper NETWORKS

ROUTER - J4300

Logged in as: regress

[Help](#) [About](#) [Logout](#)

Monitor Configuration Diagnose Manage

Configuration > Quick Configuration > Firewall/NAT

Quick Configuration

Firewall/NAT

Allow an Application Through the Firewall

Application

* Application

Source Address

Any Unicast WAN Address ☒

Source Addresses and Prefixes

Source Address	Prefix

Add Delete

To configure a stateful firewall filter and NAT with Quick Configuration:

1. In the J-Web interface, select **Configuration > Firewall/NAT**.
2. Enter information into the Firewall/NAT Quick Configuration pages, as described in Table 146.
3. Click one of the following buttons on the Firewall/NAT Quick Configuration main page:
 - To apply the configuration and stay in the Firewall/NAT Quick Configuration main page, click **Apply**.
 - To apply the configuration and return to the Quick Configuration page, click **OK**.
 - To cancel your entries and return to the Quick Configuration page, click **Cancel**.
4. Go on to one of the following procedures:

- To display the configuration, see “Displaying Firewall Filter Configurations” on page 517.
- To verify a stateless firewall filter, see “Verifying Firewall Filter Configuration” on page 517.

Table 146: Firewall/NAT Quick Configuration Pages Summary

Field	Function	Your Action
Stateful Firewall		
Enable Stateful Firewall	Enables stateful firewall filter configuration.	To enable stateful firewall filter configuration, select the check box.
Trusted Interfaces		
Trusted Interfaces	Designates the trusted and untrusted router interfaces. The stateful firewall filter is applied to the untrusted interfaces.	<p>The Trusted Interfaces box displays a list of all the interfaces configured on the router. Do either of the following:</p> <ul style="list-style-type: none"> ■ To <i>apply</i> a stateful firewall filter to an interface, click the interface in the Trusted Interfaces box to highlight it, and click the left arrow to add the interface to the Untrusted Interfaces list. You can select multiple interfaces by pressing Ctrl while you click the interface. ■ To <i>remove</i> a stateful firewall filter from an interface, click the interface in the Untrusted Interfaces box to highlight it, and click the right arrow to add the interface to the Trusted Interfaces list. You can select multiple interfaces by pressing Ctrl while you click the interface.
Network Address Translation (NAT)		
Enable NAT	Enables NAT configuration.	To enable NAT configuration, select the check box.
Low Address in Address Range (required)	Specifies the lowest address in the NAT pool address range. If a range of addresses is not specified, you can specify a single address or an IP prefix.	Type an IP address or prefix.
High Address in Address Range	Specifies the highest address in the NAT pool address range.	Type an IP address. The total range of addresses in the pool must be limited to a maximum of 32.
Outside Applications Allowed		
	Add or delete applications that are allowed to operate from the untrusted network to the trusted network.	<p>Click Add to move to the Firewall/NAT Quick Configuration application page. When you have finished entering information into this page, click OK to save it.</p> <p>To cancel your entries, click Cancel.</p>

Table 146: Firewall/NAT Quick Configuration Pages Summary (continued)

Field	Function	Your Action
Application		
Application (required)	Designate which applications are allowed to operate from the untrusted network to the trusted network.	From the list, select the application you want to operate from the untrusted network to the trusted network.
Source Address		
Any Unicast WAN Address	Specifies that any unicast source address is allowed from the untrusted network.	To allow any unicast source address, select the check box.
Source Addresses and Prefixes	Designates the source addresses and prefixes that are allowed from the untrusted network.	<p>To add an IP address and prefix, type them in the boxes above the Add button, then click Add.</p> <p>To delete an IP address and prefix, select them in the Source Addresses and Prefixes box, then click Delete.</p>
Destination Address		
Any Unicast LAN Address	Specifies that any unicast destination address is allowed from the untrusted network.	To allow any unicast destination address, select the check box.
Destination Addresses and Prefixes	Designates the destination addresses and prefixes that are allowed from the untrusted network.	<p>To add an IP address and prefix, type them in the boxes above the Add button, then click Add.</p> <p>To delete an IP address and prefix, select them in the Destination Addresses and Prefixes box, then click Delete.</p>

Configuring a Stateful Firewall Filter with a Configuration Editor

To configure a stateful firewall filter and NAT with a configuration editor, you do the following:

- Define the filter's input and output rules.



CAUTION: If a packet does not match any terms in a stateful firewall filter rule, the packet is discarded. Take care that you do not configure a firewall filter that prevents you from accessing the Services Router after you commit the configuration.

- Define an address pool and port pool for NAT.
- Define NAT input and output rules.
- Define a *service set* that includes the rules in the filter and NAT and the virtual `sp-0/0/0` services interface.
- Finally, apply the service set to any interfaces on the Services Router that lead to or from the untrusted network.

The example in this section shows how to create a stateful firewall filter and NAT with the rules described in Table 147.

Table 147: Sample Stateful Firewall Filter and NAT Rules

Rule	Type	Term or Terms
to-wan-rule	Output	<ul style="list-style-type: none"> ■ app-term—Accepts packets from any of the applications defined by the JUNOS default group <code>junos-algs-outbound</code> application set. ■ accept-all-term—Accepts packets that do not match app-term.
from-wan-rule	Input	<ul style="list-style-type: none"> ■ wan-src-addr-term—Accepts input packets with a source prefix of <code>192.168.33.0/24</code>. ■ discard-all-term—Discards all packets.
nat-to-wan-rule	Output	private-public-term —Translates the source address to an address within the pool <code>10.148.2.1</code> through <code>10.148.2.32</code> and dynamically translates the source port to a router-assigned port by means of NAPT

The example also assigns the name `public-pool` to the NAT address pool and NAPT router-assigned port.

In addition, the example creates the service set `wan-service-set` that includes the stateful firewall filter and NAT services and defines `sp-0/0/0` as its service interface. Finally, `wan-service-set` is applied to the WAN interface to the untrusted network, `t1-0/0/0`.

For stateful firewall match conditions and actions, see “Summary of Stateful Firewall Filter and NAT Match Conditions and Actions” on page 442.

To configure a stateful firewall filter and NAT and apply them to the WAN 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 148.
3. To apply the stateful firewall filter and NAT to the interface, perform the configuration tasks described in Table 149.
4. If you are finished configuring the router, commit the configuration.
5. Go on to one of the following procedures:
 - To display the configuration, see “Displaying Firewall Filter Configurations” on page 517.
 - To verify the stateful firewall filter, see “Verifying a Stateful Firewall Filter” on page 522.

Table 148: Configuring a Stateful Firewall Filter and NAT

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Stateful firewall level in the configuration hierarchy.	In the configuration editor hierarchy, select Services > Stateful firewall .	From the top of the configuration hierarchy, enter edit services stateful-firewall .
Define to-wan-rule and set its match direction.	<ol style="list-style-type: none"> 1. Next to Rule, click Add new entry. 2. In the Rule name box, type to-wan-rule. 3. From the Match direction list, select output. 	Set the rule name, match direction, term name, and match condition: set rule to-wan-rule match-direction output term app-term from application-sets junos-algs-outbound
Define app-term for the to-wan-rule rule.	<ol style="list-style-type: none"> 1. Next to Term, click Add new entry. 2. In the Term name box, type app-term. 	
Define the match condition for app-term —the default junos-algs-outbound application set.	<ol style="list-style-type: none"> 1. Next to From, click Configure. 2. Next to Application sets, click Add new entry. 3. In the Application set name box, type junos-algs-outbound. 4. Click OK twice. 	
Define an action for app-term .	<ol style="list-style-type: none"> 1. On the Term app-term page, next to Then, click Configure. 2. In the Designation list, select Accept. 3. Click OK twice. 	Set the action: set rule to-wan-rule term app-term then accept

Table 148: Configuring a Stateful Firewall Filter and NAT (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define accept-all-term for to-wan-rule .	<ol style="list-style-type: none"> 1. On the Rule to-wan-rule page, next to Term, click Add new entry. 2. In the Term name box, type accept-all-term. 	<p>Set the term name and the action:</p> <p>set rule to-wan-rule term accept-all-term then accept</p>
Define an action for accept-all-term . The action is taken only if a packet does not match app-term .	<ol style="list-style-type: none"> 1. Next to Then, click Configure. 2. From the Designation list, select Accept. 3. Next to Accept, select the check box. 4. Click OK three times. 	
Define from-wan-rule and set its match direction.	<ol style="list-style-type: none"> 1. On the Rule page, next to Rule, click Add new entry. 2. In the Rule name box, type from-wan-rule. 3. From the Match direction list, select input. 	<p>Set the rule name, match direction, term name, and the match condition:</p> <p>set rule from-wan-rule match-direction input term wan-src-addr-term from source-address 192.168.33.0/24</p>
Define wan-src-addr-term for the from-wan-rule rule.	<ol style="list-style-type: none"> 1. Next to Term, click Add new entry. 2. In the Term name box, type wan-src-addr-term. 	
Define the match condition for wan-src-addr-term .	<ol style="list-style-type: none"> 1. Next to From, click Configure. 2. Next to Source address, click Add new entry. 3. From the Address list, select Enter Specific Value—>. 4. In the Prefix box, type 192.168.33.0/24. 5. Click OK twice. 	
Define an action for wan-src-addr-term .	<ol style="list-style-type: none"> 1. On the Term wan-src-addr-term page, next to Then, click Configure. 2. In the Designation list, select Accept. 3. Click OK twice. 	<p>Set the action:</p> <p>set rule from-wan-rule term wan-src-addr-term then accept</p>
Define discard-all-term for from-wan-rule .	<ol style="list-style-type: none"> 1. On the Rule from-wan-rule page, next to Term, click Add new entry. 2. In the Term name box, type discard-all-term. 	<p>Set the term name and the action:</p> <p>set rule from-wan-rule term discard-all-term then discard</p>
Define an action for discard-all-term . The action is taken only if a packet does not match wan-src-addr-term .	<ol style="list-style-type: none"> 1. Next to Then, click Configure. 2. From the Designation list, select Discard. 3. Click OK three times. 	

Table 148: Configuring a Stateful Firewall Filter and NAT (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Nat level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Services. 2. Next to NAT, click Configure. 	From the top of the configuration hierarchy, enter <code>edit services nat</code> .
Define the public-pool address pool name and range.	<ol style="list-style-type: none"> 1. Next to Pool, click Add new entry. 2. In the Pool name box, type <code>public-pool</code>. 3. From the Address choice list, select Address range. 4. In the High box, type <code>10.148.2.32</code>. In the Low box, <code>10.148.2.1</code>. 	<p>Set the address pool name and the range:</p> <pre>set pool public-pool address-range low 10.148.2.1 high 10.148.2.32</pre>
Specify the NAT port pool to be automatically assigned by the router.	<ol style="list-style-type: none"> 1. Next to Port, click Configure. 2. From the Port choice list, select Automatic. 3. Click OK twice. 	<p>Configure the source port translation to be automatic:</p> <pre>set pool public-pool port automatic</pre>
Define nat-to-wan-rule and private-public-term .	<ol style="list-style-type: none"> 1. On the Nat page, next to Rule, click Add new entry. 2. In the Rule name box, type <code>nat-to-wan-rule</code>. 3. From the Match direction list, select output. 4. Next to Term, select Add new entry. 5. In the Term name box, type <code>private-public-term</code>. 6. Next to Then, select Configure. 7. Next to Translated, select Configure. 8. In the Source pool box, type <code>public-pool</code>. 	<p>Set the rule name, match direction, term name, and the term's pool name:</p> <pre>set rule nat-to-wan-rule match-direction output term private-public-term then translated source-pool public-pool</pre>
Set the NAT port translation type for private-public-term .	<ol style="list-style-type: none"> 1. Next to Translation type, select the check box. 2. Select Configure. 3. From the Source list, select dynamic. 4. Click OK five times. 	<p>Set the NAT translation type:</p> <pre>set rule nat-to-wan-rule match-direction output term private-public-term then translated translation-type source dynamic</pre>

Table 149: Applying a Stateful Firewall Filter and NAT to an Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Services level in the configuration hierarchy.	1. In the configuration editor hierarchy, select Services .	From the top of the configuration hierarchy, enter edit services .
Define wan-service-set and assign the stateful firewall filter rule to-wan-rule to the service set.	1. Next to Service set, click Add new entry . 2. In the Service set name box, type wan-service-set . 3. From the Stateful firewall rules choice list, select Stateful firewall rules . 4. Next to Stateful firewall rules, click Add new entry . 5. In the Rule name box, type to-wan-rule . 6. Click OK .	Define the service set and assign the rule: set service-set wan-service-set stateful-firewall-rules to-wan-rule
Assign the stateful firewall filter rule from-wan-rule to the service set.	1. Next to Stateful firewall rules, click Add new entry . 2. In the Rule name box, type from-wan-rule . 3. Click OK .	Define the service set and assign the rule: set service-set wan-service-set stateful-firewall-rules from-wan-rule
Assign the NAT rule nat-to-wan-rule to the service set.	1. From the Nat rules choice list, select Nat rules . 2. Next to Nat rules, click Add new entry . 3. In the Rule name box, type nat-to-wan-rule . 4. Click OK .	Assign the rule to the service set: set service-set wan-service-set nat-rules nat-to-wan-rule
Define the service set type and virtual interface sp-0/0/0 as the service interface for wan-service-set .	1. From the Service type choice list, select Interface service . 2. Next to Interface service, click Configure . 3. In the Service interface box, type sp-0/0/0 . 4. Click OK .	Define the service set type and the service interface: set service-set wan-service-set interface-service service-interface sp-0/0/0

Table 149: Applying a Stateful Firewall Filter and NAT to an Interface (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the sp-0/0/0 service interface.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select interfaces. 2. Next to Interface, click Add new entry. 3. In the Interface name box, type sp-0/0/0. 4. Next to Unit, click Add new entry. 5. In the Interface unit number box, type 0. 6. Next to Inet, select the check box. 7. Click Configure. 8. Click OK. 	<p>From the top of the configuration hierarchy, configure the interface:</p> <pre>set interfaces sp-0/0/0 unit 0 family inet</pre>
From the Interfaces level of the configuration hierarchy, navigate to the Inet level of the T1 interface—the untrusted interface in this example—and apply wan-service-set to the input and output sides of the t1-0/0/0 interface.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Interfaces > t1-0/0/0 > Unit > 0 > Family > Inet. 2. Next to Service, click Configure. 3. Next to Input, click Configure. 4. Next to Service set, click Add new entry. 5. In the Service set name box, type wan-service-set. 6. Click OK. 7. Next to Output, click Configure. 8. Next to Service set, click Add new entry. 9. In the Service set name box, type wan-service-set. 10. Click OK. 	<p>From the top of the configuration hierarchy, apply the service set to the interface:</p> <pre>set interfaces t1-0/0/0 unit 0 family inet service input service-set wan-service-set service output service-set wan-service-set</pre>

Configuring a Stateless Firewall Filter with Quick Configuration

The Firewall Filters Quick Configuration pages allow you to configure stateless firewall filters that examine packets traveling to or from a Services Router. You can create new filters or edit existing filters by adding terms to them. Each filter term is defined by a set of match conditions and an associated action. After you define the terms for a filter, you must associate the filter with one or more interfaces on the router.

This section contains the following topics:

- Configuring IPv4 and IPv6 Stateless Firewall Filters on page 487

- Assigning IPv4 and IPv6 Firewall Filters to Interfaces on page 499

Configuring IPv4 and IPv6 Stateless Firewall Filters

Using the Firewall Filters Quick Configuration pages, you can create filters and terms and define match conditions and actions for each filter term. For a description of match conditions, see Table 129, and for a description of actions, see Table 131.

Figure 80 shows the Initial Firewall Filters Quick Configuration page that displays existing firewall filters and allows you to add and modify filters.

Figure 81 shows the Match Conditions and Actions Quick Configuration page for configuring match conditions and resulting actions of filter terms.

Figure 80: Initial Firewall Filters Quick Configuration Page

Juniper NETWORKS ROUTER - J6300

Logged in as: regress [Help](#) [About](#) [Logout](#)

Monitor Configuration Diagnose Manage

Configuration > Quick Configuration > Firewall Filters

Quick Configuration Firewall Filters

Firewall Filters

IPv4 Filter Summary Showing Filter 1 to 1 of 1 total. (Page 1 of 1)

Filter Name
X pepsi

Legend

✓ Accept Packet ?	✗ Reject Packet ?	✗ Discard Packet ?	↓ Evaluate Next Term ?
→ Routing Instance ?	📄 Log Packet ?	📄 Syslog Packet ?	⊕ Count Packet ?
📉 Set Packet Loss Priority ?	🏠 Logical Router ?	⚖️ Load Balance Packet ?	🚦 Rate Limit (Police) Packet ?

Any firewall term match conditions that are colored red are considered negated. If a packet matches a negated condition, it is immediately considered not to match the term statement, and the next term in the filter is evaluated. Note that the order of the terms within a firewall filter is significant. Packets are tested against each term in the order in which they are listed in the configuration.

Add New IPv4 Filter

Name

Location ☒ After Final IPv4 Filter ? ☐ After IPv4 Filter ☐ Before IPv4 Filter

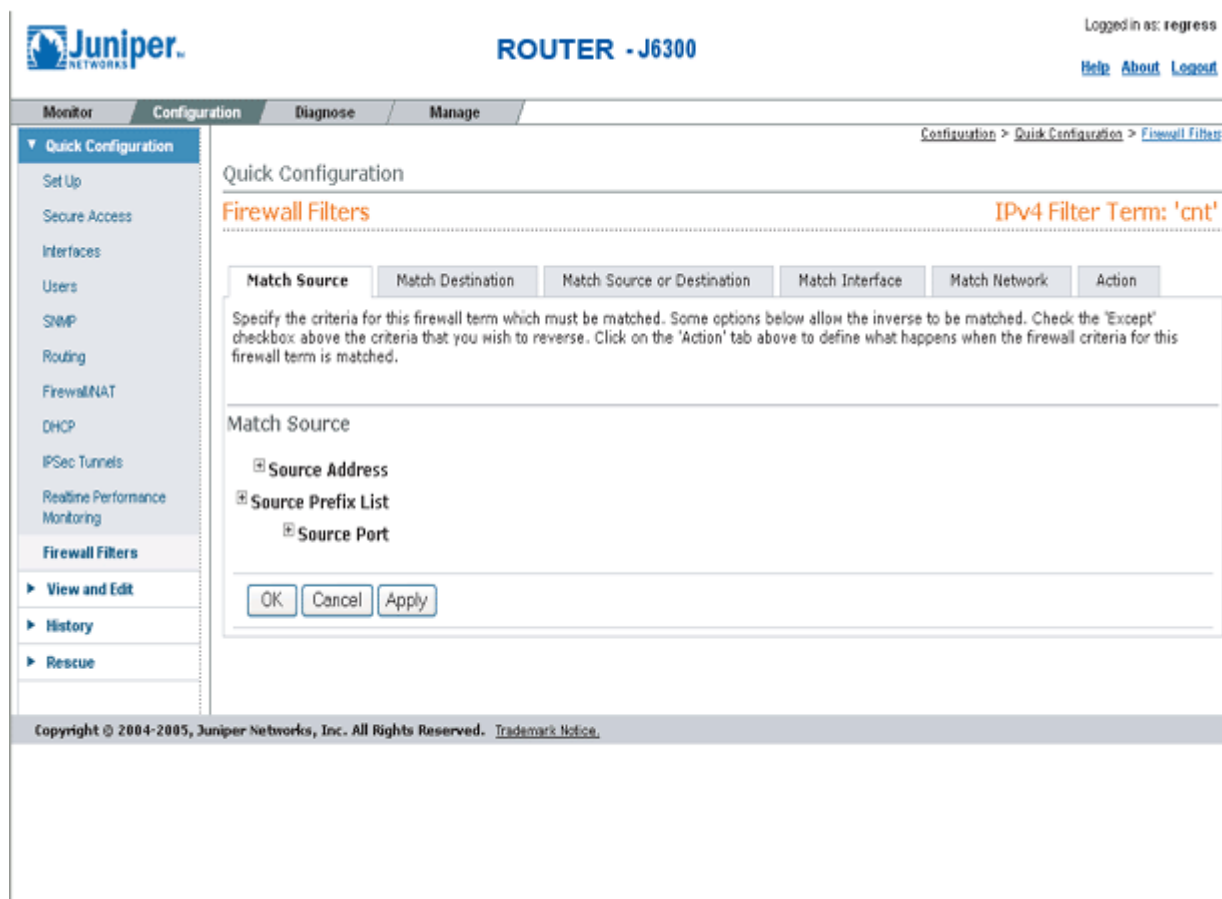
Search

IPv4 Filter Name ?

IPv4 Term Name ?

Number of Items to Display ?

Copyright © 2004-2005, Juniper Networks, Inc. All Rights Reserved. [Trademark Notice](#)

Figure 81: Match Conditions and Actions Quick Configuration Page


To configure a stateless firewall filter with Quick Configuration:

1. In the J-Web interface, select **Configuration > Firewall Filters**.
2. Select one of the following options on the Firewall Filters Quick Configuration page:
 - To edit IPv4 firewall filters and terms, select **Edit IPv4 Firewall Filters**.



NOTE: If you have existing IPv4 firewall configurations in both edit firewall filter and edit firewall family inet filter hierarchies, merge the two to one location. The J-Web Firewall Filter Quick Configuration supports configuration in one location only.

- To edit IPv6 firewall filters and terms, select **Edit IPv6 Firewall Filters**.

3. Enter information into the Firewall Filters Quick Configuration pages, as described in Table 150.
4. Click one of the following buttons on the Firewall Filters Quick Configuration main page:
 - To apply the configuration and stay in the current Firewall Filters 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**.
5. Go on to one of the following procedures:
 - If the stateless firewall filter is not already assigned to an interface, see “Assigning IPv4 and IPv6 Firewall Filters to Interfaces” on page 499.
 - To display the configuration, see “Displaying Firewall Filter Configurations” on page 517.
 - To verify a stateless firewall filter, see “Verifying Firewall Filter Configuration” on page 517.

Table 150: Firewall Filters Quick Configuration Pages Summary

Field	Function	Your Action
IPv4 Filter Summary		
Action column	Displays up and down arrows and a X, allowing you to delete or change the order of a filter or term. The order of an item is important because it determines the order in which corresponding actions are carried out.	<p>To move an item upward, locate the item and click the up arrow from the same row.</p> <p>To move an item downward, locate the item and click the down arrow from the same row.</p> <p>To delete an item, locate the item and click the X from the same row.</p>
Filter Name	<p>Displays the name of the filter and when expanded, lists the terms attached to the filter.</p> <p>Displays the match conditions and actions that are set for each term.</p> <p>Allows you to add more terms to a filter or modify filter terms.</p>	<p>To display the terms added to a filter, click the plus sign next to the filter name. This also displays the match conditions and actions set for the term.</p> <p>To edit a filter, click the filter name. To edit a term, click the name of the term.</p>
Search		

Table 150: Firewall Filters Quick Configuration Pages Summary (continued)

Field	Function	Your Action
Filter Name	Searches for existing filters by filter name.	<p>To find a specific filter, type the name of the filter in the Filter Name box.</p> <p>To list all filters with a common prefix or suffix, use the wildcard character (*) when typing the name of the filter. For example, te* lists all filters with a name starting with the characters <i>te</i>.</p>
Term Name	Searches for existing terms by term name.	<p>To find a specific term, type the name of the term in the Term Name box.</p> <p>To list all terms with a common prefix or suffix, use the wildcard character (*) when typing the name of the term. For example, ra* lists all terms with a name starting with the characters <i>ra</i>.</p>
Number of Items to Display	Specifies the number of filters or terms to display on one page.	To select the number of items to be displayed on one page, select a number from the list.
Add New IPv4 (or IPv6) Filter		
Name	Specifies the name for a new filter.	To name a filter, type a string of meaningful characters or integers that allow you to uniquely identify the filter.
Location	<p>Positions the new filter in one of the following locations:</p> <ul style="list-style-type: none"> ■ After Final IPv4 Filter—At the end of all filters. ■ After IPv4 Filter—After a specified filter. ■ Before IPv4 Filter—Before a specified filter. 	<p>To position the new filter:</p> <ul style="list-style-type: none"> ■ At the end of all filters, select After Final IPv4 Filter. ■ After a specific filter, select After IPv4 Filter then select a name from the filter name list. ■ Before a specific filter, select Before IPv4 Filter then select a name from the filter name list.
Add	<p>Adds a new filter name.</p> <p>Opens the term summary page for this filter allowing you to add new terms to this filter.</p>	To create a new filter and open the term summary page for this filter, click Add .
Add New IPv4 (or IPv6) Term		
Name	Defines a term for a specific filter.	To name a term, type a string of meaningful characters or integers that allow you to uniquely identify the term.

Table 150: Firewall Filters Quick Configuration Pages Summary (continued)

Field	Function	Your Action
Location	Positions the new term in one of the following locations: <ul style="list-style-type: none"> ■ After Final IPv4 Term—At the end of all terms. ■ After IPv4 Term—After a specified term. ■ Before IPv4 Term—Before a specified term. 	To position the new term: <ul style="list-style-type: none"> ■ At the end of all terms, select After Final IPv4 Term. ■ After a specific term, select After IPv4 Term then select a name from the term name list. ■ Before a specific term, select Before IPv4 Term then select a name from the term name list.
Add	Adds a term name for the specific filter. Opens the Filter Term page allowing you to define the match conditions and the action for this term.	To add a term name and open the Filter Term page, click Add .
Match Source		
Source Address	Specifies IP source addresses to be included in, or excluded from, the match condition. Allows you to remove source IP addresses from the match condition. If you have more than 25 addresses, this field displays a link that allows you to easily scroll through pages, change the order of addresses, and also search for them.	To specify an IP source address, type an IP address and prefix length. <ul style="list-style-type: none"> ■ To include the address in the match condition, click Add. ■ To exclude the address from the match condition, select Except then click Add. To remove an IP source address from the match condition, select it and click Delete .
Source Prefix List	Specifies source prefix lists that you have already defined, to be included in the match condition. Allows you to remove a prefix list from the match condition. For information about defining prefix lists, see the <i>JUNOS Policy Framework Configuration Guide</i> .	To include a predefined source prefix list in the match condition, type the prefix list name and click Add . To remove a prefix list from the match condition, select it and click Delete .
Source Port	Specifies the source port type to be included in, or excluded from, the match condition. Allows you to remove a source port type from the match condition. NOTE: This match condition does not check the protocol type being used on the port. Make sure to specify the protocol type (TCP or UDP) match condition in the same term.	To specify a known source port type, select the port from the port name list. To specify source port types that do not exist in the port name list, type the port name, number, or range. <ul style="list-style-type: none"> ■ To include the port in the match condition, click Add. ■ To exclude the port from the match condition, select Except then click Add. To remove a port type from the match condition, select it and click Delete .
Match Destination		

Table 150: Firewall Filters Quick Configuration Pages Summary (continued)

Field	Function	Your Action
Destination Address	<p>Specifies destination addresses to be included in, or excluded from, the match condition.</p> <p>Allows you to remove a destination IP address from the match condition.</p> <p>If you have more than 25 addresses, this field displays a link that allows you to easily scroll through pages, change the order of addresses, and also search for them.</p>	<p>To specify a destination IP address, type an IP address and prefix length.</p> <ul style="list-style-type: none"> ■ To include the address in the match condition, click Add. ■ To exclude the address from the match condition, select Except then click Add. <p>To remove an IP address from the match condition, select it and click Delete.</p>
Destination Prefix List	<p>Specifies destination prefix lists that you have already defined, to be included in the match condition.</p> <p>Allows you to remove a prefix list from the match condition.</p> <p>For information about defining prefix lists, see the <i>JUNOS Policy Framework Configuration Guide</i>.</p>	<p>To include a predefined destination prefix list, type the prefix list name and click Add.</p> <p>To remove a prefix list from the match condition, select it and click Delete.</p>
Destination Port	<p>Specifies destination port types to be included in, or excluded from, the match condition.</p> <p>Allows you to remove a destination port type from the match condition.</p> <p>NOTE: This match condition does not check the protocol type being used on the port. Make sure to specify the protocol type (TCP or UDP) match condition in the same term.</p>	<p>To specify a known destination port type, select the port from the port name list. To specify source port types that do not exist in the port name list, type the port name, number, or range.</p> <ul style="list-style-type: none"> ■ To include the port in the match condition, click Add. ■ To exclude the port from the match condition, select Except then click Add. <p>To remove a destination port type from the match condition, select it and click Delete.</p>
Match Source or Destination		
Address	<p>Specifies IP addresses to be included in, or excluded from, the match condition for a source or destination.</p> <p>Allows you to remove an IP address from the match condition.</p> <p>If you have more than 25 addresses, this field displays a link that allows you to easily scroll through pages, change the order of addresses and also search for them.</p> <p>NOTE: This address match condition cannot be specified in conjunction with the source address or destination address match conditions in the same term.</p>	<p>To specify a source or destination IP address, type the IP address and prefix length.</p> <ul style="list-style-type: none"> ■ To include the address in the match condition, click Add. ■ To exclude the address from the match condition, select Except then click Add. <p>To remove an IP address from the match condition, select it and click Delete.</p>

Table 150: Firewall Filters Quick Configuration Pages Summary (continued)

Field	Function	Your Action
Prefix List	<p>Specifies prefix lists that you have already defined, to be included in the match condition for a source or destination.</p> <p>Allows you to remove a prefix list from the match condition.</p> <p>For information about defining prefix lists, see the <i>JUNOS Policy Framework Configuration Guide</i>.</p> <p>NOTE: This prefix list match condition cannot be specified in conjunction with the source prefix list or destination prefix list match conditions in the same term.</p>	<p>To include a predefined prefix list in the match condition, type the prefix list name and click Add.</p> <p>To remove a prefix list from the match condition, select it and click Delete.</p>
Port	<p>Specifies a port type to be included in, or excluded from, a match condition for a source or destination.</p> <p>Allows you to remove a port from the match condition.</p> <p>NOTE: This match condition does not check the protocol type being used on the port. Make sure to specify the protocol type (TCP or UDP) match condition in the same term.</p> <p>Also, this port match condition cannot be specified in conjunction with the source port or destination port match conditions in the same term.</p>	<p>To specify a known port type in the match condition, select the port from the port name list. To specify port types not included in the port name list, type the port name, number, or range.</p> <ul style="list-style-type: none"> ■ To include the port in the match condition, click Add. ■ To exclude the port from the match condition, select Except then click Add. <p>To remove a port from the match condition, select it and click Delete.</p>
Match Interface		
Interface	<p>Specifies interfaces to be included in a match condition.</p> <p>Allows you to remove an interface from the match condition.</p>	<p>To include an interface in a match condition, either select a name from the interface name list or type the interface name and click Add.</p> <p>To remove an interface from the match condition, select it and click Delete.</p>
Interface Set	<p>Specifies interface sets that you have already defined, to be included in a match condition.</p> <p>Allows you to remove an interface set from the match condition.</p> <p>For information about defining interface sets, see the <i>JUNOS Policy Framework Configuration Guide</i>.</p>	<p>To include a predefined interface set in a match condition, type the interface set name and click Add.</p> <p>To remove an interface set from the match condition, select it and click Delete.</p>

Table 150: Firewall Filters Quick Configuration Pages Summary (continued)

Field	Function	Your Action
Interface Group	<p>Specifies interface groups, that you have already defined, to be included in, or excluded from, a match condition.</p> <p>Allows you to remove an interface group from the match condition.</p> <p>For information about defining interface groups, see the <i>JUNOS Policy Framework Configuration Guide</i>.</p>	<p>To specify a predefined interface group, type the name of the group.</p> <ul style="list-style-type: none"> ■ To include the group in the match condition, click Add. ■ To exclude the group from the match condition, select Except then click Add. <p>To remove an interface group from the match condition, select it and click Delete.</p>
Match Packet and Network		
First Fragment (IPv4 only)	Matches the first fragment of a fragmented packet.	To match the first fragment, select the check box.
Is Fragment (IPv4 only)	Matches trailing fragments (all but the first fragment) of a fragmented packet.	To match trailing fragments, select the check box.
Fragment Flags (IPv4 only)	Specifies fragmentation flags to be included in the match condition.	To specify fragmentation flags, type a text or numeric string defining the flag—for example, more-fragments or 0x2000 .
TCP Established	<p>Matches all TCP packets other than the first packet of a connection.</p> <p>NOTE: This match condition does not verify that the TCP protocol is used on the port. Make sure to specify the TCP protocol as a match condition in the same term.</p>	To match all TCP packets except the first of a connection, select the check box.
TCP Initial	<p>Matches the first TCP packet of a connection.</p> <p>NOTE: This match condition does not verify that the TCP protocol is used on the port. Make sure to specify the TCP protocol as a match condition in the same term.</p>	To match the first TCP packet of a connection, select the check box.
TCP Flags	<p>Specifies TCP flags to be included in the match condition.</p> <p>NOTE: This match condition does not verify that the TCP protocol is used on the port. Make sure to specify the TCP protocol as a match condition in the same term.</p>	To specify a TCP flag, type a text or numeric string defining the flag—for example, syn or 0x02 .
Protocol (IPv4 only)	<p>Specifies IPv4 protocol types to be included in, or excluded from, the match condition.</p> <p>Allows you to remove an IPv4 protocol type from the match condition.</p>	<p>To specify an IPv4 protocol type, select a protocol name from the list or type a protocol name or number—for example, ospf or 89.</p> <ul style="list-style-type: none"> ■ To include the protocol in the match condition, click Add. ■ To exclude the protocol from the match condition, select Except then click Add. <p>To remove an IPv4 protocol type from the match condition, select it and click Delete.</p>

Table 150: Firewall Filters Quick Configuration Pages Summary (continued)

Field	Function	Your Action
Next Header (IPv6 only)	<p>Specifies IPv6 protocol types to be included in, or excluded from, the match condition.</p> <p>Allows you to remove an IPv6 protocol type from the match condition.</p>	<p>To specify an IPv6 protocol type, select a protocol name from the list or type the protocol name or number—for example, igmp or 2.</p> <ul style="list-style-type: none"> ■ To include the protocol in the match condition, click Add. ■ To exclude the protocol from the match condition, select Except then click Add. <p>To remove an IPv6 protocol type from the match condition, select it and click Delete.</p>
ICMP Type	<p>Specifies ICMP packet types to be included in, or excluded from, the match condition.</p> <p>Allows you to remove an ICMP packet type from the match condition.</p> <p>NOTE: This protocol does not verify that ICMP is used on the port. Make sure to specify an ICMP type match condition in the same term.</p>	<p>To specify an ICMP packet type, select a packet type from the list or type a packet type name or number—for example, time-exceeded or 11.</p> <ul style="list-style-type: none"> ■ To include the packet type in the match condition, click Add. ■ To exclude the packet type from the match condition, select Except then click Add. <p>To remove an ICMP packet type from the match condition, select it and click Delete.</p>
ICMP Code	<p>Specifies the ICMP code to be included in, or excluded from, the match condition.</p> <p>Allows you to remove an ICMP code from the match condition.</p> <p>NOTE: The ICMP code is dependent on the ICMP type. Make sure to specify an ICMP type match condition in the same term.</p>	<p>To specify an ICMP code, select a packet code from the list or type the packet code as text or a number—for example, ip-header-bad or 0.</p> <ul style="list-style-type: none"> ■ To include the ICMP code in the match condition, click Add. ■ To exclude the ICMP code from the match condition, select Except then click Add. <p>To remove an ICMP code from the match condition, select it and click Delete.</p>
Traffic Class (IPv6 only)	<p>Specifies Differentiated Services code points (DSCPs) to be included in, or excluded from, the match condition.</p> <p>Allows you to remove a DSCP value from the match condition.</p> <p>For information about DSCPs, see “Class-of-Service Overview” on page 449.</p>	<p>To specify a DSCP, select it from the list or type the DSCP value as a keyword, decimal, or binary string—for example, af11 or 10.</p> <ul style="list-style-type: none"> ■ To include the DSCP in the match condition, click Add. ■ To exclude the DSCP from the match condition, select Except then click Add. <p>To remove a DSCP from the match condition, select it and click Delete.</p>

Table 150: Firewall Filters Quick Configuration Pages Summary (continued)

Field	Function	Your Action
Fragment Offset (IPv4 only)	<p>Specifies the fragment offset value to be included in, or excluded from, the match condition. The fragment offset value specifies the location of the fragment in the packet. For example, fragment offset zero specifies the first fragment.</p> <p>Allows you to remove a fragment offset value from the match condition.</p>	<p>To specify a fragment offset value, type the fragment offset number or range.</p> <ul style="list-style-type: none"> ■ To include the offset in the match condition, click Add. ■ To exclude the offset from the match condition, select Except then click Add. <p>To remove a fragment offset value from the match condition, select it and click Delete.</p>
Precedence (IPv4 only)	<p>Specifies IP precedences to be included in, or excluded from, the match condition.</p> <p>Allows you to remove an IP precedence entry from the match condition.</p>	<p>To specify an IP precedence, select it from the list or type the precedence as a keyword, decimal integer between 0 and 7, or binary string.</p> <ul style="list-style-type: none"> ■ To include the precedence in the match condition, click Add. ■ To exclude the precedence from the match condition, select Except then click Add. <p>To remove an IP precedence from the match condition, select it and click Delete.</p>
DSCP (IPv4 only)	<p>Specifies Differentiated Services code points (DSCPs) to be included in, or excluded from, the match condition</p> <p>Allows you to remove a DSCP entry from the match condition.</p>	<p>To specify a DSCP, select it from the list or type the DSCP value as a keyword, decimal, or binary string—for example, af11 or 10.</p> <ul style="list-style-type: none"> ■ To include the DSCP in the match condition, click Add. ■ To exclude the DSCP from the match condition, select Except then click Add. <p>To remove a DSCP, select it and click Delete.</p>
TTL (IPv4 only)	<p>Specifies the IPv4 time-to-live (TTL) value to be included in, or excluded from, the match condition.</p> <p>Allows you to remove an IPv4 TTL value from the match condition.</p>	<p>To specify an IPv4 TTL value, type a number between 1 and 255.</p> <ul style="list-style-type: none"> ■ To include the TTL in the match condition, click Add. ■ To exclude the TTL from the match condition, select Except then click Add. <p>To remove an IPv4 TTL type from the match condition, select it and click Delete.</p>

Table 150: Firewall Filters Quick Configuration Pages Summary (continued)

Field	Function	Your Action
Packet Length	<p>Specifies the length of received packets, in bytes, to be included in, or excluded from, the match condition.</p> <p>Allows you to remove a packet length value from the match condition.</p>	<p>To specify a packet length, type a value or range.</p> <ul style="list-style-type: none"> ■ To include the packet length in the match condition, click Add. ■ To exclude the packet length from the match condition, select Except then click Add. <p>To remove a packet length value from the match condition, select it and click Delete.</p>
Forwarding Class	<p>Specifies forwarding classes to be included in, or excluded from, the match condition.</p> <p>Allows you to remove a forwarding class entry from the match condition.</p> <p>For information about forwarding classes, see “Class-of-Service Overview” on page 449.</p>	<p>To specify a forwarding class, select it from the list or type it.</p> <ul style="list-style-type: none"> ■ To include the forwarding class in the match condition, click Add. ■ To exclude the forwarding class from the match condition, select Except then click Add. <p>To remove a forwarding class from the match condition, select it and click Delete.</p>
IP Options (IPv4 only)	<p>Specifies IP options to be included in, or excluded from, the match condition.</p> <p>Allows you to remove an IP option from the match condition.</p>	<p>To specify an IP option, select it from the list or type a text or numeric string identifying the option.</p> <ul style="list-style-type: none"> ■ To include the IP option in the match condition, click Add. ■ To exclude the IP option from the match condition, select Except then click Add. <p>To remove an IP option from the match condition, select it and click Delete.</p>
IPSec ESP SPI (IPv4 only)	<p>Specifies IPSec Encapsulating Security Payload (ESP) Security Parameter Index (SPI) values to be included in, or excluded from, the match condition.</p> <p>Allows you to remove an ESP SPI value from the match condition.</p>	<p>To specify an ESP SPI value, type a binary, hexadecimal, or decimal SPI value or range.</p> <ul style="list-style-type: none"> ■ To include the value in the match condition, click Add. ■ To exclude the value from the match condition, select Except then click Add. <p>To remove an ESP SPI value from the match condition, select it and click Delete.</p>
Action		
Nothing	<p>No action is performed. By default, a packet is accepted if it meets the match conditions of the term, and packets that do not match any conditions in the firewall filter are dropped.</p>	<p>To specify no action (or the default action), select Nothing.</p>

Table 150: Firewall Filters Quick Configuration Pages Summary (continued)

Field	Function	Your Action
Accept	Accepts the packet if it meets the match conditions of the term.	To accept the packet, select Accept .
Discard	Discards the packet if it meets the match conditions of the term. Names a discard collector for packets.	To discard a packet, select Discard . To name a discard collector, type a filename in the Accounting box.
Reject	Rejects the packet and returns a rejection message if the packet meets the match conditions in the term. Allows you to specify a message type that denotes the reason the packet was rejected. NOTE: To log and sample rejected packets, specify Log and Sample action modifiers in conjunction with this action.	To reject a packet, select Reject . To specify a message type, select the message from the Reason list.
Next Term	Evaluates the packet with the next term in the filter if the packet meets the match conditions in this term. This action makes sure that the next term is used for evaluation even when the packet matches the conditions of a term. When this action is not specified, the filter stops evaluating the packet after it matches the conditions of a term, and takes the associated action.	To continue to the next term, select Next Term .
Routing Instance	Accepts the packet and forwards it to the specified routing instance if the packet meets the match conditions of the term.	To specify a routing instance, select Routing Instance and type the routing instance name in the box next to Routing Instance.
Action Modifiers		
Forwarding Class	Classifies the packet as a specific forwarding class. For information about forwarding classes, see “Class-of-Service Overview” on page 449.	To specify a forwarding class, select it from the list.
Count	Counts the packets passing this term. Allows you to name a counter, which is specific to this filter. This means that every time a packet transits any interface that uses this filter, it increments the specified counter.	To count packets passing this term, select Count . To specify a counter name, type a 24-character string containing letters, numbers, or hyphens.
Virtual Channel (IPv4 only)	Specifies the virtual channel to be set on a particular logical interface.	To specify the virtual channel, type a string identifying the virtual channel.
Log	Logs the packet header information in the Routing Engine.	To log packet header information, select Log .
Syslog	Records packet information in the system log.	To record information in the system log, select Syslog .

Table 150: Firewall Filters Quick Configuration Pages Summary (continued)


Field	Function	Your Action
Sample (IPv4 only)	<p>Samples traffic on the interface.</p> <p>NOTE: Make sure to enable traffic sampling for this action to work. For more information, see the <i>JUNOS Policy Framework Configuration Guide</i>.</p>	To sample traffic on an interface, select Sample .
Loss Priority	<p>Sets the loss priority of the packet. This is the priority of dropping a packet before it is sent. The loss priority affects the scheduling priority of the packet.</p> <p>For more information, see the <i>JUNOS Class of Service Configuration Guide</i>.</p>	To set the loss priority of the packet, select a loss priority from the list.

Assigning IPv4 and IPv6 Firewall Filters to Interfaces

For a firewall filter to work, you must assign it to an interface. Use the Firewall Filters Quick Configuration pages to assign IPv4 and IPv6 filters to interfaces. Using these pages you can select a firewall filter to evaluate packets that are received or transmitted on a specific interface.

When assigning firewall filters to interfaces, remember that you can assign only one input and one output firewall filter to each interface. However, you can assign the same filter to multiple interfaces.

Figure 82 shows the Firewall Filters Quick Configuration page that displays the Services Router interfaces available for filter assignment and the status of existing filter assignments.

Figure 82: Firewall Filters Interface Assignment Quick Configuration Page

To assign IPv4 and IPv6 firewall filters to interfaces with Quick Configuration:

1. In the J-Web interface, select **Configuration > Firewall Filters > Assign Firewall Filters to Interfaces**.
2. Enter information into the Firewall Filters Quick Configuration pages, as described in Table 151.
3. Click one of the following buttons on the Firewall Filters Quick Configuration main page:
 - To apply the configuration and stay in current the Firewall Filters 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**.
4. Go on to one of the following procedures:
 - To display the configuration, see “Displaying Firewall Filter Configurations” on page 517.
 - To verify a stateless firewall filter, see “Verifying Firewall Filter Configuration” on page 517.

Table 151: Assigning Firewall Filters Quick Configuration Pages Summary

Field	Function	Your Action
Firewall Filters		
Logical Interface Name	Displays the logical interfaces on a router. Allows you to apply IPv4 and IPv6 firewall filters to packets received on the interface and packets transmitted from the interface.	To apply firewall filters to an interface, click the interface name <ul style="list-style-type: none">■ To apply an input firewall filter, follow instructions in the input firewall filters section.■ To apply an output firewall filter, follow instructions in the ouput firewall filters section.
Link State	Displays the status of the logical interface.	None.
Input Firewall Filters	Displays the input firewall filter applied on an interface. This filter evaluates all packets received on the interface.	None.
Output Firewall Filters	Displays the output firewall filter applied on an interface. This filter evaluates all packets transmitted from the interface.	None.
Input Firewall Filters		
IPv4 Input Filter	Allows you to apply an input firewall filter to an interface. This filter evaluates all packets received on the interface.	To apply an input firewall filter to an interface, select the name of the firewall filter from the list.
IPv6 Input Filter		
Output Firewall Filters		
IPv4 Output Filter	Allows you to apply an output firewall filter to an interface. This filter evaluates all packets transmitted on the interface.	To apply an output firewall filter to an interface, select the name of the firewall filter from the list.
IPv6 Output Filter		

Configuring a Stateless Firewall Filter with a Configuration Editor

The section contains the following topics. For stateless firewall match conditions, actions, and modifiers, see “Stateless Firewall Filter Match Conditions, Actions, and Action Modifiers” on page 445.

- Stateless Firewall Filter Strategies on page 502
- Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources on page 502
- Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods on page 506
- Configuring a Routing Engine Firewall Filter to Handle Fragments on page 511
- Applying a Stateless Firewall Filter to an Interface on page 516

Stateless Firewall Filter Strategies

For best results, use the following sections to plan the purpose and contents of a stateless firewall filter before starting configuration.



CAUTION: If a packet does not match any terms in a stateless firewall filter rule, the packet is discarded. Take care that you do not configure a firewall filter that prevents you from accessing the Services Router after you commit the configuration. For example, if you configure a firewall filter that does not match HTTP or HTTPS packets, you cannot access the router with the J-Web interface.

Strategy for a Typical Stateless Firewall Filter

A primary goal of a typical stateless firewall filter is to protect the Routing Engine processes and resources from malicious or untrusted packets. You can configure a stateless firewall filter like the sample filter `protect-RE` to restrict traffic destined for the Routing Engine based on its source, protocol, and application. In addition, you can limit the traffic rate of packets destined for the Routing Engine to protect against flood, or *denial-of-service* (DoS), attacks.

For details, see “Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources” on page 502 and “Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods” on page 506.

Strategy for Handling Packet Fragments

You can configure a stateless firewall filter like the sample filter `fragment-filter` to address special circumstances associated with fragmented packets destined for the Routing Engine. Because the Services Router evaluates every packet against a firewall filter (including fragments), you must configure the filter to accommodate fragments that do not contain packet header information. Otherwise, the filter discards all but the first fragment of a fragmented packet.

For details, see “Configuring a Routing Engine Firewall Filter to Handle Fragments” on page 511.

Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources

The following example shows how to create a stateless firewall filter, `protect-RE`, that discards all traffic destined for the Routing Engine, except SSH and BGP protocol packets from specified trusted sources. Table 152 lists the terms that are configured in this sample filter.

Table 152: Sample Stateless Firewall Filter protect-RE Terms to Allow Packets from Trusted Sources

Term	Purpose
ssh-term	Accepts TCP packets with a source address of 192.168.122.0/24 and a destination port that specifies SSH.
bgp-term	Accepts TCP packets with a source address of 10.2.1.0/24 and a destination port that specifies the BGP protocol.
discard-rest-term	For all packets that are not accepted by ssh-term or bgp-term , creates a firewall filter log and system logging records, then discards all packets. To view the log, enter the show firewall log operational mode command. (For more information, see “Displaying Firewall Filter Logs” on page 523.)

By applying firewall filter **protect-RE** to the Routing Engine, you specify which protocols and services, or applications, are allowed to reach the Routing Engine, and you ensure the packets are from a trusted source. This protects processes running on the Routing Engine from an external attack.

To use the configuration editor to configure the stateless firewall filter:

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 153.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To display the configuration, see “Displaying Firewall Filter Configurations” on page 517.
 - To apply the firewall filter to the Routing Engine, see “Applying a Stateless Firewall Filter to an Interface” on page 516.
 - To verify the firewall filter, see “Verifying a Services, Protocols, and Trusted Sources Firewall Filter” on page 525.

Table 153: Configuring a Protocols and Services Firewall Filter for the Routing Engine

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Firewall level in the configuration hierarchy.	In the configuration editor hierarchy, select Firewall .	From the top of the configuration hierarchy, enter edit firewall .

Table 153: Configuring a Protocols and Services Firewall Filter for the Routing Engine (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define protect-RE and ssh-term , and define the protocol, destination port, and source address match conditions.	<ol style="list-style-type: none"> Next to Filter, click Add new entry. In the Filter name box, type protect-RE. Next to Term, click Add New Entry. In the Rule name box, type ssh-term. Next to From, click Configure. In the Protocol choice list, select Protocol. Next to Protocol, click Add new entry. In the Value keyword list, select tcp. Click OK. In the Destination port choice list, select Destination port. Next to Destination port, click Add new entry. In the Value keyword list, select ssh. Click OK. Next to Source address, click Add new entry. In the Address box, type 192.168.122.0/24. Click OK twice. 	<p>Set the term name and define the match conditions:</p> <pre>set family inet filter protect-RE term ssh-term from protocol tcp destination-port ssh source-address 192.168.122.0/24</pre>
Define the actions for ssh-term .	<ol style="list-style-type: none"> On the Term ssh-term page, next to Then, click Configure. In the Designation list, select Accept. Click OK twice. 	<p>Set the actions:</p> <pre>set family inet filter protect-RE term ssh-term then accept</pre>

Table 153: Configuring a Protocols and Services Firewall Filter for the Routing Engine (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define bgp-term , and define the protocol, destination port, and source address match conditions.	<ol style="list-style-type: none"> On the Filter protect-RE page, next to Term, click Add New Entry. In the Rule name box, type bgp-term. Next to From, click Configure. In the Protocol choice list, select Protocol. Next to Protocol, click Add new entry. In the Value keyword list, select tcp. Click OK. In the Destination port choice list, select Destination port. Next to Destination port, click Add new entry. In the Value keyword list, select bgp. Click OK. Next to Source address, click Add new entry. In the Address box, type 10.2.1.0/24. Click OK twice. 	<p>Set the term name and define the match conditions:</p> <pre>set family inet filter protect-RE term bgp-term from protocol tcp destination-port bgp source-address 10.2.1.0/24</pre>
Define the action for bgp-term .	<ol style="list-style-type: none"> On the Term bgp-term page, next to Then, click Configure. In the Designation list, select Accept. Click OK twice. 	<p>Set the action:</p> <pre>set family inet filter protect-RE term bgp-term then accept</pre>
Define discard-rest-term and its action.	<ol style="list-style-type: none"> On the Filter protect-RE page, next to Term, click Add New Entry. In the Rule name box, type discard-rest-term. Next to Then, click Configure. Next to Log, select the check box. Next to Syslog, select the check box. In the Designation list, select Discard. Click OK four times. 	<p>Set the term name and define its actions:</p> <pre>set family inet filter protect-RE term discard-rest-term then log syslog discard</pre>

Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods

The procedure in this section creates a sample stateless firewall filter, `protect-RE`, that limits certain TCP and ICMP traffic destined for the Routing Engine. A router without this kind of protection is vulnerable to TCP and ICMP flood attacks—also called denial-of-service (DoS) attacks. For example:

- A TCP flood attack of SYN packets initiating connection requests can so overwhelm the Services Router that it can no longer process legitimate connection requests, resulting in denial of service.
- An ICMP flood can overload the Services Router with so many echo requests (ping requests) that it expends all its resources responding and can no longer process valid network traffic, also resulting in denial of service.

Applying a firewall filter like `protect-RE` to the Routing Engine protects against these types of attacks.

For each term in the sample filter, you first create a policer and then incorporate it into the action of the term. For more information about firewall filter policers, see the *JUNOS Policy Framework Configuration Guide*.

If you want to include the terms created in this procedure in the `protect-RE` firewall filter configured in the previous section (see “Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources” on page 502), perform the configuration tasks in this section first, then configure the terms as described in the previous section. This approach ensures that the rate-limiting terms are included as the first two terms in the firewall filter.



NOTE: You can move terms within a firewall filter by using the `insert` CLI command. For more information, see “Inserting an Identifier” on page 28.

Table 154 lists the terms that are configured in this sample filter.

Table 154: Sample Stateless Firewall Filter protect-RE Terms to Protect Against Floods

Term	Purpose	Policer
tcp-connection-term	<p>Polices the following types of TCP packets with a source address of 192.168.122.0/24 or 10.2.1.0/24:</p> <ul style="list-style-type: none"> ■ Connection request packets (SYN and ACK flag bits equal 1 and 0) ■ Connection release packets (FIN flag bit equals 1) ■ Connection reset packets (RST flag bit equals 1) 	tcp-connection-policer—Limits the traffic rate and burst size of these TCP packets to 500,000 bps and 15,000 bytes. Packets that exceed the traffic rate are discarded.
icmp-term	<p>Polices the following types of ICMP packets. All are counted in counter icmp-counter.</p> <ul style="list-style-type: none"> ■ Echo request packets ■ Echo response packets ■ Unreachable packets ■ Time-exceeded packets 	icmp-policer—Limits the traffic rate and burst size of these ICMP packets to 1,000,000 bps and 15,000 bytes. Packets that exceed the traffic rate are discarded.

To use the configuration editor to configure the policers and the stateless firewall filter:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. To configure the firewall filter policers, perform the configuration tasks described in Table 155.
3. To configure the prefix lists and the firewall filter, perform the configuration tasks described in Table 156.
4. If you are finished configuring the router, commit the configuration.
5. Go on to one of the following procedures:
 - To display the configuration, see “Displaying Firewall Filter Configurations” on page 517.
 - To apply the firewall filter to the Routing Engine, see “Applying a Stateless Firewall Filter to an Interface” on page 516.
 - To verify the firewall filter, see “Verifying a TCP and ICMP Flood Firewall Filter” on page 526.

Table 155: Configuring Policers for TCP and ICMP

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Firewall level in the configuration hierarchy.	In the configuration editor hierarchy, select Firewall .	From the top of the configuration hierarchy, enter edit firewall .
<p>Define tcp-connection-policer and set its rate limits.</p> <p>The burst size limit can be from 1,500 bytes through 100,000,000 bytes.</p> <p>The bandwidth limit can be from 32,000 bps through 32,000,000,000 bps.</p> <p>Use the following abbreviations when specifying these limits:</p> <ul style="list-style-type: none"> ■ k (1000) ■ m (1,000,000) ■ g (1,000,000,000) 	<ol style="list-style-type: none"> Next to Policer, click Add new entry. In the Policer name box, type tcp-connection-policer. Next to Filter specific, select the check box. Next to If Exceeding, select the check box and click Configure. In the Burst size limit box, type 15k. In the Bandwidth list, select Bandwidth limit. In the Bandwidth limit box, type 500k. Click OK. 	<p>Set the policer name and its rate limits:</p> <pre>set policer tcp-connection-policer filter-specific if-exceeding burst-size-limit 15k bandwidth-limit 500k</pre>
Define the policer action for tcp-connection-policer .	<ol style="list-style-type: none"> On the Policer tcp-connection-policer page, next to Then, click Configure. Next to Discard, select the check box. Click OK twice. 	<p>Set the policer action:</p> <pre>set policer tcp-connection-policer then discard</pre>

Table 155: Configuring Policers for TCP and ICMP (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define <code>icmp-policer</code> and set its rate limits.	1. On the Firewall page, next to Policer, click Add new entry .	Set the policer name and its rate limits:
The burst size limit can be from 1,500 bytes through 100,000,000 bytes.	2. In the Policer name box, type <code>icmp-policer</code> .	<code>set policer icmp-policer filter-specific</code>
The bandwidth limit can be from 32,000 bps through 32,000,000,000 bps.	3. Next to Filter specific, select the check box.	<code>if-exceeding burst-size-limit 15k</code>
Use the following abbreviations when specifying these limits:	4. Next to If Exceeding, select the check box and click Configure .	<code>bandwidth-limit 1m</code>
<ul style="list-style-type: none"> ■ k (1000) ■ m (1,000,000) ■ g (1,000,000,000) 	5. In the Burst size limit box, type 15k .	
	6. In the Bandwidth list, select Bandwidth limit .	
	7. In the Bandwidth limit box, type 1m .	
	8. Click OK .	
Define the policer action for <code>icmp-policer</code> .	1. On the Policer <code>icmp-policer</code> page, next to Then, click Configure .	Set the policer action:
	2. Next to Discard, select the check box.	<code>set policer icmp-policer then discard</code>
	3. Click OK three times.	

Table 156: Configuring a TCP and ICMP Flood Firewall Filter for the Routing Engine

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Policy options level in the configuration hierarchy.	In the configuration editor hierarchy, select Policy options .	From the top of the configuration hierarchy, enter <code>edit policy-options</code> .
Define the prefix list <code>trusted-addresses</code> .	1. Next to Prefix list, click Add new entry .	Set the prefix list:
	2. In the Name box, type <code>trusted-addresses</code> .	<code>set prefix-list trusted-addresses</code>
	3. Next to Prefix list item, click Add new entry .	<code>192.168.122.0/24</code>
	4. In the Prefix box, type <code>192.168.122.0/24</code> .	<code>set prefix-list trusted-addresses 10.2.1.0/24</code>
	5. Click OK .	
	6. Next to Prefix list item, click Add new entry .	
	7. In the Prefix box, type <code>10.2.1.0/24</code> .	
	8. Click OK three times.	

Table 156: Configuring a TCP and ICMP Flood Firewall Filter for the Routing Engine (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Firewall level in the configuration hierarchy.	In the configuration editor hierarchy, select Firewall .	From the top of the configuration hierarchy, enter <code>edit firewall</code> .
Define protect-RE and tcp-connection-term , and define the source prefix list match condition.	<ol style="list-style-type: none"> Next to Filter, click Add new entry. In the Filter name box, type protect-RE. Next to Term, click Add New Entry. In the Rule name box, type tcp-connection-term. Next to From, click Configure. Next to Source prefix list, click Add new entry. In the Name box, type trusted-addresses. Click OK. 	<p>Set the term name and define the source address match condition:</p> <pre>set family inet filter protect-RE term tcp-connection-term from source-prefix-list trusted-addresses</pre>
Define the TCP flags and protocol match conditions for tcp-connection-term .	<ol style="list-style-type: none"> In the TCP flags box, type (syn & !ack) fin rst. In the Protocol choice list, select Protocol. Next to Protocol, click Add new entry. In the Value keyword list, select tcp. Click OK. 	<p>Set the TCP flags and protocol and protocol match conditions for the term:</p> <pre>set family inet filter protect-RE term tcp-connection-term from protocol tcp tcp-flags "(syn & !ack) fin rst"</pre>
Define the actions for tcp-connection-term .	<ol style="list-style-type: none"> On the Term tcp-connection-term page, next to Then, click Configure. In the Policer box, type tcp-connection-policer. In the Designation list, select Accept. Click OK twice. 	<p>Set the actions:</p> <pre>set family inet filter protect-RE term tcp-connection-term then policer tcp-connection-policer accept</pre>
Define icmp-term , and define the protocol.	<ol style="list-style-type: none"> On the Filter protect-RE page, next to Term, click Add New Entry. In the Rule name box, type icmp-term. Next to From, click Configure. In the Protocol choice list, select Protocol. Next to Protocol, click Add new entry. In the Value keyword list, select icmp. Click OK. 	<p>Set the term name and define the protocol:</p> <pre>set family inet filter protect-RE term icmp-term from protocol icmp</pre>

Table 156: Configuring a TCP and ICMP Flood Firewall Filter for the Routing Engine (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define the ICMP type match conditions.	<ol style="list-style-type: none"> 1. In the <code>Icmp</code> type choice list, select Icmp type. 2. Next to <code>Icmp</code> type, click Add new entry. 3. In the Value keyword list, select echo-request. 4. Click OK. 5. Next to <code>Icmp</code> type, click Add new entry. 6. In the Value keyword list, select echo-reply. 7. Click OK. 8. Next to <code>Icmp</code> type, click Add new entry. 9. In the Value keyword list, select unreachable. 10. Click OK. 11. Next to <code>Icmp</code> type, click Add new entry. 12. In the Value keyword list, select time-exceeded. 13. Click OK. 	<p>Set the ICMP type match conditions:</p> <pre>set family inet filter protect-RE term icmp-term from icmp-type [echo-request echo-reply unreachable time-exceeded]</pre>
Define the actions for <code>icmp-term</code> .	<ol style="list-style-type: none"> 1. On the <code>icmp-term</code> page, next to Then, click Configure. 2. In the Count box, type <code>icmp-counter</code>. 3. In the Policer box, type <code>icmp-policer</code>. 4. In the Designation list, select Accept. 5. Click OK four times. 	<p>Set the actions:</p> <pre>set family inet filter protect-RE term icmp-term then policer icmp-policer count icmp-counter accept</pre>

Configuring a Routing Engine Firewall Filter to Handle Fragments

The procedure in this section creates a sample stateless firewall filter, `fragment-RE`, that handles fragmented packets destined for the Routing Engine. By applying `fragment-RE` to the Routing Engine, you protect against the use of IP fragmentation as a means to disguise TCP packets from a firewall filter.

Table 157 lists the terms that are configured in this sample filter.

Table 157: Sample Stateless Firewall Filter fragment-RE Terms

Term	Purpose
small-offset-term	Discards IP packets with a fragment offset of 1 through 5, and adds a record to the system logging facility.
not-fragmented-term	Accepts unfragmented TCP packets with a source address of 10.2.1.0/24 and a destination port that specifies the BGP protocol. A packet is considered unfragmented if its MF flag and its fragment offset in the TCP header equal 0.
first-fragment-term	Accepts the first fragment of a fragmented TCP packet with a source address of 10.2.1.0/24 and a destination port that specifies the BGP protocol.
fragment-term	Accepts all packet fragments with an offset of 6 through 8191.

For example, consider an IP packet that is fragmented into the smallest allowable fragment size of 8 bytes (a 20-byte IP header plus an 8-byte payload). If this IP packet carries a TCP packet, the first fragment (fragment offset of 0) that arrives at the Services Router contains only the TCP source and destination ports (first 4 bytes), and the sequence number (next 4 bytes). The TCP flags, which are contained in the next 8 bytes of the TCP header, arrive in the second fragment (fragment offset of 1). The fragment-RE filter works as follows:

- Term **small-offset-term** discards small offset packets to ensure that subsequent terms in the firewall filter can be matched against all the headers in the packet.
- Term **fragment-term** accepts all fragments that were not discarded by **small-offset-term**. However, only those fragments that are part of a packet containing a first fragment accepted by **first-fragment-term** are reassembled by the Services Router.

For more information about IP fragment filtering, see RFC 1858, *Security Considerations for IP Fragment Filtering*.

To use the configuration editor to configure the stateless firewall filter:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. To configure the firewall filter, perform the configuration tasks described in Table 158.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:
 - To display the configuration, see “Displaying Firewall Filter Configurations” on page 517.
 - To apply the firewall filter to the Routing Engine, see “Applying a Stateless Firewall Filter to an Interface” on page 516.
 - To verify the firewall filter, see “Verifying a Firewall Filter That Handles Fragments” on page 527.

Table 158: Configuring a Fragments Firewall Filter for the Routing Engine

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Firewall level in the configuration hierarchy.	In the configuration editor hierarchy, select Firewall .	From the top of the configuration hierarchy, enter edit firewall .
Define fragment-RE and small-offset-term , and define the fragment offset match condition. The fragment offset can be from 1 through 8191.	<ol style="list-style-type: none"> Next to Filter, click Add new entry. In the Filter name box, type fragment-RE. Next to Term, click Add New Entry. In the Rule name box, type small-offset-term. Next to From, click Configure. In the Fragment offset choice list, select Fragment offset. Next to Fragment offset, select Add New Entry. In the Range box, type 1-5. Click OK twice. 	<p>Set the term name and define the fragment offset match condition:</p> <pre>set family inet filter fragment-RE term small-offset-term from fragment-offset 1-5</pre>
Define the action for small-offset-term .	<ol style="list-style-type: none"> On the Term small-offset-term page, next to Then, click Configure. Next to Syslog, select the check box. In the Designation list, select Discard. Click OK twice. 	<p>Set the action:</p> <pre>set family inet filter fragment-RE term small-offset-term then syslog discard</pre>

Table 158: Configuring a Fragments Firewall Filter for the Routing Engine (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define not-fragmented-term , and define the fragment, protocol, destination port, and source address match conditions.	<ol style="list-style-type: none"> On the Filter fragment-RE page, next to Term, click Add New Entry. In the Term name box, type not-fragmented-term. Next to From, click Configure. In the Fragment flags box, type 0x0. In the Fragment offset choice list, select Fragment offset. Next to Fragment offset, select Add New Entry. In the Range box, type 0. Click OK. In the Protocol choice list, select Protocol. Next to Protocol, click Add new entry. In the Value keyword list, select tcp. Click OK. In the Destination port choice list, select Destination port. Next to Destination port, click Add new entry. In the Value keyword list, select bgp. Click OK. Next to Source address, click Add new entry. In the Address box, type 10.2.1.0/24. Click OK twice. 	<p>Set the term name and define match conditions:</p> <pre>set family inet filter fragment-RE term not-fragmented-term from fragment-flags 0x0 fragment-offset 0 protocol tcp destination-port bgp source-address 10.2.1.0/24</pre>
Define the action for not-fragmented-term .	<ol style="list-style-type: none"> On the Term not-fragmented-term page, next to Then, click Configure. In the Designation list, select Accept. Click OK twice. 	<p>Set the action:</p> <pre>set family inet filter fragment-RE term not-fragmented-term then accept</pre>

Table 158: Configuring a Fragments Firewall Filter for the Routing Engine (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define <code>first-fragment-term</code> , and define the fragment, protocol, destination port, and source address match conditions.	<ol style="list-style-type: none"> 1. On the Filter <code>fragment-RE</code> page, next to Term, click Add New Entry. 2. In the Rule name box, type <code>first-fragment-term</code>. 3. Next to From, click Configure. 4. Next to First fragment, select the check box. 5. In the Protocol choice list, select Protocol. 6. Next to Protocol, click Add new entry. 7. In the Value keyword list, select tcp. 8. Click OK. 9. In the Destination port choice list, select Destination port. 10. Next to Destination port, click Add new entry. 11. In the Value keyword list, select bgp. 12. Click OK. 13. Next to Source address, click Add new entry. 14. In the Address box, type <code>10.2.1.0/24</code>. 15. Click OK twice. 	<p>Set the term name and define match conditions:</p> <pre>set family inet filter fragment-RE term first-fragment-term from first-fragment protocol tcp destination-port bgp source-address 10.2.1.0/24</pre>
Define the action for <code>first-fragment-term</code> .	<ol style="list-style-type: none"> 1. On the Term <code>first-fragment-term</code> page, next to Then, click Configure. 2. In the Designation list, select Accept. 3. Click OK twice. 	<p>Set the action:</p> <pre>set family inet filter fragment-RE term first-fragment-term then accept</pre>

Table 158: Configuring a Fragments Firewall Filter for the Routing Engine (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define <code>fragment-term</code> and define the fragment match condition.	<ol style="list-style-type: none"> 1. On the Filter <code>fragment-RE</code> page, next to Term, click Add New Entry. 2. In the Rule name box, type <code>fragment-term</code>. 3. Next to From, click Configure. 4. In the Fragment offset choice list, select Fragment offset. 5. Next to Fragment offset, select Add New Entry. 6. In the Range box, type <code>6-8191</code>. 7. Click OK twice. 	Set the term name and define match conditions: <pre>set family inet filter fragment-RE term fragment-term from fragment-offset 6-8191</pre>
Define the action for <code>fragment-term</code> .	<ol style="list-style-type: none"> 1. On the Term <code>fragment-term</code> page, next to Then, click Configure. 2. In the Designation list, select Accept. 3. Click OK four times. 	Set the action: <pre>set family inet filter fragment-RE term fragment-term then accept</pre>

Applying a Stateless Firewall Filter to an Interface

You can apply a stateless firewall to the input or output sides, or both, of an interface. To filter packets transiting the router, apply the firewall filter to any non-Routing Engine interface. To filter packets originating from, or destined for, the Routing Engine, apply the firewall filter to the loopback (lo0) interface.

For example, to apply a stateless firewall filter `protect-RE` to the input side of the Routing Engine interface, follow this procedure:

1. Perform the configuration tasks described in Table 159.
2. If you are finished configuring the router, commit the configuration.

Table 159: Applying a Firewall Filter to the Routing Engine Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Inet level in the configuration hierarchy.	In the configuration editor hierarchy, select Interfaces > lo0 > Unit > 0 > Family > Inet .	From the top of the configuration hierarchy, apply the filter to the interface:
Apply protect-RE as an input filter to the lo0 interface.	<ol style="list-style-type: none"> Next to Filter, click Configure. In the Input box, type protect-RE. Click OK five times. 	set interfaces lo0 unit 0 family inet filter input protect-RE

To view the configuration of the Routing Engine interface, enter the `show interfaces lo0` command. For example:

```
user@host# show interfaces lo0
unit 0 {
    family inet {
        filter {
            input protect-RE;
        }
        address 127.0.0.1/32;
    }
}
```

Verifying Firewall Filter Configuration

To verify a firewall filter configuration, perform these tasks:

- Displaying Firewall Filter Configurations on page 517
- Verifying a Stateful Firewall Filter on page 522
- Displaying Firewall Filter Logs on page 523
- Displaying Firewall Filter Statistics on page 524
- Verifying a Services, Protocols, and Trusted Sources Firewall Filter on page 525
- Verifying a TCP and ICMP Flood Firewall Filter on page 526
- Verifying a Firewall Filter That Handles Fragments on page 527

Displaying Firewall Filter Configurations

Purpose	Verify the configuration of the firewall filter. You can analyze the flow of the firewall filter terms by displaying the entire 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 services` or `show firewall` command for stateful and stateless firewall filters.

The sample output in this section displays the following firewall filters (in order):

- Stateful firewall filter and NAT configured in “Configuring a Stateful Firewall Filter with a Configuration Editor” on page 480
- Stateless `protect-RE` filter configured in “Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources” on page 502
- Stateless `protect-RE` filter configured in “Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods” on page 506
- Stateless `fragment-RE` filter configured in “Configuring a Routing Engine Firewall Filter to Handle Fragments” on page 511

Sample Output

```
[edit]
user@host# show services
stateful-firewall {
  rule to-wan-rule {
    match-direction output;
    term app-term {
      from {
        application-sets junos-algs-outbound;
      }
      then {
        accept;
      }
    }
    term accept-all-term {
      then {
        accept;
      }
    }
  }
  rule from-wan-rule {
    match-direction input;
    term wan-src-addr-term {
      from {
        source-address {
          192.168.33.0/24;
        }
      }
      then {
        accept;
      }
    }
    term discard-all-term {
      then {
        discard;
      }
    }
  }
}
```



```

nat {
  pool public-pool {
    address-range low 10.148.2.1 high 10.148.2.32;
    port automatic;
  }
  rule nat-to-wan-rule {
    match-direction output;
    term private-public-term {
      then {
        translated {
          source-pool public-pool;
          translation-type source dynamic;
        }
      }
    }
  }
}
service-set wan-service-set {
  stateful-firewall-rules to-wan-rule;
  stateful-firewall-rules from-wan-rule;
  nat-rules nat-to-wan-rule;
  interface-service {
    service-interface sp-0/0/0;
  }
}

```

```

[edit]
user@host# show firewall
firewall {
  family inet {
    filter protect-RE {
      term ssh-term {
        from {
          source-address {
            192.168.122.0/24;
          }
          protocol tcp;
          destination-port ssh;
        }
        then accept;
      }
      term bgp-term {
        from {
          source-address {
            10.2.1.0/24;
          }
          protocol tcp;
          destination-port bgp;
        }
        then accept;
      }
      term discard-rest-term {
        then {
          log;
        }
      }
    }
  }
}

```

```

        syslog;
        discard;
    }
}
}
}
}

```

```

[edit]
user@host# show firewall
firewall {
    policer tcp-connection-policer {
        filter-specific;
        if-exceeding {
            bandwidth-limit 500k;
            burst-size-limit 15k;
        }
        then discard;
    }
    policer icmp-policer {
        filter-specific;
        if-exceeding {
            bandwidth-limit 1m;
            burst-size-limit 15k;
        }
        then discard;
    }
    family inet {
        filter protect-RE {
            term tcp-connection-term {
                from {
                    source-prefix-list {
                        trusted-addresses;
                    }
                    protocol tcp;
                    tcp-flags "(syn & !ack) | fin | rst";
                }
                then {
                    policer tcp-connection-policer;
                    accept;
                }
            }
            term icmp-term {
                from {
                    protocol icmp;
                    icmp-type [ echo-request echo-reply unreachable time-exceeded ];
                }
                then {
                    policer icmp-policer;
                    count icmp-counter;
                    accept;
                }
            }
        }
        additional terms ...
    }
}

```

```

    }
  }
}

[edit]
user@host# show firewall
firewall {
    family inet {
        filter fragment-RE {
            term small-offset-term {
                from {
                    fragment-offset 1-5;
                }
                then {
                    syslog;
                    discard;
                }
            }
            term not-fragmented-term {
                from {
                    source-address {
                        10.2.1.0/24;
                    }
                    fragment-offset 0;
                    fragment-flags 0x0;
                    protocol tcp;
                    destination-port bgp;
                }
                then accept;
            }
            term first-fragment-term {
                from {
                    source-address {
                        10.2.1.0/24;
                    }
                    first-fragment;
                    protocol tcp;
                    destination-port bgp;
                }
                then accept;
            }
            term fragment-term {
                from {
                    fragment-offset 6-8191;
                }
                then accept;
            }
            additional terms ...
        }
    }
}

```

- What It Means** Verify that the output shows the intended configuration of the firewall filter. For more information about the format of a configuration file, see “Viewing the Configuration Text” on page 12.
- Verify that the terms are listed in the order in which you want the packets to be tested. You can move terms within a firewall filter by using the `insert` CLI command. For more information, see “Inserting an Identifier” on page 28.

Verifying a Stateful Firewall Filter

Purpose Verify the firewall filter configured in “Configuring a Stateful Firewall Filter with a Configuration Editor” on page 480.

Action To verify that the actions of the firewall filter terms are taken, send packets to and from the untrusted network that match the terms. In addition, verify that actions are *not* taken for packets that do not match.

- Send packets—associated with the `junos-algs-outbound` application set—from a host in the trusted network to a host in the untrusted network. Verify that packets received from the host in the untrusted network are responses only to the session originated by the host in the trusted network. To ensure that packets from the host are not accepted because of rule `from-wan-rule`, do not send packets to the host in the untrusted network with an IP address that matches `192.168.33.0/24`.

For example, send a ping request from host `trusted-nw-trusted-host` to host `untrusted-nw-untrusted-host`, and verify that a ping response is returned. Ping requests and responses use ICMP, which belongs to the `junos-algs-outbound` application set.



NOTE: To view the configuration of `junos-algs-outbound`, enter the `show groups junos-defaults applications application-set junos-algs-outbound` configuration mode command.

- Send packets from a host in the untrusted network to a host in the trusted network. Verify that the host in the trusted network receives packets only from the host in the untrusted network with an IP address that matches `192.168.33.0/24`.

For example, send a ping request from host `untrusted-nw-trusted-host` with an IP address that matches `192.168.33.0/24` to host `trusted-nw-trusted-host`, and verify that a ping response is returned.

Verify that the ping response displays an IP address from the configured NAT pool.

Sample Output `user@trusted-nw-trusted-host> ping untrusted-nw-untrusted-host`

```
PING untrusted-nw-untrusted-host.acme.net (172.69.13.5): 56 data bytes
64 bytes from 192.169.13.5: icmp_seq=0 ttl=22 time=8.238 ms
64 bytes from 192.169.13.5: icmp_seq=1 ttl=22 time=9.116 ms
```

```
64 bytes from 192.169.13.5: icmp_seq=2 ttl=22 time=10.875 ms
...
```

```
user@untrusted-nw-trusted-host> ping trusted-nw-trusted-host
```

```
PING trusted-nw-trusted-host-fe-000.acme.net (112.148.2.3): 56 data bytes
64 bytes from 10.148.2.3: icmp_seq=0 ttl=253 time=18.248 ms
64 bytes from 10.148.2.3: icmp_seq=1 ttl=253 time=10.906 ms
64 bytes from 10.148.2.3: icmp_seq=2 ttl=253 time=12.845 ms
...
```

What It Means Verify the following information:

- A ping request from host `trusted-nw-trusted-host` returns a ping response from host `untrusted-nw-untrusted-host`.
- A ping request from host `untrusted-nw-trusted-host` returns a ping response from host `trusted-nw-trusted-host`. Verify that the ping response displays an IP address from the configured NAT pool of 10.148.2.1 through 10.148.2.32.

For information about using the J-Web interface to ping a host, see the *J-series Services Router Administration Guide*.

For more information about the `ping` command, see the *J-series Services Router Administration Guide* or the *JUNOS System Basics and Services Command Reference*.

Displaying Firewall Filter Logs

Purpose Verify that packets are being logged. If you included the `log` or `syslog` action in a term, verify that packets matching the term are recorded in the firewall log or your system logging facility.

Action From operational mode in the CLI, enter the `show firewall log` command.

The log of discarded packets generated from the firewall filter configured in “Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources” on page 502 is displayed in the following sample output.

Sample Output

```
user@host> show firewall log
```

```
Log :
```

Time	Filter	Action	Interface	Protocol	Src Addr	Dest Addr
15:11:02	pfe	D	fe-0/0/0.0	TCP	172.17.28.19	192.168.70.71
15:11:01	pfe	D	fe-0/0/0.0	TCP	172.17.28.19	192.168.70.71
15:11:01	pfe	D	fe-0/0/0.0	TCP	172.17.28.19	192.168.70.71
15:11:01	pfe	D	fe-0/0/0.0	TCP	172.17.28.19	192.168.70.71
...						

- What It Means** Each record of the output contains information about the logged packet. Verify the following information:
- Under **Time**, the time of day the packet was filtered is shown.
 - The **Filter** output is always **pfe**.
 - Under **Action**, the configured action of the term matches the action taken on the packet—A (accept), D (discard), R (reject).
 - Under **Interface**, the inbound (ingress) interface on which the packet arrived is appropriate for the filter.
 - Under **Protocol**, the protocol in the IP header of the packet is appropriate for the filter.
 - Under **Src Addr**, the source address in the IP header of the packet is appropriate for the filter.
 - Under **Dest Addr**, the destination address in the IP header of the packet is appropriate for the filter.

For more information about the `show firewall log` command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Displaying Firewall Filter Statistics

Purpose Verify that packets are being policed and counted.

Action From operational mode in the CLI, enter the `show firewall filter filter-name` command.

The value of the counter, `icmp-counter`, and the number of packets discarded by the policers in the firewall filter configured in “Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods” on page 506 are displayed in the following sample output.

```

Sample Output  user@host> show firewall filter protect-RE

Filter: protect-RE
Counters:
Name              Bytes              Packets
icmp-counter      1040000            5600
Policers:
Name              Packets
tcp-connection-policer 643254873
icmp-policer      7391

```

What It Means Verify the following information:

- Next to Filter, the name of the firewall filter is correct.
- Under Counters:
 - Under Name, the names of any counters configured in the firewall filter are correct.
 - Under Bytes, the number of bytes that match the filter term containing the count *counter-name* action are shown.
 - Under Packets, the number of packets that match the filter term containing the count *counter-name* action are shown.
- Under Policers:
 - Under Name, the names of any policers configured in the firewall filter are correct.
 - Under Packets, the number of packets that match the conditions specified for the policer are shown.

For more information about the `show firewall filter` command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying a Services, Protocols, and Trusted Sources Firewall Filter

Purpose Verify the firewall filter configured in “Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources” on page 502.

Action To verify that the actions of the firewall filter terms are taken, send packets to the Services Router that match the terms. In addition, verify that the filter actions are *not* taken for packets that do not match.

- Use the `ssh host-name` command from a host at an IP address that matches 192.168.122.0/24 to verify that you can log in to the Services Router using only SSH from a host with this address prefix.
- Use the `show route summary` command to verify that the routing table on the Services Router does not contain any entries with a protocol other than Direct, Local, BGP, or Static.

Sample Output `% ssh 192.168.249.71`

```
%ssh host
user@host's password:
--- JUNOS 6.4-20040518.0 (JSERIES) #0: 2004-05-18 09:27:50 UTC

user@host>

user@host> show route summary
```

```

Router ID: 192.168.249.71

inet.0: 34 destinations, 34 routes (33 active, 0 holddown, 1 hidden)
      Direct:        10 routes,        9 active
        Local:         9 routes,        9 active
         BGP:         10 routes,       10 active
        Static:         5 routes,        5 active
...

```

What It Means Verify the following information:

- You can successfully log in to the Services Router using SSH.
- The show route summary command does not display a protocol other than Direct, Local, BGP, or Static.

For more information about the `show route summary` command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Verifying a TCP and ICMP Flood Firewall Filter

Purpose Verify the firewall filter configured in “Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods” on page 506.

Action To verify that the actions of the firewall filter terms are taken, send packets to the Services Router that match the terms. In addition, verify that the filter actions are *not* taken for packets that do not match.

- Verify that the Services Router can establish only TCP sessions with a host at an IP address that matches 192.168.122.0/24 or 10.2.1.0/24. For example, log in to the router with the `telnet host-name` command from another host with one of these address prefixes.
- Use the `ping host-name` command to verify that the Services Router responds only to ICMP packets (such as ping requests) that do not exceed the policer traffic rates.
- Use the `ping host-name size bytes` command to exceed the policer traffic rates by sending ping requests with large data payloads.

Sample Output user@host> **telnet 192.168.249.71**

```

Trying 192.168.249.71...
Connected to host.acme.net.
Escape character is '^]'.

host (ttyp0)

login: user
Password:

--- JUNOS 6.4-20040521.1 built 2004-05-21 09:38:12 UTC

user@host>

```



```

user@host> ping 192.168.249.71

PING host-fe-000.acme.net (192.168.249.71): 56 data bytes
64 bytes from 192.168.249.71: icmp_seq=0 ttl=253 time=11.946 ms
64 bytes from 192.168.249.71: icmp_seq=1 ttl=253 time=19.474 ms
64 bytes from 192.168.249.71: icmp_seq=2 ttl=253 time=14.639 ms
...

user@host> ping 192.168.249.71 size 20000

PING host-fe-000.acme.net (192.168.249.71): 20000 data bytes
^C
--- host-fe-000.acme.net ping statistics ---
12 packets transmitted, 0 packets received, 100% packet loss

```

What It Means Verify the following information:

- You can successfully log in to the Services Router using Telnet.
- The Services Router sends responses to the `ping host` command.
- The Services Router does not send responses to the `ping host size 20000` command.

For more information about the `ping` command, see the *J-series Services Router Administration Guide* or the *JUNOS System Basics and Services Command Reference*.

For information about using the J-Web interface to ping a host, see the *J-series Services Router Administration Guide*.

For more information about the `telnet` command, see the *J-series Services Router Administration Guide* or the *JUNOS System Basics and Services Command Reference*.

Verifying a Firewall Filter That Handles Fragments

Purpose Verify the firewall filter configured in “Configuring a Routing Engine Firewall Filter to Handle Fragments” on page 511.

Action To verify that the actions of the firewall filter terms are taken, send packets to the Services Router that match the terms. In addition, verify that the filter actions are *not* taken for packets that do not match.

- Verify that packets with small fragment offsets are recorded in the router’s system logging facility.
- Use the `show route summary` command to verify that the routing table does not contain any entries with a protocol other than Direct, Local, BGP, or Static.

Sample Output user@host> `show route summary`

```

Router ID: 192.168.249.71

inet.0: 34 destinations, 34 routes (33 active, 0 holddown, 1 hidden)
      Direct:      10 routes,          9 active

```

```
Local:      9 routes,      9 active
BGP:       10 routes,     10 active
Static:     5 routes,      5 active
...
```

What It Means Verify that the `show route summary` command does not display a protocol other than `Direct`, `Local`, `BGP`, or `Static`. For more information about the `show route summary` command, see the *JUNOS Routing Protocols and Policies Command Reference*.

Chapter 23

Configuring Class of Service with DiffServ

You configure class of service (CoS) with Differentiated Services (DiffServ) when you need to override the default packet forwarding behavior of a Services Router—especially in the three areas identified in Table 160.

Table 160: Reasons to Configure Class of Service (Cos) with DiffServ

Default Behavior to Override with CoS	CoS Configuration Area
Packet classification—By default, the Services Router does not use DiffServ to classify packets. Packet classification applies to incoming traffic.	Classifiers
Scheduling queues—By default, the Services Router has only two queues enabled. Scheduling queues apply to outgoing traffic.	Schedulers
Packet headers—By default, the Services Router does not rewrite CoS bits in packet headers. Rewriting packet headers applies to outgoing traffic.	Rewrite rules

You can use either the J-Web configuration editor or CLI configuration editor to configure CoS with DiffServ. The J-Web interface does not include Quick Configuration pages for CoS or DiffServ.

This chapter contains the following topics. For more information about CoS and DiffServ, see the *JUNOS Class of Service Configuration Guide*.

- Before You Begin on page 530
- Configuring CoS with DiffServ with a Configuration Editor on page 530
- Verifying a DiffServ Configuration on page 560

Before You Begin

Before you begin configuring a Services Router for CoS with DiffServ, complete the following tasks:

- If you do not already have a basic understanding of CoS and DiffServ, read “Policy, Firewall Filter, and Class-of-Service Overview” on page 433.
- Determine whether the Services Router needs to support different traffic streams, such as voice or video. If so, CoS with DiffServ helps to make sure this traffic receives more than basic best-effort packet delivery service.
- Determine whether the Services Router is directly attached to any applications that send DiffServ packets. If no sources are enabled for DiffServ, you must configure and apply rewrite rules on the interfaces to the sources.
- Determine whether the Services Router must support DiffServ assured forwarding (AF) classes. Assured forwarding usually requires random early detection (RED) drop profiles to be configured and applied.
- Determine whether the Services Router must support DiffServ expedited forwarding (EF) classes with a policer. Policers require you to apply a burst size and bandwidth limit to the traffic flow, and set a consequence for packets that exceed these limits—usually a high loss priority, so that packets exceeding the policer limits are discarded first.

Configuring CoS with DiffServ with a Configuration Editor

To configure the Services Router as a node in a network supporting CoS with DiffServ, you must perform the following tasks marked (*Required*). For information about using the J-Web and CLI configuration editors, see “Using J-Web Configuration Tools” on page 3.

- Configuring a Policer for a Firewall Filter (*Required*) on page 531
- Configuring and Applying a Firewall Filter for a Multifield Classifier (*Required*) on page 532
- Assigning Forwarding Classes to Output Queues (*Required*) on page 535
- Configuring and Applying Rewrite Rules (*Required*) on page 537
- Configuring and Applying Behavior Aggregate Classifiers (*Required*) on page 542
- Configuring RED Drop Profiles for Assured Forwarding Congestion Control (*Required*) on page 546
- Configuring Schedulers (*Optional*) on page 548
- Configuring and Applying Scheduler Maps (*Optional*) on page 552

- Configuring and Applying Virtual Channels (Optional) on page 555
- Configuring and Applying Adaptive Shaping (Optional) on page 559

Configuring a Policer for a Firewall Filter (Required)

You configure a policer to detect packets that exceed the limits established for DiffServ expedited forwarding. For DiffServ, packets that exceed these limits are given a higher loss priority than packets within the bandwidth and burst size limits.

The following example shows how to configure a policer called `ef-policer` that identifies for likely discard expedited forwarding packets with a burst size greater than 2000 bytes and a bandwidth greater than 10 percent.

For more information about firewall filters, see “Configuring Firewall Filters and NAT” on page 475 and the *JUNOS Policy Framework Configuration Guide*.

To configure an expedited forwarding policer for a firewall filter 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 161.
3. Go on to “Configuring and Applying a Firewall Filter for a Multifield Classifier (Required)” on page 532.

Table 161: Configuring a Policer for a Firewall Filter

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Firewall level in the configuration hierarchy.	In the configuration editor hierarchy, select Firewall .	From the top of the configuration hierarchy, enter edit firewall
Create the policer for expedited forwarding, and give the policer a name—for example, <code>ef-policer</code> .	<div>1. Click Add new entry next to Policer.</div> <div>2. In the Policer name box, type <code>ef-policer</code>.</div>	Enter edit policer ef-policer

Table 161: Configuring a Policer for a Firewall Filter (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Set the burst limit for the policer—for example, 2k.	1. Click Configure next to If exceeding.	Enter set if-exceeding burst-limit-size 2k
Set the bandwidth limit or percentage for the bandwidth allowed for this type of traffic—for example, use a bandwidth percent of 10.	2. In the Burst size limit box, type a limit for the burst size allowed—for example, 2k. 3. From the Bandwidth list, select bandwidth-percent . 4. In the Bandwidth percent box, type 10. 5. Click OK .	set if-exceeding bandwidth-percent 10
Enter the loss priority for packets exceeding the limits established by the policer—for example, high.	1. Click Configure next to Then. 2. From the Loss priority list, select high . 3. Click OK three times.	Enter set then loss-priority high

Configuring and Applying a Firewall Filter for a Multifield Classifier (Required)

You configure a multifield (MF) classifier to detect packets of interest to CoS and assign the packet to the proper forwarding class independently of the DiffServ code point (DSCP). To configure a multifield classifier on a customer-facing or host-facing link, configure a firewall filter to classify traffic. Packets are classified as they arrive on an interface.

One common way to detect packets of CoS interest is by source or destination address. The destination address is used in this example, but many other matching criteria for packet detection are available to firewall filters.

This example shows how to configure the firewall filter `mf-classifier` and apply it to the Services Router's Fast Ethernet interface `fe-0/0/0`. The firewall filter consists of the rules (terms) listed in Table 162.

Table 162: Sample mf-classifier Firewall Filter Terms

Rule (Term)	Purpose	Contents
assured forwarding	Detects packets destined for 192.168.44.55, assigns them to an assured forwarding class, and gives them a low likelihood of being dropped.	Match condition: destination address 192.168.44.55 Forwarding class: <code>af-class</code> Loss priority: low

Table 162: Sample mf-classifier Firewall Filter Terms (continued)

Rule (Term)	Purpose	Contents
expedited-forwarding	Detects packets destined for 192.168.66.77, assigns them to an expedited forwarding class, and subjects them to the EF policer configured in “Configuring a Policer for a Firewall Filter (Required)” on page 531.	Match condition: destination address 192.168.66.77 Forwarding class: ef-class Policer: ef-policer
network control	Detects packets with a network control precedence and forwards them to the network control class.	Match condition: precedence net-control Forwarding class: nc-class
best-effort-data	Detects all other packets and assigns them to the best effort class.	Forwarding class: be-class

For more information about firewalls filters see “Configuring Firewall Filters and NAT” on page 475 and the *JUNOS Policy Framework Configuration Guide*.

To configure a firewall filter for a multifold classifier 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 163.
3. Go on to “Assigning Forwarding Classes to Output Queues (Required)” on page 535.

Table 163: Configuring and Applying a Firewall Filter for a Multifold Classifier

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Firewall level in the configuration hierarchy.	In the configuration editor hierarchy, select Firewall .	From the top of the configuration hierarchy, enter edit firewall
Create the multifold classifier filter and name it—for example, mf-classifier.	<ol style="list-style-type: none"> 1. Click Add new entry next to Filter. 2. In the Filter name box, type mf-classifier. 3. Select the check box next to Interface specific. 	Enter edit filter mf-classifier set interface-specific
Create the term for the assured forwarding traffic class, and give it a name—for example, assured-forwarding.	<ol style="list-style-type: none"> 1. Click Add new entry next to Term. 2. In the Rule name box, type assured-forwarding. 	Enter edit term assured-forwarding

Table 163: Configuring and Applying a Firewall Filter for a Multifield Classifier (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the match condition for the assured forwarding traffic class. Use the destination address for assured forwarding traffic—for example, 192.168.44.55.	<ol style="list-style-type: none"> 1. Click Configure next to From. 2. Click Add new entry next to Destination address. 3. In the Address box, type 192.168.44.55. 4. Click OK three times. 	<p>Enter</p> <p>set from destination-address 192.168.44.55</p>
Create the forwarding class for assured forwarding DiffServ traffic—for example, af-class. Set the loss priority for the assured forwarding traffic class—for example, low.	<ol style="list-style-type: none"> 1. Click Configure next to Then. 2. In the Forwarding class box, type af-class. 3. From the Loss priority list, select low. 4. Click OK twice. 	<p>From the top of the configuration hierarchy, enter</p> <p>edit firewall filter mf-classifier term assured-forwarding</p> <p>set then forwarding-class af-class</p> <p>set then loss-priority low</p>
Create the term for the expedited forwarding traffic class, and give it a name—for example, expedited-forwarding.	<ol style="list-style-type: none"> 1. Click Add new entry next to Term. 2. In the Rule name box, type expedited-forwarding. 	<p>Enter</p> <p>edit term expedited-forwarding</p>
Create the match condition for the assured forwarding traffic class. Use the destination address for assured forwarding traffic—for example, 192.168.66.77.	<ol style="list-style-type: none"> 1. Click Configure next to From. 2. Click Add new entry next to Destination address. 3. In the Address box, type 192.168.66.77. 4. Click OK twice. 	<p>Enter</p> <p>set from destination-address 192.168.66.77</p>
Create the forwarding class for expedited forwarding DiffServ traffic—for example, ef-class. Apply the policer for the expedited forwarding traffic class. Use the EF policer previously configured for expedited forwarding DiffServ traffic—ef-policer.	<ol style="list-style-type: none"> 1. Click Configure next to Then. 2. In the Forwarding class box, type ef-class. 3. In the Policer box, type ef-policer. 4. Click OK twice. 	<p>From the top of the configuration hierarchy, enter</p> <p>edit firewall filter mf-classifier term expedited-forwarding</p> <p>set then forwarding-class ef-class</p> <p>set then policer ef-policer</p>
(See “Configuring a Policer for a Firewall Filter (Required)” on page 531.)		
Create the term for the network control traffic class, and give it a name—for example, network-control.	<ol style="list-style-type: none"> 1. Click Add new entry next to Term. 2. In the Rule name box, type network-control. 	<p>Enter</p> <p>edit term network-control</p>

Table 163: Configuring and Applying a Firewall Filter for a Multifield Classifier (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the match condition for the network control traffic class.	<ol style="list-style-type: none"> 1. Click Configure next to From. 2. From the Precedence choice list, select Precedence. 3. Click Add new entry next to Precedence. 4. From the Value keyword list, select net-control. 5. Click OK twice. 	<p>Enter</p> <p>set from traffic-class net-control</p>
Create the forwarding class for the network control traffic class, and give it a name—for example, nc-class .	<ol style="list-style-type: none"> 1. Click Configure next to Then. 2. In the Forwarding class box, type nc-class. 3. Click OK twice. 	<p>From the top of the configuration hierarchy, enter</p> <p>edit firewall filter mf-classifier term network-control</p> <p>set then forwarding-class nc-class</p>
Create the term for the best-effort traffic class, and give it a name—for example, best-effort-data .	<ol style="list-style-type: none"> 1. Click Add new entry next to Term. 2. In the Rule name box, type best-effort-data. 	<p>Enter</p> <p>edit term best-effort-data</p>
Create the forwarding class for the best-effort traffic class, and give it a name—for example, be-class . (Because this is the last term in the filter, it has no match condition.)	<ol style="list-style-type: none"> 1. Click Configure next to Then. 2. In the Forwarding class box, type be-class. 3. Click OK four times. 	<p>From the top of the configuration hierarchy, enter</p> <p>set then forwarding-class be-class</p>
Navigate to the Interfaces level in the configuration hierarchy.	In the configuration editor hierarchy, select Interfaces .	<p>From the top of the configuration hierarchy, enter</p> <p>edit interfaces</p>
Apply the multifield classifier firewall filter mf-classifier as an input filter on each customer-facing or host-facing interface that needs the filter—for example, on fe-0/0/0 , unit 0.	<ol style="list-style-type: none"> 1. Click the Interface fe-0/0/0 and Unit 0. 2. Click Configure next to Inet. 3. Click Configure next to Filter. 4. In the Input box, type mf-classifier. 5. Click OK five times. 	<p>Enter</p> <p>set fe-0/0/0 unit 0 family inet filter input mf-classifier</p>

Assigning Forwarding Classes to Output Queues (Required)

You must assign the forwarding classes established by the mf-classifier multifield classifier to output queues. This example assigns output queues as shown in Table 164.

Table 164: Sample Output Queue Assignments for mf-classifier Forwarding Queues

mf-classifier Forwarding Class	For Traffic Type	Output Queue
be-class	Best-effort traffic	Queue 0
ef-class	Expedited forwarding traffic	Queue 1
af-class	Assured forwarding traffic	Queue 2
nc-class	Network control traffic	Queue 3

For multifield classifier details, see “Configuring and Applying a Firewall Filter for a Multifield Classifier (Required)” on page 532.

To assign forwarding classes to output queues 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 165.
3. Go on to “Configuring and Applying Rewrite Rules (Required)” on page 537.

Table 165: Assigning Forwarding Classes to Output Queues

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Class of service level in the configuration hierarchy.	In the configuration editor hierarchy, select Class of service .	From the top of the configuration hierarchy, enter edit class-of-service
Assign best-effort traffic to queue 0.	<ol style="list-style-type: none"> 1. Click Configure next to Forwarding classes. 2. Click Add new entry next to Queue. 3. In the Queue num box, type 0. 4. In the Class name box, type the previously configured name of the best-effort class—be-class. 5. Click OK. 	Enter set forwarding-classes queue 0 be-class
Assign expedited forwarding traffic to queue 1.	<ol style="list-style-type: none"> 1. Click Add new entry next to Queue. 2. In the Queue num box, type 1. 3. In the Class name box, type the previously configured name of the expedited forwarding class—ef-class. 4. Click OK. 	Enter set forwarding-classes queue 1 ef-class

Table 165: Assigning Forwarding Classes to Output Queues (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Assign assured forwarding traffic to queue 2.	<ol style="list-style-type: none"> 1. Click Add new entry next to Queue. 2. In the Queue num box, type 2. 3. In the Class name box, type the previously configured name of the assured forwarding class—af-class. 4. Click OK. 	Enter set forwarding-classes queue 2 af-class
Assign network control traffic to queue 3.	<ol style="list-style-type: none"> 1. Click Add new entry next to Queue. 2. In the Queue num box, type 3. 3. In the Class name box, type the previously configured name of the expedited forwarding class—nc-class. 4. Click OK twice. 	Enter set forwarding-classes queue 3 nc-class

Configuring and Applying Rewrite Rules (Required)

You optionally configure rewrite rules to replace DiffServ code points (DSCPs) on packets received from the customer or host with the values expected by other routers. You do not have to configure rewrite rules if the received packets already contain valid DSCPs. Rewrite rules apply the forwarding class information and packet loss priority used internally by the Services Router to establish the DSCP on outbound packets. Once configured, you must apply the rewrite rules to the correct interfaces.

The following example shows how to create the rewrite rules **rewrite-dscps**, and apply them to the Services Router's Fast Ethernet interface **fe-0/0/0**. The rewrite rules replace the DSCPs on packets in the four forwarding classes, as shown in Table 166.

Table 166: Sample rewrite-dscps Rewrite Rules to Replace DSCPs

mf-classifier Forwarding Class	For CoS Traffic Type	rewrite-dscps Rewrite Rules
be-class	Best-effort traffic	Low-priority code point: 000000
		High-priority code point: 000001
ef-class	Expedited forwarding traffic	Low-priority code point: 101110
		High-priority code point: 101111

Table 166: Sample rewrite-dscps Rewrite Rules to Replace DSCPs (continued)

mf-classifier Forwarding Class	For CoS Traffic Type	rewrite-dscps Rewrite Rules
af-class	Assured forwarding traffic	Low-priority code point: 001010
		High-priority code point: 001100
nc-class	Network control traffic	Low-priority code point: 110000
		High-priority code point: 110001

To configure and apply rewrite rules 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 167.
3. Go on to “Configuring and Applying Behavior Aggregate Classifiers (Required)” on page 542.

Table 167: Configuring and Applying Rewrite Rules

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Class of service level in the configuration hierarchy.	In the configuration editor hierarchy, select Class of service .	From the top of the configuration hierarchy, enter edit class-of-service
Configure rewrite rules for DiffServ CoS.	<ol style="list-style-type: none"> 1. Click Configure next to Rewrite rules. 2. Click Add new entry next to Dscp. 3. In the Name box, type the name of the rewrite rules—for example, rewrite-dscps. 	Enter edit rewrite-rules dscp rewrite-dscps

Table 167: Configuring and Applying Rewrite Rules (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure best-effort forwarding class rewrite rules.	<ol style="list-style-type: none"> 1. Click Add new entry next to Forwarding class. 2. In the Class name box, type the name of the previously configured best-effort forwarding class—be-class. 3. Click Add new entry next to Loss priority. 4. From the Loss val list, select low. 5. In the Code point box, type the value of the low-priority code point for best-effort traffic—for example, 000000. 6. Click OK. 7. Click Add new entry next to Loss priority. 8. From the Loss val list, select high. 9. In the Code point box, type the value of the high-priority code point for best-effort traffic—for example, 000001. 10. Click OK twice. 	<p>Enter</p> <pre>set forwarding-class be-class loss-priority low code points 000000 set forwarding-class be-class loss-priority high code points 000001</pre>

Table 167: Configuring and Applying Rewrite Rules (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure expedited forwarding class rewrite rules.	<ol style="list-style-type: none"> 1. Click Add new entry next to Forwarding class. 2. In the Class name box, type the name of the previously configured expedited forwarding class—ef-class. 3. Click Add new entry next to Loss priority. 4. From the Loss val list, select low. 5. In the Code point box, type the value of the low-priority code point for expedited forwarding traffic—for example, 101110. 6. Click OK. 7. Click Add new entry next to Loss priority. 8. From the Loss val list, select high. 9. In the Code point box, type the value of the high-priority code point for expedited forwarding traffic—for example, 101111. 10. Click OK twice. 	<p>Enter</p> <pre>set forwarding-class ef-class loss-priority low code points 101110 set forwarding-class ef-class loss-priority high code points 101111</pre>

Table 167: Configuring and Applying Rewrite Rules (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure assured forwarding class rewrite rules.	<ol style="list-style-type: none"> 1. Click Add new entry next to Forwarding class. 2. In the Class name box, type the name of the previously configured assured forwarding class—af-class. 3. Click Add new entry next to Loss priority. 4. From the Loss val list, select low. 5. In the Code point box, type the value of the low-priority code point for assured forwarding traffic—for example, 001010. 6. Click OK. 7. Click Add new entry next to Loss priority. 8. From the Loss val list, select high. 9. In the Code point box, type the value of the high-priority code point for assured forwarding traffic—for example, 001100. 10. Click OK twice. 	<p>Enter</p> <pre>set forwarding-class af-class loss-priority low code points 001010 set forwarding-class af-class loss-priority high code points 001100</pre>

Table 167: Configuring and Applying Rewrite Rules (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure network control class rewrite rules.	<ol style="list-style-type: none"> 1. Click Add new entry next to Forwarding class. 2. In the Class name box, type the name of the previously configured network control forwarding class—nc-class. 3. Click Add new entry next to Loss priority. 4. From the Loss val list, select low. 5. In the Code point box, type the value of the low-priority code point for network control traffic—for example, 110000. 6. Click OK. 7. Click Add new entry next to Loss priority. 8. From the Loss val list, select high. 9. In the Code point box, type the value of the high-priority code point for network control traffic—for example, 110001. 10. Click OK twice. 	<p>Enter</p> <pre>set forwarding-class nc-class loss-priority low code points 110000 set forwarding-class nc-class loss-priority high code points 110001</pre>
Apply rewrite rules to an interface.	<ol style="list-style-type: none"> 1. Click Add new entry next to Interfaces 2. In the Interface name box, type the name of the interface—for example, fe-0/0/0. 3. In the Rewrite rules box, type the name of the previously configured rewrite rules—rewrite-dscps. 4. Click OK. 	<p>Enter</p> <pre>set interfaces fe-0/0/0 unit 0 rewrite-rules rewrite-dscps</pre>

Configuring and Applying Behavior Aggregate Classifiers (Required)

You configure DiffServ behavior aggregate (BA) classifiers to classify packets that contain valid DSCPs to appropriate queues. Once configured, you must apply the BA classifier to the correct interfaces.

The following example shows how to configure the DSCP BA classifier **ba-classifier** as the default DSCP map, and apply it to the Services Router's Fast Ethernet

interface fe-0/0/0. The BA classifier assigns loss priorities, as shown in Table 168, to incoming packets in the four forwarding classes.

Table 168: Sample ba-classifier Loss Priority Assignments

mf-classifier Forwarding Class	For CoS Traffic Type	ba-classifier Assignments
be-class	Best-effort traffic	High-priority code point: 000001
ef-class	Expedited forwarding traffic	High-priority code point: 101111
af-class	Assured forwarding traffic	High-priority code point: 001100
nc-class	Network control traffic	High-priority code point: 110001

To configure and apply BA classifiers 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 169.
3. Go on to “Configuring RED Drop Profiles for Assured Forwarding Congestion Control (Required)” on page 546.

Table 169: Configuring and Applying Behavior Aggregate Classifiers

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Class of service level in the configuration hierarchy.	In the configuration editor hierarchy, select Class of service .	From the top of the configuration hierarchy, enter edit class-of-service
Configure BA classifiers for DiffServ CoS.	<ol style="list-style-type: none"> 1. Click Configure next to Classifiers. 2. Click Add new entry next to Dscp. 3. In the Name box, type the name of the BA classifier—for example, ba-classifier. 4. In the Import box, type the name of the default DSCP map, default. 	Enter edit classifiers dscp ba-classifier set import default

Table 169: Configuring and Applying Behavior Aggregate Classifiers (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a best-effort forwarding class classifier.	<ol style="list-style-type: none"> 1. Click Add new entry next to Forwarding class. 2. In the Class name box, type the name of the previously configured best-effort forwarding class—be-class. 3. Click Add new entry next to Loss priority. 4. From the Loss val list, select high. 5. In the Code point box, type the value of the high-priority code point for best-effort traffic—for example, 00001. 6. Click OK three times. 	<p>Enter</p> <pre>set forwarding-class be-class loss-priority high code points 000001</pre>
Configure an expedited forwarding class classifier.	<ol style="list-style-type: none"> 1. Click Add new entry next to Forwarding class. 2. In the Class name box, type the name of the previously configured expedited forwarding class—ef-class. 3. Click Add new entry next to Loss priority. 4. From the Loss val list, select high. 5. In the Code point box, type the value of the high-priority code point for expedited forwarding traffic—for example, 101111. 6. Click OK three times. 	<p>Enter</p> <pre>set forwarding-class ef-class loss-priority high code points 101111</pre>

Table 169: Configuring and Applying Behavior Aggregate Classifiers (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure an assured forwarding class classifier.	<ol style="list-style-type: none"> 1. Click Add new entry next to Forwarding class. 2. In the Class name box, type the name of the previously configured assured forwarding class—af-class. 3. Click Add new entry next to Loss priority. 4. From the Loss val list, select high. 5. In the Code point box, type the value of the high-priority code point for assured forwarding traffic—for example, 001100. 6. Click OK three times. 	<p>Enter</p> <pre>set forwarding-class af-class loss-priority high code points 001100</pre>
Configure a network control class classifier.	<ol style="list-style-type: none"> 1. Click Add new entry next to Forwarding class. 2. In the Class name box, type the name of the previously configured network control forwarding class—nc-class. 3. Click Add new entry next to Loss priority. 4. From the Loss val list, select high. 5. In the Code point box, type the value of the high-priority code point for network control traffic—for example, 110001. 6. Click OK three times. 	<p>Enter</p> <pre>set forwarding-class nc-class loss-priority high code points 110001</pre>
Apply the BA classifier to an interface.	<ol style="list-style-type: none"> 1. Click Add new entry next to Interfaces. 2. In the Interface name box, type the name of the interface—for example, fe-0/0/0. 3. In the Classifiers box, type the name of the previously configured BA classifier—ba-classifier. 4. Click OK. 	<p>Enter</p> <pre>set interfaces fe-0/0/0 unit 0 classifiers dscp ba-classifier</pre>

Configuring RED Drop Profiles for Assured Forwarding Congestion Control (Required)

If the Services Router must support DiffServ assured forwarding (AF), you can control congestion by configuring random early detection (RED) drop profiles. RED drop profiles use drop probabilities for different levels of buffer fullness to determine which scheduling queue on the router is likely to drop DiffServ assured forwarding (AF) packets under congested conditions. The router can drop packets when the queue buffer becomes filled to the configured percentage.

Assured forwarding traffic with the PLP (packet loss priority) bit set is more likely to be discarded than traffic without the PLP bit set. This example shows how to configure a drop probability and a queue fill level for both PLP and non-PLP assured forwarding traffic. It is only one example of how to use RED drop profiles.

The example shows how to configure the RED drop profiles listed in Table 170.

Table 170: Sample RED Drop Profiles

Drop Profile	Drop Probability	Queue Fill Level
af-normal—For non-PLP (normal) assured forwarding traffic	Between 0 (never dropped) and 100 percent (always dropped)	Between 95 and 100 percent
af-with-plp—For PLP (aggressive packet dropping) assured forwarding traffic	Between 95 and 100 percent (always dropped)	Between 80 and 95 percent

To configure RED drop profiles for assured forwarding congestion control on 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 171.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following tasks:
 - To assign resources, priorities, and profiles to output queues, see “Configuring Schedulers (Optional)” on page 548.
 - To apply rules to logical interfaces, see “Configuring and Applying Virtual Channels (Optional)” on page 555.
 - To use adaptive shapers to limit bandwidth for Frame Relay, see “Configuring and Applying Adaptive Shaping (Optional)” on page 559.
 - To check the configuration, see “Verifying a DiffServ Configuration” on page 560.

Table 171: Configuring RED Drop Profiles for Assured Forwarding Congestion Control

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Class of service level in the configuration hierarchy.	In the configuration editor hierarchy, select Class of service .	From the top of the configuration hierarchy, enter edit class-of-service
Configure the lower drop probability for normal, non-PLP traffic.	<ol style="list-style-type: none"> 1. Click Add new entry next to Drop profiles. 2. In the Profile name box, type the name of the drop profile—for example, af-normal. 3. Click Configure next to Interpolate. 4. Click Add new entry next to Drop probability. 5. In the Value box, type a number for the first drop point—for example, 0. 6. Click OK. 7. Click Add new entry next to Drop probability again. 8. In the Value box, type a number for the next drop point—for example, 100. 9. Click OK. 	Enter edit drop-profiles af-normal interpolate set drop-probability 0 set drop-probability 100
Configure a queue fill level for the lower non-PLP drop probability.	<ol style="list-style-type: none"> 1. Click Add new entry next to Fill level. 2. In the Value box, type a number for the first fill level—for example, 95. 3. Click OK. 4. In the Value box, type a number for the next fill level—for example, 100. 5. Click OK three times. 	Enter set fill-level 95 set fill-level 100

Table 171: Configuring RED Drop Profiles for Assured Forwarding Congestion Control (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the higher drop probability for PLP traffic.	<ol style="list-style-type: none"> 1. Click Add new entry next to Drop profiles. 2. In the Profile name box, type the name of the drop profile—for example, af-with-plp. 3. Click Configure next to Interpolate. 4. Click Add new entry next to Drop probability. 5. In the Value box, type a number for the first drop point—for example, 95. 6. Click OK. 7. In the Value box, type a number for the next drop point—for example, 100. 8. Click OK. 	<p>Enter</p> <p>edit drop-profiles af-with-PLP interpolate</p> <p>set drop-probability 95</p> <p>set drop-probability 100</p>
Configure a queue fill level for the higher PLP drop probability.	<ol style="list-style-type: none"> 1. Click Add new entry next to Fill level. 2. In the Value box, type a number for the first fill level—for example, 80. 3. Click OK. 4. In the Value box, type a number for the next fill level—for example, 95. 5. Click OK. 	<p>Enter</p> <p>set fill-level 80</p> <p>set fill-level 95</p>

Configuring Schedulers (Optional)

You configure schedulers to assign resources, priorities, and drop profiles to output queues. By default, only queues 0 and 4 have resources assigned.

This example creates the schedulers listed in Table 172.

Table 172: Sample Schedulers

Scheduler	For CoS Traffic Type	Assigned Priority	Allocated Portion of Queue Buffer	Assigned Bandwidth (Transmit Rate)
be-scheduler	Best-effort traffic	Low	40 percent	10 percent
ef-scheduler	Expedited forwarding traffic	High	10 percent	10 percent
af-scheduler	Assured forwarding traffic	High	45 percent	45 percent
nc-scheduler	Network control traffic	Low	5 percent	5 percent

To configure schedulers 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 173.
3. Go on to “Configuring and Applying Scheduler Maps (Optional)” on page 552.

Table 173: Configuring Schedulers

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Class of service level in the configuration hierarchy.	In the configuration editor hierarchy, select Class of service .	From the top of the configuration hierarchy, enter edit class-of-service
Configure a best-effort scheduler.	<ol style="list-style-type: none"> 1. Click Add new entry next to Schedulers. 2. In the Scheduler name box, type the name of the best-effort scheduler—for example, be-scheduler. 	Enter edit schedulers be-scheduler
Configure a best-effort scheduler priority and buffer size.	<ol style="list-style-type: none"> 1. In the Priority box, type low. 2. Click Configure next to Buffer size. 3. From the Buffer size choice list, select the basis for the buffer allocation method—for example, percent. 4. In the Percent box, type the percentage of the buffer to be used by the best-effort scheduler—for example, 40. 5. Click OK. 	Enter set priority low set buffer-size percent 40

Table 173: Configuring Schedulers (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a best-effort scheduler transmit rate.	<ol style="list-style-type: none"> 1. Click Configure next to Transmit rate. 2. From the Transmit rate choice list, select the basis for the transmit rate method—for example, percent. 3. In the Percent box, type the percentage of the bandwidth to be used by the best-effort scheduler—for example, 10. 4. Click OK twice. 	Enter set transmit-rate percent 10
Configure an expedited forwarding scheduler.	<ol style="list-style-type: none"> 1. Click Add new entry next to Schedulers. 2. In the Scheduler name box, type the name of the expedited forwarding scheduler—for example, ef-scheduler. 	Enter edit schedulers ef-scheduler
Configure an expedited forwarding scheduler priority and buffer size.	<ol style="list-style-type: none"> 1. In the Priority box, type high. 2. Click Configure next to Buffer size. 3. From the Buffer size choice list, select the basis for the buffer allocation method—for example, percent. 4. In the Percent box, type the percentage of the buffer to be used by the expedited forwarding scheduler—for example, 10. 5. Click OK. 	Enter set priority high set buffer-size percent 10
Configure an expedited forwarding scheduler transmit rate.	<ol style="list-style-type: none"> 1. Click Configure next to Transmit rate. 2. From the Transmit rate choice list, select the basis for the transmit rate method—for example, percent. 3. In the Percent box, type the percentage of the bandwidth to be used by the expedited forwarding scheduler—for example, 10. 4. Click OK twice. 	Enter set transmit-rate percent 10
Configure an assured forwarding scheduler.	<ol style="list-style-type: none"> 1. Click Add new entry next to Schedulers. 2. In the Scheduler name box, type the name of the assured forwarding scheduler—for example, af-scheduler. 	Enter edit schedulers af-scheduler

Table 173: Configuring Schedulers (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure an assured forwarding scheduler priority and buffer size.	<ol style="list-style-type: none"> 1. In the Priority box, type high. 2. Click Configure next to Buffer size. 3. From the Buffer size choice list, select the basis for the buffer allocation method—for example, percent. 4. In the Percent box, type the percentage of the buffer to be used by the assured forwarding scheduler—for example, 45. 5. Click OK. 	<p>Enter</p> <p>set priority high</p> <p>set buffer-size percent 45</p>
Configure an assured forwarding scheduler transmit rate.	<ol style="list-style-type: none"> 1. Click Configure next to Transmit rate. 2. From the Transmit rate choice list, select the basis for the transmit rate method—for example, percent. 3. In the Percent box, type the percentage of the bandwidth to be used by the assured forwarding scheduler—for example, 45. 4. Click OK. 	<p>Enter</p> <p>set transmit-rate percent 45</p>
(Optional) Configure a drop profile map for assured forwarding low and high priority. (DiffServ can have a RED drop profile associated with assured forwarding.)	<ol style="list-style-type: none"> 1. Click Add new entry next to Drop profile map. 2. From the Loss priority box, select Low. 3. From the Protocol box, select Any. 4. In the Drop profile box, type the name of the drop profile—for example, af-normal. 5. Click OK. 6. Click Add new entry next to Drop profile map. 7. From the Loss priority box, select High. 8. From the Protocol box, select Any. 9. In the Drop profile box, type the name of the drop profile—for example, af-with-PLP. 10. Click OK. 	<p>Enter</p> <p>set drop-profile-map loss-priority low protocol any drop-profile af-normal</p> <p>set drop-profile-map loss-priority high protocol any drop-profile af-with-PLP</p>

Table 173: Configuring Schedulers (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a network control scheduler.	<ol style="list-style-type: none"> 1. Click Add new entry next to Schedulers. 2. In the Scheduler name box, type the name of the network control scheduler—for example, nc-scheduler. 	<p>Enter</p> <p>edit schedulers nc-scheduler</p>
Configure a network control scheduler priority and buffer size.	<ol style="list-style-type: none"> 1. In the Priority box, type low. 2. Click Configure next to Buffer size. 3. From the Buffer size choice list, select the basis for the buffer allocation method—for example, percent. 4. In the Percent box, type the percentage of the buffer to be used by the network control scheduler—for example, 5. 5. Click OK. 	<p>Enter</p> <p>set priority low</p> <p>set buffer-size percent 5</p>
Configure a network control scheduler transmit rate.	<ol style="list-style-type: none"> 1. Click Configure next to Transmit rate. 2. From the Transmit rate choice list, select the basis for the transmit rate method—for example, percent. 3. In the Percent box, type the percentage of the bandwidth to be used by the network control scheduler—for example, 5. 4. Click OK twice. 	<p>Enter</p> <p>set transmit-rate percent 5</p>

Configuring and Applying Scheduler Maps (Optional)

You configure a scheduler map to assign a forwarding class to a scheduler, then apply the scheduler map to any interface that must enforce DiffServ CoS.

The following example shows how to create the scheduler map **diffserv-cos-map** and apply it to the Services Router's Fast Ethernet interface **fe-0/0/0**. The map associates the **mf-classifier** forwarding classes configured in “Configuring and Applying a Firewall Filter for a Multifield Classifier (Required)” on page 532 to the schedulers configured in “Configuring Schedulers (Optional)” on page 548, as shown in Table 174.

Table 174: Sample diffserv-cos-map Scheduler Mapping

mf-classifier Forwarding Class	For CoS Traffic Type	diffserv-cos-map Scheduler
be-class	Best-effort traffic	be-scheduler
ef-class	Expedited forwarding traffic	ef-scheduler
af-class	Assured forwarding traffic	af-scheduler
nc-class	Network control traffic	nc-scheduler

To configure and apply scheduler maps 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 175.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following tasks:
 - To apply rules to logical interfaces, see “Configuring and Applying Virtual Channels (Optional)” on page 555.
 - To use adaptive shapers to limit bandwidth for Frame Relay, see “Configuring and Applying Adaptive Shaping (Optional)” on page 559.
 - To check the configuration, see “Verifying a DiffServ Configuration” on page 560.

Table 175: Configuring Scheduler Maps

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Class of service level in the configuration hierarchy.	In the configuration editor hierarchy, select Class of service .	From the top of the configuration hierarchy, enter edit class-of-service
Configure a scheduler map for DiffServ CoS.	<ol style="list-style-type: none"> 1. Click Add new entry next to Scheduler maps. 2. In the Map name box, type the name of the scheduler map—for example, diffserv-cos-map. 	Enter edit scheduler-maps diffserv-cos-map

Table 175: Configuring Scheduler Maps (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a best-effort forwarding class and scheduler.	<ol style="list-style-type: none"> 1. Click Add new entry next to Forwarding class. 2. In the Class name box, type the name of the previously configured best-effort forwarding class—be-class. 3. In the Scheduler box, type the name of the previously configured best-effort scheduler—be-scheduler. 4. Click OK. 	<p>Enter</p> <p>set forwarding-class be-class scheduler be-scheduler</p>
Configure an expedited forwarding class and scheduler.	<ol style="list-style-type: none"> 1. Click Add new entry next to Forwarding class. 2. In the Class name box, type the name of the previously configured expedited forwarding class—ef-class. 3. In the Scheduler box, type the name of the previously configured expedited forwarding scheduler—ef-scheduler. 4. Click OK. 	<p>Enter</p> <p>set forwarding-class ef-class scheduler ef-scheduler</p>
Configure an assured forwarding class and scheduler.	<ol style="list-style-type: none"> 1. Click Add new entry next to Forwarding class. 2. In the Class name box, type the name of the previously configured assured forwarding class—af-class. 3. In the Scheduler box, type the name of the previously configured assured forwarding scheduler—af-scheduler. 4. Click OK. 	<p>Enter</p> <p>set forwarding-class af-class scheduler af-scheduler</p>

Table 175: Configuring Scheduler Maps (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure a network control class and scheduler.	<ol style="list-style-type: none"> 1. Click Add new entry next to Forwarding class. 2. In the Class name box, type the name of the previously configured network control class—nc-class. 3. In the Scheduler box, type the name of the previously configured network control scheduler—nc-scheduler. 4. Click OK twice. 	<p>Enter</p> <pre>set forwarding-class nc-class scheduler nc-scheduler</pre>
Apply the scheduler map to an interface.	<ol style="list-style-type: none"> 1. Click Add new entry next to Interfaces. 2. In the Interface name box, type the name of the interface—for example, fe-0/0/0. 3. In the Scheduler map box, type the name of the previously configured scheduler map—diffserv-cos-map. 4. Click OK. 	<p>Enter</p> <pre>set interfaces fe-0/0/0 scheduler-map diffserv-cos-map</pre>

Configuring and Applying Virtual Channels (Optional)

You configure a virtual channel to set up queuing, packet scheduling, and accounting rules to be applied to one or more logical interfaces. You then must apply the virtual channel to a particular logical interface.

The following example shows how to create the virtual channels **branch1-vc**, **branch2-vc**, and **branch3-vc** and apply them in the firewall filter **choose-vc** to the Services Router's T3 interface **t3-1/0/0**.

To configure and apply virtual channels 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 176.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following tasks:
 - To assign resources, priorities, and profiles to output queues, see “Configuring Schedulers (Optional)” on page 548.

- To use adaptive shapers to limit bandwidth for Frame Relay, see “Configuring and Applying Adaptive Shaping (Optional)” on page 559.
- To check the configuration, see “Verifying a DiffServ Configuration” on page 560.

Table 176: Configuring and Applying Virtual Channels

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Class of service level in the configuration hierarchy.	In the configuration editor hierarchy, select Class of service .	From the top of the configuration hierarchy, enter edit class-of-service
Define the virtual channels branch1-vc , branch2-vc , branch3-vc , and the default virtual channel. You must specify a default virtual channel.	<ol style="list-style-type: none"> 1. Click Add new entry next to Virtual channels. 2. In the Channel name box, type the name of the virtual channel—for example, branch1-vc. 3. Click OK. 4. Create additional virtual channels for branch2-vc, branch3-vc, and default-vc. 	Enter set virtual-channels branch1-vc Repeat this statement for branch2-vc , branch3-vc , and default-vc .
Define the virtual channel group wan-vc-group to include the four virtual channels, and assign each virtual channel the scheduler map bestscheduler .	<ol style="list-style-type: none"> 1. Click Add new entry next to Virtual channel groups. 2. In the Group name box, type the name of the virtual channel group—wan-vc-group. 3. Click Add new entry next to Channel. 4. In the Channel name box, enter the name of the previously configured virtual channels—branch1-vc. 5. In the Scheduler map box, enter the name of the previously configured scheduler map—bestscheduler. 6. Click OK. 7. Add the virtual channels branch2-vc, branch3-vc, and default-vc. Select the Default box when adding the virtual channel default-vc. 	<ol style="list-style-type: none"> 1. Enter set virtual-channel-groups wan-vc-group branch1-vc scheduler-map bestscheduler 2. Repeat this statement for branch2-vc, branch3-vc, and default-vc. 3. Enter set virtual-channel-groups wan-vc-group default-vc default

Table 176: Configuring and Applying Virtual Channels (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Specify a shaping rate of 1.5 Mbps for each virtual channel within the virtual channel group.	<ol style="list-style-type: none"> 1. Click branch1-vc in the list of virtual channels. 2. Select the Shaping rate box. 3. Click Configure. 4. Select Absolute rate from the Rate choice box.. 5. In the Absolute rate box, enter the shaping rate—1.5m. 6. Add the shaping rate for the branch2-vc and branch3-vc virtual channels. 7. Click OK. 	<ol style="list-style-type: none"> 1. Enter <pre>set virtual-channel-groups wan-vc-group branch1-vc shaping-rate 1.5m</pre> 2. Repeat this statement for branch2-vc and branch3-vc.
Apply the virtual channel group to the logical interface t3-1/0/0.0 .	<ol style="list-style-type: none"> 1. Click Add new entry next to Interfaces. 2. In the Interface name box, type the name of the interface—t3-1/0/0. 3. Click Add new entry next to Unit. 4. In the Unit number box, type the logical interface unit number—0. 5. In the Virtual channel group box, type the name of the previously configured virtual channel group—wan-vc-group. 6. Click OK. 	<pre>Enter set interfaces t3-1/0/0 unit 0 virtual-channel-group wan-vc-group</pre>

Table 176: Configuring and Applying Virtual Channels (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the firewall filter choose-vc to select the traffic that is transmitted on a particular virtual channel.	<ol style="list-style-type: none"> 1. Navigate to the top of the configuration hierarchy and select Firewall. 2. Click Add new entry next to Filter. 3. In the Filter name box, enter the name of the firewall filter—choose-vc. 4. Click Add new entry next to Term. 5. In the Rule name box, enter the name of the firewall term—branch1. 6. Click Configure next to From. 7. Click Add new entry next to Destination address. 8. In the Address box, enter the IP address of the destination host—192.168.10.0/24. 9. Click OK twice. 10. On the firewall term page, click Configure next to Then. 11. Select Accept from the Designation box. 12. In the Virtual channel box, enter the name of the previously configured virtual channel—branch1-vc. 13. Click OK. 14. Repeat these steps for the virtual channels branch2-vc and branch3-vc. 	<ol style="list-style-type: none"> 1. From the top of the configuration hierarchy, enter <code>edit firewall</code> 2. Enter <code>set family inet filter choose-vc term branch1 from destination 192.168.10.0/24</code> 3. Enter <code>set family inet filter choose-vc term branch1 then accept</code> 4. Enter <code>set family inet filter choose-vc term branch1 then virtual-channel branch1-vc</code> 5. Repeat these steps for virtual channels branch2-vc and branch3-vc.
Apply the firewall filter choose-vc to output traffic on the t3-1/0/0.0 interface.	<ol style="list-style-type: none"> 1. Navigate to the top of the configuration hierarchy and select Interfaces. 2. Click t3-1/0/0 in the list of configured interfaces. 3. Click 0 in the list of configured logical units for the interface. 4. Click Edit next to Inet. 5. Click Configure next to Filter. 6. In the Output box, enter the name of the previously configured firewall filter—choose-vc. 7. Click OK. 	<ol style="list-style-type: none"> 1. From the top of the configuration hierarchy, enter <code>edit interfaces</code> 2. Enter <code>set t3-1/0/0 unit 0 family inet filter output choose-vc</code>

Configuring and Applying Adaptive Shaping (Optional)

You can use adaptive shaping to limit the bandwidth of traffic flowing on a Frame Relay logical interface. If you configure and apply adaptive shaping, the Services Router checks the backward explicit congestion notification (BECN) bit within the last inbound (ingress) packet received on the interface. If the BECN bit is set, the router limits the transmit bandwidth on the interface to the configured adaptive shaper maximum transmit rate. If the BECN bit is not set, the transmit bandwidth is not limited and is allowed to exceed the adaptive shaper rate.

For more information about adaptive shapers for a Frame Relay interface, see the *JUNOS Class of Service Configuration Guide*.

The following example shows how to create adaptive shaper `fr-shaper` and apply it to the Services Router's T1 interface `t1-0/0/2`. The adapter shaper limits the transmit bandwidth on the interface to 64 Kbps.

To configure and apply an adaptive shaper 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 177.
- 3. If you are finished configuring the router, commit the configuration.
- 4. Go on to one of the following tasks:
 - To assign resources, priorities, and profiles to output queues, see “Configuring Schedulers (Optional)” on page 548.
 - To apply rules to logical interfaces, see “Configuring and Applying Virtual Channels (Optional)” on page 555.
 - To check the configuration, see “Verifying a DiffServ Configuration” on page 560.

Table 177: Configuring and Applying an Adaptive Shaper

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Class of service level in the configuration hierarchy.	In the configuration editor hierarchy, select Class of service .	From the top of the configuration hierarchy, enter edit class-of-service

Table 177: Configuring and Applying an Adaptive Shaper (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Define the adaptive shaper name and maximum transmit rate.	<ol style="list-style-type: none"> Next to Adaptive Shapers, click Add new entry. In the Adaptive shaper name box, type fr-shaper. Next to Trigger, click Add new entry. Next to Becn, select the check box. Next to Shaping rate, select the check box and click Configure. From the Rate choice list, select Absolute rate. In the Absolute rate box, type 64k. Click OK three times. 	<p>Enter</p> <pre>set adaptive-shapers fr-shaper trigger becn shaping-rate 64k</pre>
Apply the adaptive shaper to the logical interface t1-0/0/2.0 .	<ol style="list-style-type: none"> Next to Interfaces, click Add new entry. In the Interface name box, type the name of the interface—t1-0/0/2. Next to Unit, click Add new entry. In the Unit number box, type the logical interface unit number—0. In the Adaptive shaper box, type the name of the adaptive shaper—fr-shaper. Click OK. 	<p>Enter</p> <pre>set interfaces t1-0/0/2 unit 0 adaptive-shaper fr-shaper</pre>

Verifying a DiffServ Configuration

To verify a DiffServ configuration, perform the following tasks:

- Verifying Multicast Session Announcements on page 561
- Verifying an Adaptive Shaper Configuration on page 561
- Verifying a Virtual Channel Configuration on page 562
- Verifying a Virtual Channel Group Configuration on page 562

Verifying Multicast Session Announcements

Purpose	Verify that the Services Router is listening to the appropriate groups for multicast Session Announcement Protocol (SAP) session announcements.
Action	From the CLI, enter the show sap listen command.
Sample Output	<pre>user@host> show sap listen Group Address Port 224.2.127.254 9875</pre>
What It Means	<p>The output shows a list of the group addresses and ports that SAP and SDP listen on. Verify the following information:</p> <ul style="list-style-type: none">■ Each group address configured, especially the default 224.2.127.254, is listed.■ Each port configured, especially the default 9875, is listed. <p>For more information about show sap listen, see the <i>JUNOS Routing Protocols and Policies Command Reference</i>.</p>

Verifying an Adaptive Shaper Configuration

Purpose	Verify the adaptive shaper trigger point and its associated transmit rate. Verify the class-of-service (CoS) configuration associated with an interface.
Action	From the CLI, enter the show class-of-service adaptive-shaper and show class-of-service interface t1-0/0/2 commands.
Sample Output	<pre>user@host> show class-of-service adaptive-shaper Adaptive shaper: fr-shaper, Index: 35320 Trigger type Shaping rate BECN 64000 bps user@host> show class-of-service interface t1-0/0/2 Physical interface: t1-0/0/2, Index: 137 Queues supported: 8, Queues in use: 4 Scheduler map: <default>, Index: 2 Logical interface: t1-0/0/2.0, Index: 69 Object Name Type Index Adaptive-shaper fr-shaper 35320 Classifier ipprec-compatibility ip 11</pre>
What It Means	<p>Verify the following information:</p> <ul style="list-style-type: none">■ The trigger type and shaping rate are consistent with the configured adaptive shaper.■ The adaptive shaper applied to the logical interface is displayed under Name.

Verifying a Virtual Channel Configuration

Purpose	Verify the virtual channel configuration on a logical interface. Verify the class-of-service (CoS) configuration associated with an interface.
Action	From the CLI, enter the <code>show class-of-service virtual-channel</code> command.
Sample Output	<pre>user@host> show class-of-service virtual-channel</pre>
What It Means	<pre>Virtual channel: vc-1 Index: 1</pre> <p>Verify that the name of the configured virtual channel is displayed in the output.</p>

Verifying a Virtual Channel Group Configuration

Purpose	Verify the virtual channel group configuration on a logical interface. Verify the class-of-service (CoS) configuration associated with an interface.
Action	From the CLI, enter the <code>show class-of-service virtual-channel-group</code> command.
Sample Output	<pre>user@host> show class-of-service virtual-channel-group</pre> <pre>Virtual channel group: vc-group, Index: 16321 Virtual channel: vc-1 Scheduler map: sc-map</pre>
What It Means	Verify that the name of the configured virtual channel group is displayed in the output.

Part 8

Index

Index

Symbols

[], in configuration statements	xxiv
{ }, in configuration statements	xxiv
(), in syntax descriptions	xxiv
< >, in syntax descriptions	xxiv
(pipe), in syntax descriptions	xxiv
#, comments in configuration statements	xxiv
1-port four-wire mode, SHDSL <i>See</i> ATM-over-SHDSL interfaces	
2-port two-wire mode, SHDSL <i>See</i> ATM-over-SHDSL interfaces	

A

ABM (Asynchronous Balance Mode), HDLC	88
ABRs <i>See</i> area border routers	
accept, filter action	498
access concentrator	
as a PPPoE server	152
naming for PPPoE	159
access control lists (ACLs) <i>See</i> stateless firewall filters	
ACLs <i>See</i> stateless firewall filters	
action modifiers, stateless firewall filters	
list of	448
setting	498
<i>See also</i> actions	
Action tab, stateless firewall filters	497
actions	
accept, setting	498
count modifier, setting	498
default, routing policy	440
discard, setting	498
final, routing policy	440
forwarding class modifier, setting	498
log modifier, setting	498
loss priority modifier, setting	499
modifiers, list of	448
NAT, list of	443
next term, setting	498
no action, setting	497
reject, setting	498
route list match types	462
routing instance, setting	498
routing policy	438
routing policy, summary of	439

sample modifier, setting	499
stateful firewall filters, list of	443
stateless firewall filters, list of	448
stateless firewall filters, setting actions (Quick Configuration)	497
stateless firewall filters, setting modifiers (Quick Configuration)	498
syslog modifier, setting	498
virtual channel modifier, setting	498
active routes, versus passive routes	239
adaptive shaping, applying CoS rules to logical interfaces	559
Add button	8
Add new entry link	10
address match conditions	447
address translation <i>See</i> NAT	
addresses	
BGP external peer address (configuration editor)	301
BGP internal peer address (configuration editor)	303
BGP local address (Quick Configuration)	298
BGP peer address (Quick Configuration)	298
IS-IS NETs	223
<i>See also</i> NETs	
IS-IS NSAP addresses	287
multicast ranges	399
physical, in data link layer	50
translating <i>See</i> NAT	
adjacencies, IS-IS	
hello PDUs for	224
<i>See also</i> IS-IS	
verifying	292
verifying (detail)	293
administrative groups, for MPLS path selection	325
administrative scoping	401
ADSL ports <i>See</i> ATM-over-ADSL interfaces	
advertisements <i>See</i> LSAs; route advertisements	
AF <i>See</i> DiffServ, assured forwarding	
aggregation, route	211
alternate mark inversion <i>See</i> AMI encoding	
AMI (alternate mark inversion) encoding	
E1	108
overview	58

T1.....	118	asynchronous networks	
Annex A PIMs		data stream clocking	79
ATM-over-ADSL interfaces	128	explicit clocking signal transmission	79
<i>See also</i> ATM-over-ADSL interfaces		overview	79
ATM-over-SHDSL interfaces	135	Asynchronous Response Mode (ARM), HDLC	88
<i>See also</i> ATM-over-SHDSL interfaces		Asynchronous Transfer Mode (ATM) interfaces	
ATM-over-SHDSL modes.....	135	<i>See</i> ATM-over-ADSL interfaces; ATM-over-SHDSL	
G.SHDSL PIMs, setting annex type on	137	interfaces	
operating modes.....	131	at-0/0/0 <i>See</i> ATM-over-ADSL interfaces;	
standards supported.....	72	ATM-over-SHDSL interfaces	
Annex B PIMs		ATM interfaces <i>See</i> ATM-over-ADSL interfaces;	
ATM-over-ADSL interfaces	128	ATM-over-SHDSL interfaces	
<i>See also</i> ATM-over-ADSL interfaces		ATM NLPID encapsulation	132, 139
ATM-over-SDSL interfaces	135	ATM PPP over AAL5 LLC encapsulation	132, 139
<i>See also</i> ATM-over-SHDSL interfaces		ATM PVC encapsulation	131, 137
ATM-over-SHDSL modes.....	135	ATM SNAP encapsulation.....	132, 139
G.SHDSL PIMs, setting annex type on	137	ATM VC multiplex encapsulation.....	132, 139
operating modes.....	131	ATM-over-ADSL interfaces	128
standards supported.....	72	adding.....	128
ANSI DMT operating mode.....	131	ADSL overview.....	71
anycast IPv6 addresses.....	94	ADSL systems.....	72
Apply button	8	ADSL topology	73
area border routers		ADSL2.....	73
adding interfaces	275	ADSL2 +	73
area ID (configuration editor)	275	ATM interface type.....	73
backbone area <i>See</i> backbone area		CHAP for PPPoA	133
backbone area interface.....	275	CHAP for PPPoE	161
description	219	encapsulation types, logical	132
areas <i>See</i> area border routers; backbone area; IS-IS,		encapsulation types, physical	131
areas; NSSAs; stub areas		logical properties	132
ARM (Asynchronous Response Mode), HDLC	88	operating modes.....	131
AS path		physical properties	129
description	229	PPPoE configuration.....	158
forcing by MED	230	PPPoE encapsulation	157
prepending.....	467	PPPoE session on.....	153
role in route selection	228	statistics	149
ASs (autonomous systems)		VCI	133
area border routers	219	verifying.....	145
AS number (configuration editor).....	301	verifying a PPPoA configuration	149
AS number (Quick Configuration)	298	verifying a PPPoE configuration	162–163
AS number, in VPNs	352	<i>See also</i> PPPoE; PPPoE over ATM-over-ADSL	
breaking into confederations.....	234	ATM-over-SHDSL interfaces	74
description	208	1-port four-wire mode.....	135
group AS number (configuration editor).....	301	1-port four-wire mode, setting	136
individual AS number (configuration editor)	301	2-port two-wire mode, overview	135
IS-IS networks.....	223	2-port two-wire mode, setting.....	136
LSPs through.....	318	adding.....	135
sample BGP confederation	307	annex type, setting	137
stub areas <i>See</i> stub areas		ATM interface type.....	73
sub-AS number	307	encapsulation types, logical	139
assured forwarding	547	encapsulation types, physical	137
asymmetrical digital subscriber line (ADSL)		logical properties	138
<i>See</i> ATM-over-ADSL interfaces		overview	74
Asynchronous Balance Mode (ABM), HDLC	88	VCI	140
		<i>See also</i> G.SHDSL PIMs	

authentication

- CHAP, for PPPoE interfaces 155
- OSPF, MD5 280
- OSPF, plain-text passwords 280
- RIPv2, MD5 260
- RIPv2, plain-text passwords 259

auto operating mode 131

Auto-RP 402

autonomous systems *See* ASs

B

B-channels

- description 74
- naming convention 171
- verifying 194

B8ZS encoding 58

BA classifiers *See* classifiers

backbone area

- area ID (configuration editor) 272
- area ID (Quick Configuration) 268
- area type (Quick Configuration) 269
- configuring 270
- description 220
- interface 275

backoff algorithm, collision detection 53

backup connection, ISDN 169

backward-explicit congestion notification (BECN)

- bits 82

bandwidth, for RSVP-signaled LSPs 338

bandwidth-on-demand

- dialer interface 181
- ISDN interface 183
- overview 181

bc-0/0/0 171

BECN (backward-explicit congestion notification)

- bits 82

behavior aggregate classifiers *See* classifiers

BERTs (bit error rate tests) 78

best-effort service 449

BGP (Border Gateway Protocol)

- AS number (Quick Configuration) 298
- See also* ASs (autonomous systems), AS number
- AS path 229
- See also* AS path
- confederations *See* BGP confederations
- enabling (Quick Configuration) 298
- export policy for CLNS 372
- external 227
- See also* EBGp
- external group type (configuration editor) 301
- external neighbor (peer) address (configuration editor) 301
- for CLNS VPN NLRI 375
- full mesh requirement 227, 296

- injecting OSPF routes into BGP 464
- internal 227
- See also* IBGP
- internal group type (configuration editor) 303
- internal neighbor (peer) address (configuration editor) 303
- local address (Quick Configuration) 298
- local preference 228
- MED metric 230
- origin value 230
- overview 225, 295
- peer address (Quick Configuration) 298
- peer AS number (Quick Configuration) 298
- peering sessions *See* BGP peers; BGP sessions
- point-to-point internal peer session (configuration editor) 302
- point-to-point peer session (configuration editor) 299
- policy to make routes less preferable 467
- Quick Configuration 297
- requirements 297
- route reflectors *See* BGP route reflectors
- route selection process 228
- See also* route selection
- route-flap damping 470
- router ID (Quick Configuration) 298
- routing policy (configuration editor) 303
- See also* routing policies
- sample BGP peer network 300
- sample confederation 307
- sample full mesh 302
- sample route reflector 304
- scaling techniques 231
- session establishment 226
- session maintenance 227
- verifying BGP configuration 310
- verifying BGP groups 309
- verifying BGP peers (neighbors) 308
- verifying peer reachability 311
- VPNs 351

BGP confederations

- confederation members 308
- confederation number 307
- creating (configuration editor) 306
- description 234, 296
- route-flap damping 470
- sample network 307
- sub-AS number 307

BGP groups

- cluster identifier (configuration editor) 305
- confederations (configuration editor) 306
- external group type (configuration editor) 301
- external, creating (configuration editor) 301
- group AS number (configuration editor) 301
- internal group type (configuration editor) 303

internal, creating (configuration editor)	303
internal, creating for a route reflector (configuration editor)	305
verifying	309
BGP messages	
to establish sessions	226
update, to maintain sessions	227
BGP page	297
BGP peers	
directing traffic by local preference	228
external (configuration editor)	299
internal (configuration editor)	302
internal, sample full mesh	302
internal, sample route reflector	304
peer address (Quick Configuration)	298
peer AS number (Quick Configuration)	298
point-to-point connections	226
routing policy (configuration editor)	303
<i>See also</i> routing policies	
sample peer network	300
sessions between peers	295
verifying	308, 310
verifying reachability	311
BGP route reflectors	
cluster (configuration editor)	305
cluster identifier (configuration editor)	305
cluster of clusters	233
creating (configuration editor)	303
description	231, 296
group type (configuration editor)	305
multiple clusters	232
sample IBGP network	304
BGP sessions	
configured at both ends	295
establishment	226
maintenance	227
point-to-point external (configuration editor)	299
point-to-point internal (configuration editor)	302
sample peering session	226
types	296
bipolar with 8-zero substitution (B8ZS) encoding	58
bit error rate tests (BERTs)	78
bit stuffing	61
bit-field logical operators, stateless firewall filters	448
bit-field match conditions	447
bit-field synonym match conditions	447
bootstrap router	402
Border Gateway Protocol <i>See</i> BGP	
br-0/0/0	171
braces, in configuration statements	xxiv
brackets	
angle, in syntax descriptions	xxiv
square, in configuration statements	xxiv
branches	398
<i>See also</i> multicast	

bridges, on LAN segments	54
BSR (bootstrap router)	402
buttons	11
Add (Quick Configuration)	8
Apply (Quick Configuration)	8
Cancel (J-Web configuration editor)	11
Cancel (Quick Configuration)	8
Commit (J-Web configuration editor)	11
CONFIG <i>See</i> CONFIG button	
Delete (Quick Configuration)	8
Discard (J-Web configuration editor)	11
OK (J-Web configuration editor)	11
OK (Quick Configuration)	8
Refresh (J-Web configuration editor)	11
<i>See also</i> radio buttons	

C

C-bit parity frame format	
enable or disable on T3 ports	122
overview	63
cables	
T1 cable length	119
T3 cable length	122
call setup, ISDN	76
Cancel button	
J-Web configuration editor	11
Quick Configuration	8
canceled a commit	31–32
canreach message, DLSw	420
capabilities exchange, DLSw	419
carrier sense multiple access with collision detection (CSMA/CD)	52
ccc protocol family	90
CE (customer edge) routers	
description	327
VPN task overview	346
VPN topology	344
<i>See also</i> VPNs	
Challenge Handshake Authentication Protocol	
<i>See</i> CHAP	
channel number, in interface name	49
channel service unit (CSU) device	85
CHAP (Challenge Handshake Authentication Protocol)	
E1 local identity	107
E3 local identity	110
enabling for PPPoA	133
enabling for PPPoE	161
enabling on ATM-over-ADSL interfaces	133
enabling on E1	107
enabling on E3	110
enabling on serial interfaces	124
enabling on T1	117
enabling on T3	121
overview	84
PPP links	84

- PPPoE 155
- serial interface local identity 124
- T1 local identity..... 117
- T3 local identity..... 121
- CHAP secret *See* CHAP, local identity
- checksum
 - E1 frame 108
 - E3 frame 112
 - overview 80
 - T1 frame 119
 - T3 frame 122
- circuit *See* Layer 2 circuits
- circuit establishment, DLSw 420
- Cisco NLPID encapsulation..... 132, 139
- class of service *See* CoS
- classful addressing..... 90
- classifiers
 - applying BA classifiers..... 542–543
 - default BA classifiers..... 455
 - default DSCP CoS classifier for DLSw..... 426
 - description 452
 - sample BA classification..... 456
 - sample BA classifier assignments..... 543
 - sample, for firewall filter 533
- clear system commit command 32
- CLI configuration editor
 - activating a configuration 31
 - ATM-over-ADSL interfaces 128
 - ATM-over-SHDLSL interfaces 135
 - BGP 299
 - CHAP on ATM-over-ADSL interfaces 133
 - CLNS 369
 - command summary..... 5
 - committing files 30
 - confirming a configuration..... 31
 - CoS with DiffServ 530
 - CRTTP 140
 - DLSw (basic) 420
 - DLSw CoS 424
 - exiting..... 22
 - IPSec tunnels 382
 - IS-IS 289
 - ISDN connections..... 172
 - managing files 34
 - modifying a configuration..... 25
 - MPLS traffic engineering 333
 - multicast network 406
 - network interfaces..... 126
 - network interfaces, adding..... 126
 - network interfaces, deleting..... 142
 - OSPF 269
 - PPPoE 155
 - PPPoE over ATM-over-ADSL 155
 - RIP..... 253
 - routing policies 460
 - saving files 37
 - starting 22
 - stateful firewall filters..... 480
 - stateless firewall filters 501
 - static routes 242
 - using show commands with 34
 - verifying a configuration 30
 - VPNs 346
- clickable configuration..... 8
 - committing..... 12
 - discarding changes 11
 - viewing and editing..... 8
 - See also* J-Web configuration editor
- CLNS (Connectionless Network Service) VPNs
 - BGP export policy..... 372
 - BGP, to carry CLNS VPN NLRI..... 375
 - displaying configurations..... 376
 - ES-IS..... 371
 - IS-IS 372
 - linking hosts 367
 - overview 368
 - requirements..... 369
 - static routes (without IS-IS)..... 374
 - verifying configuration 376
 - VPN routing instance 370
- clock rate, serial interface
 - DTE default reduction 68
 - values 125
- clocking
 - data stream clocking 79
 - E1 107
 - E3 110
 - explicit clocking signal transmission 79
 - overview 78
 - serial interface 125
 - serial interface, inverting the transmit clock .. 68, 125
 - serial interface, modes 67
 - T1..... 117
 - T3..... 121
- clusters *See* BGP route reflectors
- collision detection
 - backoff algorithm..... 53
 - overview 53
- coloring, link, for MPLS path selection 325
- combined stations, HDLC 88
- comments, in configuration statements xxiv
- commit and-quit command..... 31
- commit at command 31
- Commit button..... 11
- commit check command..... 30
- commit command..... 30
- commit confirmed command..... 31
- committed configuration
 - activating (CLI configuration editor) 31

canceling a commit (CLI configuration editor)	32
comparing two configurations	18
confirming (CLI configuration editor)	31
description	4
methods	17
replacing (CLI configuration editor)	32
rescue configuration (CLI configuration editor)	33
rescue configuration (J-Web)	21
scheduling (CLI configuration editor)	31
storage location	5
summaries	17
verifying (CLI configuration editor)	30
viewing previous (CLI configuration editor)	33
complete sequence number PDU (CSNP)	225
Compressed Real-Time Transport Protocol <i>See</i> CRTP	
confederations <i>See</i> BGP confederations	
CONFIG button	
default behavior	21, 33
disabling	33
return to factory configuration	21, 33
config-button <no-rescue> <no-clear> statement	33
configuration	
activating (CLI configuration editor)	31
adding a statement (CLI configuration editor)	26
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)	12
committing as a text file, with caution (J-Web)	13
confirming (CLI configuration editor)	31
copying a statement	27
deactivating a statement	29
deleting a statement	26
deleting with the CONFIG button	21, 33
disabling CONFIG button	33
discarding changes (J-Web)	11
downloading (J-Web)	20
editing (J-Web)	8
editing as a text file, with caution (J-Web)	13
history	16
<i>See also</i> configuration history	
inserting an identifier	28
J-Web options	5
loading new (CLI configuration editor)	34
loading previous (CLI configuration editor)	32
loading previous (J-Web)	21
locked, with the configure exclusive command	23
managing files (CLI configuration editor)	34
managing files (J-Web)	15
merging (CLI configuration editor)	35
modifying (CLI configuration editor)	25
modifying a statement (CLI configuration editor)	26
overriding (CLI configuration editor)	35
renaming an identifier	27
replacing configuration statements (CLI configuration editor)	35
requirements	6
rescuing (CLI configuration editor)	33
rescuing (J-Web)	21
rollback (CLI configuration editor)	32
rollback (J-Web)	21
saving (CLI configuration editor)	37
uploading (J-Web)	14
users-editors, viewing	18
verifying (CLI configuration editor)	30
viewing as a text file (J-Web)	12
configuration database, summary	17
configuration hierarchy, navigating	24
configuration history	
comparing files	18
database summary	17
displaying	16
downloading files	20
summary	17
users-editors, viewing	18
Configuration History page	16
configuration mode	
entering and exiting	22
using show commands in	34
configuration text	
editing and committing, with caution	13
viewing	12
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	10
configure private command	23
confirming a configuration	31
congestion control	
for Frame Relay, with DE bits	82
with DiffServ assured forwarding	546
connection process	
ISDN BRI interfaces	76
LCP, for PPP	83
serial interfaces	66
Connectionless Network Service <i>See</i> CLNS	
connectivity	
bidirectional (BGP)	225
bidirectional (OSPF)	217
unidirectional (RIP)	216
Constrained Shortest Path First <i>See</i> CSPF	
conventions	
for interface names	47
how to use this guide	xxii

notice iconsxxiii
text and syntaxxxiii
copy command27
CoS (class of service)
adaptive shaping for rules 559
assigning forwarding classes to output queues... 535
BA classifiers 542
configuration tasks 530
default BA classifiers 455
default DSCP classifier for DLSw 426
default forwarding class queue assignments 453
default scheduler settings 454
DiffServ benefits 449
See also DiffServ
DSCP rewrites 456
DSCPs 450
See also DSCPs
firewall filter for a multifield classifier 532
JUNOS components 452
JUNOS implementation 451
policer for firewall filter 531
preparation 530
RED drop profiles 546
rewrite rules 537
sample BA classification 456
scheduler maps 552
schedulers 548
ToS value for DLSw 426
uses 529
verifying adaptive shaper configuration 561
verifying multicast session announcements 561
verifying virtual channel configuration 562
verifying virtual channel group configuration 562
virtual channels for rules 555
cost, of a network path *See* path cost metrics
count, filter action modifier 498
CPE device, Services Router as, with PPPoE 151
See also PPPoE
CRC (cyclic redundancy check) 80
CRTP (Compressed Real-Time Transport Protocol)
E1 interfaces 140
overview 100
T1 interfaces 140
CSMA/CD (carrier sense multiple access with collision
detection) 52
CSNP (complete sequence number PDU) 225
CSPF (Constrained Shortest Path First)
constraints 325
disabling 338
link coloring 325
rules 325
CSPF algorithm *See* CSPF
CSU (channel service unit) device 85
curly braces, in configuration statements xxiv
customer edge routers *See* CE routers

customer premises equipment (CPE) device, Services
Router as, with PPPoE 151
See also PPPoE
customer support xxvi
contacting JTAC xxvi
cyclic redundancy check (CRC) 80

D

D-channel
description 74
naming convention 171
verifying 196
D4 framing 59
data communications equipment *See* DCE
data inversion
E1 108
T1 118
data link layer
error notification 50
flow control 51
frame sequencing 50
MAC addresses 50
network topology 50
physical addressing 50
purpose 50
sublayers 51
data link switching *See* DLSw
data service unit (DSU) device 85
data stream clocking 79
data terminal equipment *See* DTE
data-link connection identifiers (DLCIs) 82
Database Information page 16
dc-0/0/0 171
DCE (data communications equipment)
serial connection process 66
serial device 65
DCE clocking mode 67
DDR *See* dial-on-demand routing, ISDN
DE (discard eligibility) bits
BECE bits 82
FECE bits 82
deactivate command 29
deactivating configuration statements or identifiers 29
default gateway, static routing 241
defaults
BA classifiers 455
CoS forwarding class assignments 454
DSCP classifier for DLSw 426
junos-algs-outbound group, stateful firewall
filters 442
routing policy actions 440
setting for static routes 246
Delete button 8
delete command 26
Delete Configuration Below This Point radio button 11

Delete link	10	verifying ISDN interfaces	193
deleting		verifying ISDN status	193
current rescue configuration (CLI configuration editor)	33	verifying MPLS traffic engineering	338
current rescue configuration (J-Web)	22	verifying multicast IGMP versions	412
network interfaces	142	verifying multicast SAP and SDP configuration ..	412
denial-of-service attacks, preventing	506	verifying OSPF host reachability	284
dense routing mode, caution for use	400	verifying OSPF neighbors	282
<i>See also</i> multicast routing modes		verifying OSPF routes	283
designated router, OSPF		verifying OSPF-enabled interfaces	281
controlling election	280	verifying PIM mode and interface configuration ..	413
description	218	verifying PIM RPF routing table	414
destination prefix lengths	92	verifying PIM RPs	413
Deutsche Telekom UR-2 operating mode	131	verifying PPPoA for ATM-over-ADSL	
diagnosis		configuration	149
BERT	78	verifying PPPoE interfaces	164
displaying CLNS VPN configurations	376	verifying PPPoE over ATM-over-ADSL	
displaying firewall filter configurations	517	configuration	162–163
displaying firewall filter statistics	524	verifying PPPoE sessions	165
displaying IS-IS-enabled interfaces	291	verifying PPPoE statistics	166
displaying IS-IS-enabled interfaces (detail)	291	verifying PPPoE version information	166
displaying static routes in the routing table	247	verifying RIP host reachability	262
IS-IS adjacencies	292	verifying RIP message exchange	261
IS-IS adjacencies (detail)	293	verifying RIP-enabled interfaces	261
IS-IS neighbors	292	verifying stateful firewall filters	522
IS-IS neighbors (detail)	293	verifying virtual channel configuration	562
LDP neighbors	338	verifying VPN connectivity	364
LDP sessions	339	dial-on-demand connectivity, ISDN	
LDP-signaled LSP	340	dialer (dial-on-demand) filter, configuring	179
PPP magic numbers	85	dialer filter, applying	180
RSVP neighbors	341	overview	179
RSVP sessions	341	dial-on-demand filter <i>See</i> dialer filter	
RSVP-signaled LSP	342	dial-on-demand routing, ISDN	
traffic forwarding over LDP-signaled LSPs	340	dialer (dial-on-demand) filter, applying	185
verifying adaptive shaper configuration	561	dialer (dial-on-demand) filter, configuring	184
verifying B-channels	194	overview	184
verifying BGP configuration	310	dialer filter, ISDN	
verifying BGP groups	309	applying, for dial-on-demand connectivity	180
verifying BGP peer reachability	311	applying, for dial-on-demand routing	185
verifying BGP peers (neighbors)	308, 561	for dial-on-demand connectivity	179
verifying D-channel	196	for dial-on-demand routing	184
verifying dialer interfaces	197	dialer interface, ISDN	
verifying DLSw capabilities	427	adding	175
verifying DLSw circuit state	427	applying dialer (dial-on-demand) filter	180
verifying DLSw circuit state (detail)	428	bandwidth-on-demand	181
verifying DLSw LLC2 properties	427	dial-on-demand routing	184
verifying DLSw peers	428	dialer profiles	191
verifying DLSw peers (detail)	428	dialer watch	186
verifying DLSw reachability	429	limitations	171
verifying firewall filter actions	525	naming convention	171
verifying firewall filter DoS protection	526	restrictions	171
verifying firewall filter flood protection	526	verifying	197
verifying firewall filter handles fragments	527	dialer profiles, ISDN	191
verifying firewall filters with packet logs	523	dialer watch	
verifying IPsec tunnel operation	391	adding dialer watch interface	186
		ISDN interface	189

- overview 186
- Differentiated Services *See* DiffServ
- DiffServ (Differentiated Services)
 - assigning forwarding classes to output queues... 535
 - assured forwarding 546
 - BA classifiers 542
 - benefits for CoS 449
 - code points 450
 - See also* DSCPs
 - configuration tasks 530
 - default BA classifiers 455
 - default forwarding class queue assignments 453
 - default scheduler settings 454
 - DSCP rewrites 456
 - firewall filter for a multifield classifier 532
 - forwarding service classes 450
 - interoperability 450
 - JUNOS implementation 451
 - policer for firewall filter 531
 - preparation 530
 - RED drop profiles 546
 - rewrite rules 537
 - sample BA classification 456
 - scheduler maps 552
 - schedulers 548
 - uses 529
 - virtual channels for rules 555
- digital subscriber line (DSL) *See* ATM-over-ADSL
- interfaces; ATM-over-SHDSL interfaces; DSLAM connection
- Discard All Changes radio button 11
- Discard button 11
- Discard Changes Below This Point radio button 11
- discard eligibility bits *See* DE bits
- discard interface 98
- discard rule
 - firewall filters 441
 - stateful firewall filters 442
 - stateless firewall filters 444
- discard, filter action 498
- discarding configuration changes 11
- discovery packets, PPPoE 86, 154
- Distance Vector Multicast Routing Protocol 401
- distance-vector routing protocols 213
 - See also* RIP
- dl0 171
- DLCIs (data-link connection identifiers) 82
- DLSw (data link switching)
 - basic configuration 420
 - canureach message 420
 - capabilities exchange 419
 - circuit establishment 420
 - CoS classification of DLSw packets 424
 - icanreach message 420
 - LLC2 properties on Ethernet interfaces 421
 - local router configuration 422
 - overview 418
 - peers *See* DLSw peers
 - preparation 420
 - remote router configuration 423
 - sample DLSw network 419
 - sample peer router values 422
 - SNA forwarding 418
 - SSP 419
 - stages of operation 419
 - ToS precedence for DLSw packets 424
 - verifying capabilities 427
 - verifying DLSw circuit state 427
 - verifying DLSw circuit state (detail) 428
 - verifying DLSw peers 428
 - verifying DLSw peers (detail) 428
 - verifying DLSw reachability 429
 - verifying LLC2 properties 427
- DLSw peers
 - local 422
 - remote 423
 - verifying 428
 - verifying (detail) 428
- documentation set
 - comments on xxvi
- domains
 - broadcast domains 55
 - collision domains 54
- DoS (denial-of-service) attacks, preventing 506
- dotted decimal notation 91
- downloading, configuration files (J-Web) 20
- downstream interfaces 398
 - See also* multicast
- DS1 ports *See* T1 ports
- DS1 signals
 - E1 and T1 57
 - See also* E1 interfaces; T1 interfaces
 - multiplexing into DS2 signal 60
- DS2 signals
 - bit stuffing 61
 - frame format 61
- DS3 ports *See* T3 ports
- DS3 signals
 - DS3 C-bit parity frame format 63
 - M13 frame format 62
- dsc interface 98
- DSCPs (DiffServ code points)
 - corresponding forwarding service classes 450
 - default classifier for DLSw 426
 - default forwarding class queue assignments 453
 - description 450
 - matching with a filter 495
 - matching with an IPv4 filter 496
 - replacing with rewrite rules 538
 - rewrites 456

sample BA classification.....	456
DSL <i>See</i> ATM-over-ADSL interfaces; ATM-over-SHDSDL interfaces; DSLAM connection	
DSL access multiplexer (DSLAM) connection <i>See</i> DSLAM connection	
DSLAM connection	
ATM-over-ADSL interface for.....	128
ATM-over-SHDSDL interface for.....	135
PPPoE over ATM-over-ADSL topology.....	153
DSU (data service unit) device	85
DTE (data terminal equipment)	
default clock rate reduction	68
serial connection process	66
serial device.....	65
DTE clocking mode	67
DVMRP (Distance Vector Multicast Routing Protocol) ..	401
dynamic LSPs	321
dynamic routing	210

E

E1 interfaces	56
AMI encoding.....	58
data stream	57
encoding	58
framing.....	59
HDB3 encoding.....	58
loopback	60
overview	57
signals.....	57
<i>See also</i> E1 ports	
E1 ports	56
adding CRTP	140
CHAP	107
clocking	107
configuring.....	105
data inversion.....	108
encapsulation type	107
fractional, channel number	49
frame checksum	108
framing.....	108
logical interfaces.....	107
MTU	107
overview	57
time slots.....	108
<i>See also</i> E1 interfaces	
E3 interfaces	60
bit stuffing	61
data stream	60
DS3 framing	61
multiplexing on	61
overview	60
<i>See also</i> E3 ports	
E3 ports	60
CHAP	110
clocking	110

configuring.....	108
encapsulation type	110
frame checksum.....	112
logical interfaces.....	110
MTU	110
overview	60
<i>See also</i> E3 interfaces	
EBGP (external BGP)	
description	227
route-flap damping	470
sample network.....	302
edit command	24
Edit Configuration page	9
Edit Configuration Text page	14
Edit link	10
EGPs (exterior gateway protocols)	208
egress router <i>See</i> LSPs; outbound router	
EIA-232.....	69
EIA-422.....	70
EIA-449.....	70
EIA-530.....	69
Encapsulating Security Payload Security Parameter Index (ESP SPI) values, matching with a filter	497
encapsulation type.....	107
ATM-over-ADSL logical interfaces	132
ATM-over-ADSL physical interfaces	131
ATM-over-SHDSDL logical interfaces	139
ATM-over-SHDSDL physical interfaces.....	137
E1	107
E3	110
Frame Relay	81
HDLC	87
overview	81
PPP	83
PPPoE	151
PPPoE for Ethernet	156
PPPoE, over ATM for ADSL.....	157
PPPoE, overview.....	86
serial interfaces.....	124
T1.....	117
T3.....	121
<i>See also</i> packet encapsulation	
encoding	
AMI	58
B8ZS.....	58
HDB3.....	58
End System-to-Intermediate System <i>See</i> ES-IS	
EROs (Explicit Route Objects)	
loose hops.....	324
strict hops	324
error notification, in the data link layer.....	50
ES-IS (End System-to-Intermediate System)	
for a PE router in a CLNS island	371
overview	368
ESF (extended superframe) framing.....	59

ESP SPI (Encapsulating Security Payload Security
Parameter Index) values, matching with a filter 497

Ethernet interfaces 52

- access control 52
- broadcast domains 55
- collision detection 53
- collision domains 54
- CSMA/CD 52
- DLSw, LLC2 properties for 421
- frame format 55
- IS-IS, NET address 290
- overview 52
- See also* Fast Ethernet ports

Ethernet over ATM encapsulation 131, 137

Ethernet over LLC encapsulation 132, 139

Ethernet ports *See* Fast Ethernet ports

ETSI operating mode 131

EU-64 addresses 51

exact route list match type 462

exit command

- to leave configuration mode 23
- to navigate the configuration hierarchy 24

exit configuration-mode command 23

explicit clocking signal transmission 79

Explicit Route Objects *See* EROs

export routing policy, for Layer 2 VPNs 361

export statement, for routing policies 440

extended superframe (ESF) framing 59

exterior gateway protocols 208

external BGP *See* EBGp

F

failover connection, ISDN 169

Fast Ethernet ports 52

- CHAP for PPPoA 133
- CHAP for PPPoE 161
- configuring 113
- LLC2 properties for DLSw 422
- logical interfaces 115
- MTU 115
- overview 52
- PPPoE configuration 158
- PPPoE encapsulation 156
- PPPoE session on 153
- See also* Ethernet interfaces

FCS (frame check sequence)

- checksums 80
- CRCs 80
- overview 80
- two-dimensional parity 80

fe-0/0/0

- disabling PIM on 409
- management interface 99

FEAC C-bit condition indicators 65

FECN (forward-explicit congestion notification) bits 82

file management

- configuration files (CLI configuration editor) 34
- configuration files (J-Web) 15

filters *See* firewall filters; stateful firewall filters;
stateless firewall filters

firewall filters

- applying CoS rules to logical interfaces 555
- displaying configurations 517
- displaying statistics 524
- multifield classifier filter terms 532
- overview 440
- policer for 531
- sample classifier terms 533
- stateful firewall filters 441
- See also* stateful firewall filters
- stateless firewall filters 441
- See also* stateless firewall filters
- term number caution 442
- verifying configuration 517
- verifying flood protection 526
- verifying fragment handling 527
- verifying packet logging 523

Firewall Filters configuration pages

- field summary 489
- initial page 487
- match conditions and actions page 488

Firewall Filters interface assignment pages

- available interfaces and filter status page 500
- field summary 501

Firewall/NAT application page 478

- field summary 479

Firewall/NAT main page 477

- field summary 479

flap damping 470

- parameters 470

flooding, preventing 506

flow control

- actions in routing policies 439
- data link layer 51

font conventions xxiii

forward-explicit congestion notification (FECN) bits 82

forwarding classes

- assigning to output queues 536
- default queue assignments 453
- description 452
- filter action modifier, setting 498
- mapping to schedulers 553
- matching with a filter 497
- policy to group source and destination prefixes .. 466
- sample BA classification 456
- sample mappings 553

forwarding policy options 452

forwarding states, multicast notation 399

forwarding table

- controlling OSPF routes in 278

controlling static routes in	238, 245
description	209
MED to determine routes in	230
four-wire mode (1 port), SHDSL <i>See</i> ATM-over-SHDSL interfaces	
FPC (PIM slot on a Services Router) <i>See</i> PIMs	
fragment offset, matching with a filter	496
frame check sequence <i>See</i> FCS	
Frame Relay encapsulation	
congestion control	82
DLCIs	82
overview	81
PVCs	82
SVCs	82
virtual circuits	82
Frame Relay network, typical	81
frames	
DS2 M-frame format	61
DS3 C-bit parity frame format	63
DS3 M13 frame format	62
Ethernet frame format	55
sequencing, data link layer	50
framing	
E1	108
T1	118
T3	122
FRF.15 and FRF.16	100
from statement, routing policy match conditions	436
full mesh requirement	
description	227
fulfilling with confederations	234
fulfilling with route reflectors	231
sample network	302
fxp0 interface (not supported)	96

G

*,G notation, for multicast forwarding states	399
G.SHDSL PIMs	74
Annex A or Annex B modes	135
configuring	135
default mode	136
standard supported	74
<i>See also</i> ATM-over-SHDSL interfaces	
gateway, local and remote, for IPSec service sets	385
global unicast IPv6 addresses	94
glossary	
CLNS	367
configuration	3
CoS	433
DLSw	417
firewall filters	433
interfaces	42
ISDN	169
MPLS	315
multicast	395

ports	42
PPPoE	151
routing policies	433
routing protocols	203
VPNs	315
gr-0/0/0 interface	96
gre interface	96
groups	
BGP <i>See</i> BGP groups	
default junos-algs-outbound group, for stateful firewall filters	442
OSPF areas	272
RIP routers	253

H

handling packet fragments	513
HDB3 encoding	58
HDLC (High-Level Data Link Control)	
encapsulation	87
HDLC operational modes	88
HDLC stations	87
hello PDUs	224
hierarchy, configuration	24
high-density bipolar 3 code (HDB3) encoding	58
High-Level Data Link Control <i>See</i> HDLC	
history <i>See</i> configuration history	
hold time, to maintain a session	227
hop count, maximizing	214
<i>See also</i> RIP	
host reachability	
verifying for a RIP network	262
verifying for an OSPF network	284
hostname	
for PPPoA CHAP	134
for PPPoE CHAP	162
IS-IS identifier-to-hostname mapping	288
how to use this guide	xxii

I

IBGP (internal BGP)	
description	227
full mesh (configuration editor)	302
full mesh requirement	296
sample network	302
sample route reflector	304
IBM networking <i>See</i> DLSw	
icanreach message, DLSw	420
ICMP (Internet Control Message Protocol), policers	508
ICMP packets, matching with a filter	495
identifier link	10
identifiers, configuration	
adding or modifying	26
deactivating	29
deleting	26
inserting	28

- renaming27
- IGMP (Internet Group Management Protocol)
 - IGMPv1 402
 - IGMPv2 402
 - IGMPv3 402
 - setting the version 407
 - verifying the version 412
- IGPs (interior gateway protocols) 353
 - overview 208
 - VPNs 353
 - See also* OSPF; RIP
- IKE (Internet Key Exchange)
 - description 380
 - preshared key (configuration editor) 386
 - preshared key (Quick Configuration) 382
- import routing policy, for Layer 2 VPNs 360
- import statement, for routing policies 440
- inbound router, in an LSP 319
- incoming metric (RIP)
 - description 250
 - modifying 257
- inet protocol family 89
 - MTU value for PPPoE 160
- inet routing table 410
- inet6 protocol family 89
 - MTU value for PPPoE 160
- ingress router *See* inbound router; LSPs
- injecting routes 465
- input filters, assigning to interfaces 501
- insert command 28
- inserting configuration identifiers 28
- Integrated Services Digital Network *See* ISDN
- interface groups, matching with a filter 494
- interface naming conventions 47
- interface sets, matching with a filter 493
- interfaces 41
 - ATM-over-ADSL interfaces 71
 - ATM-over-SHDSL interfaces 74
 - clocking 78
 - data link layer 50
 - E1 interfaces 56
 - E3 interfaces 60
 - Ethernet interfaces 52
 - FCS 80
 - G.SHDSL interfaces 74
 - IPv4 addressing 90
 - IPv6 addressing 93
 - ISDN interfaces 74
 - logical properties 88
 - overview 41
 - physical encapsulation 81
 - See also* encapsulation types
 - physical properties 77
 - protocol families 89
 - serial interfaces 65
 - special interfaces 96
 - T1 interfaces 56
 - T3 interfaces 60
 - VLANs 95
 - See also* ATM-over-ADSL interfaces;
 - ATM-over-SHDSL interfaces; ISDN interfaces;
 - loopback interface; management interfaces;
 - network interfaces; services interfaces; special
 - interfaces; ports
- Interfaces page 104
 - for E1 106
 - for E3 109
 - for Fast Ethernet 114
 - for serial interfaces 123
 - for T1 116
 - for T3 (DS3) 120
- interior gateway protocols *See* IGPs
- Intermediate System-to-Intermediate System *See* IS-IS
- internal BGP *See* IBGP
- Internet Control Message Protocol policers 508
- Internet Group Management Protocol *See* IGMP
- Internet Key Exchange *See* IKE
- Internet routing, with BGP 295
- invalid configuration, replacing
 - with J-Web 21
 - with the CLI 33
- invalid routes, rejecting 464
- inverting the transmit clock 125
- IP addresses 90
 - as IS-IS system identifiers 288
 - See also* addresses; IPv4 addressing; IPv6 addressing
- IP options, matching with a filter 497
- IP Security *See* IPSec
- ip-0/0/0 interface 97
- ip-ip interface 97
- IPSec (IP Security)
 - ESP SPI values, matching with a filter 497
 - IKE *See* IKE
 - security associations 380
 - tunnels *See* IPSec tunnels
 - verifying tunnels 391
- IPSec security associations 380
 - See also* IKE
- IPSec tunnels
 - IKE key (configuration editor) 386
 - IKE key (Quick Configuration) 382
 - IPSec rule (configuration editor) 387
 - local endpoint (Quick Configuration) 382
 - NAT pools (configuration editor) 389
 - outgoing traffic filters 380
 - overview 379
 - private addresses (Quick Configuration) 382
 - Quick Configuration 380
 - remote endpoint (Quick Configuration) 382

requirements.....	380	verifying neighbors.....	292
services interfaces (configuration editor).....	383	verifying neighbors (detail).....	293
services sets (configuration editor).....	384	with CLNS.....	368
stateful firewall filter rules (configuration editor).....	387	ISDN BRI interfaces.....	74
verifying.....	391	adding an interface.....	172
IPSec Tunnels page.....	381	B-channel interface.....	171
field summary.....	382	bandwidth-on-demand.....	181
IPv4 addressing.....		call setup.....	76
assigning for PPPoE.....	160	connection initialization.....	76
classful addressing.....	90	D-channel interface.....	171
dotted decimal notation.....	91	dial-on-demand connectivity.....	179
MAC-48 address format.....	51	dial-on-demand routing.....	184
overview.....	90	dial-on-demand routing, with OSPF.....	190
subnets.....	91	dialer interface.....	171
VLSMs.....	92	dialer interface, adding.....	175
IPv4 filters.....		dialer profiles.....	191
assigning to interfaces (Quick Configuration)....	499	dialer watch.....	186
creating and editing (Quick Configuration).....	487	ISDN channels.....	74
<i>See also</i> stateless firewall filters		naming conventions.....	171
IPv4 MTU value, PPPoE.....	160	NT1 devices.....	75
IPv6 addressing.....		overview.....	74
address format.....	93	PIMs supported.....	171
address scope.....	94	requirements.....	172
address structure.....	94	S/T interfaces.....	75, 171
address types.....	94	secondary (backup) connection.....	178
assigning for PPPoE.....	160	session establishment.....	76
overview.....	93	switch types supported.....	174
IPv6 filters.....		typical network.....	75
assigning to interfaces (Quick Configuration)....	499	U interface.....	76, 171
creating and editing (Quick Configuration).....	487	verifying B-channels.....	194
<i>See also</i> stateless firewall filters		verifying D-channel.....	196
IPv6 MTU value, PPPoE.....	160	verifying dialer interfaces.....	197
IPv6 support.....	203	verifying ISDN interfaces.....	193
IS-IS (Intermediate System-to-Intermediate System)		verifying ISDN status.....	193
adjacency establishment with hello PDUs.....	224	<i>See also</i> ISDN connections	
areas.....	223	ISDN connections.....	169
ASs.....	223	adding an interface.....	172
CSNPs.....	225	bandwidth-on-demand.....	181
enabling on router interfaces.....	289	configuring.....	169
for CLNS route exchange.....	372	dial-on-demand connectivity.....	179
hello PDUs.....	224	dial-on-demand routing.....	184
LSPs.....	224	dial-on-demand routing, with OSPF.....	190
NETs.....	223	dialer (dial-on-demand) filter, applying.....	180, 185
NSAP addresses.....	287	dialer (dial-on-demand) filter, configuring..	179, 184
overview.....	222, 287	dialer interface.....	171
path selection.....	224	dialer interface for bandwidth-on-demand.....	181
preparation.....	289	dialer interface, adding.....	175
PSNPs.....	225	dialer profiles.....	191
system identifiers.....	224	dialer watch.....	186
<i>See also</i> system identifiers		dialer watch interface.....	186
verifying adjacencies.....	292	interface naming conventions.....	171
verifying adjacencies (detail).....	293	ISDN interface for bandwidth-on-demand.....	183
verifying interface configuration.....	291	ISDN interface for dialer watch.....	189
verifying interface configuration (detail).....	291	ISDN interface types.....	171
		overview.....	170

- requirements..... 172
 - secondary (backup) connection..... 178
 - switch types supported..... 174
 - verifying B-channels..... 194
 - verifying D-channel..... 196
 - verifying dialer interfaces..... 197
 - verifying ISDN interfaces..... 193
 - verifying ISDN status..... 193
 - See also* ISDN BRI interfaces
 - ISO network addresses, for IS-IS routers..... 287
 - ISO protocol family..... 89
 - ITU Annex B non-UR-2 operating mode..... 131
 - ITU Annex B UR-2 operating mode..... 131
 - ITU DMT BIS operating mode..... 131
 - ITU DMT operating mode..... 131
- J**
- J-series
- BGP routing..... 295
 - CLNS VPNs..... 367
 - configuration tools..... 3
 - CoS overview..... 449
 - CoS with DiffServ..... 529
 - DLSw..... 417
 - firewall filter overview..... 440
 - firewall filters..... 475
 - IBM networking..... 417
 - interfaces overview..... 41
 - IPSec tunnels..... 379
 - IS-IS protocol..... 287
 - ISDN connections..... 169
 - MPLS for VPNs overview..... 315
 - MPLS traffic engineering..... 331
 - multicast..... 405
 - multicast overview..... 395
 - NAT..... 475
 - network interfaces..... 103
 - OSPF routing..... 265
 - PPPoE..... 151
 - release notes, URL..... xxi
 - RIP routing..... 249
 - routing policies..... 459
 - routing policy overview..... 435
 - routing protocols overview..... 203
 - static routing..... 237
 - VPNs..... 343
- J-Web configuration editor
- ATM-over-ADSL interfaces..... 128
 - ATM-over-SHDSL interfaces..... 135
 - BGP..... 299
 - CHAP on ATM-over-ADSL interfaces..... 133
 - clickable configuration, committing..... 12
 - clickable configuration, discarding changes..... 11
 - clickable configuration, editing..... 8
 - CLNS..... 369
 - committing a text file, with caution..... 13
 - configuration text, viewing..... 12
 - CoS with DiffServ..... 530
 - CRTTP..... 140
 - DLSw (basic)..... 420
 - DLSw CoS..... 424
 - editing a text file, with caution..... 13
 - IPSec tunnels..... 382
 - IS-IS..... 289
 - ISDN connections..... 172
 - managing files..... 15
 - MPLS traffic engineering..... 333
 - multicast network..... 406
 - network interfaces..... 126
 - network interfaces, adding..... 126
 - network interfaces, deleting..... 142
 - OSPF..... 269
 - PPPoE..... 155
 - PPPoE over ATM-over-ADSL..... 155
 - RIP..... 253
 - routing policies..... 460
 - stateful firewall filters..... 480
 - stateless firewall filters..... 501
 - static routes..... 242
 - uploading a file..... 14
 - VPNs..... 346
- J-Web interface..... 5
- comparing configuration differences..... 18
 - configuration history..... 16
 - See also* configuration history
 - configuration options..... 5
 - See also* J-Web configuration editor
- JTAC (Juniper Networks Technical Assistance Center)
- See* technical support
- Juniper Networks Technical Assistance Center
- See* technical support
- JUNOS Internet software
- CoS components..... 452
 - CoS functions..... 451
 - DiffServ implementation..... 451
 - ISDN connections..... 169
 - release notes, URL..... xxi
- junos-algs-outbound group, for stateful firewall filters..... 442
- K**
- keepalive interval, for LDP-signaled LSPs..... 335
 - keepalive messages, for session hold time..... 227
- L**
- Label Distribution Protocol *See* LDP
 - label switching..... 318
 - label-switched paths *See* LSPs
 - label-switching routers (LSRs)..... 319

labels, MPLS.....	320	operation.....	322
label operations.....	320	overview.....	332
PHP.....	321	requirements.....	332
LANs		verifying LSPs.....	340
bridges on LAN segments.....	54	verifying neighbors.....	338
collision domains.....	54	verifying sessions.....	339
repeaters on LAN segments.....	54	verifying traffic forwarding.....	340
topology.....	95	LDP neighbors, verifying.....	338
Layer 2 circuits		LDP-signaled LSP <i>See</i> LDP	
AS number.....	352	leaves.....	398
basic, description.....	345	<i>See also</i> multicast	
encapsulation.....	348	Level 1 areas, IS-IS.....	223
IGPs.....	353	Level 2 areas, IS-IS.....	223
MPLS.....	349	line buildout	
neighbor address.....	356	T1.....	119
participating interfaces.....	347	T3.....	122
signaling protocols.....	353	line speed, serial interface.....	125
task overview.....	346	line timing.....	67
verifying PE router connections.....	365	link coloring, for MPLS path selection.....	325
verifying PE router interfaces.....	365	link services.....	100
virtual circuit ID.....	356	<i>See also</i> ls-0/0/0	
Layer 2 VPNs		link states, verifying.....	143
AS number.....	352	link-local unicast IPv6 addresses.....	94
basic, description.....	344	link-state advertisements <i>See</i> LSAs	
BGP.....	351	link-state PDUs <i>See</i> LSPs	
encapsulation.....	348	LLC (Logical Link Control) properties for DLSw	
export routing policies.....	361	configuration.....	421
IGPs.....	353	verification.....	427
import routing policies.....	360	lo0 interface functions.....	99
MPLS.....	349	<i>See also</i> loopback interface	
overview.....	329	lo0.16385, internal loopback address.....	97
participating interfaces.....	347	load command.....	35
routing instance.....	357	load merge command.....	35
signaling protocols.....	353	load override command.....	35
task overview.....	346	load patch command.....	35
verifying PE router connections.....	365	load replace command.....	35
verifying PE router interfaces.....	365	loading a configuration file	
Layer 3 VPNs		CLI configuration editor.....	34
AS number.....	352	downloading (J-Web).....	20
basic, description.....	345	rollback (J-Web).....	21
BGP.....	351	rollback command.....	32
IGPs.....	353	uploading (J-Web).....	14
overview.....	329	without specifying full hierarchy.....	35
participating interfaces.....	347	local preference	
route target.....	357	description.....	228
routing instance.....	357	high value preferred.....	229
routing policies.....	363	role in route selection.....	228
signaling protocols.....	353	local router, DLSw.....	422
task overview.....	346	<i>See also</i> DLSw peers	
verifying PE router connections.....	365	local tunnel endpoint, IPSec.....	382
LCP (Link Control Protocol), connection process.....	83	locked configuration.....	23
LDP (Label Distribution Protocol)		logging packet header information.....	498
and OSPF for VPNs.....	353	logical interfaces	
LDP-signaled LSPs.....	333	adaptive shaping for.....	559
messages.....	322	adding (configuration editor).....	128

- ATM-over-ADSL 132
 - ATM-over-SHDSL 138
 - CoS rules for 555, 559
 - E1 107
 - E3 110
 - Fast Ethernet 115
 - inside services interface, IPSec 383
 - outside services interface, IPSec 383
 - serial 124
 - T1 117
 - T3 121
 - virtual channels for 555
 - Logical Link Control (LLC) properties for DLSw *See* LLC
 - logical units
 - adding (configuration editor) 128
 - E1 interface 107
 - E3 interface 110
 - Fast Ethernet interface 115
 - number in interface name 49
 - pp0 interface 158
 - PPPoE encapsulation 156
 - PPPoE over ATM-over-ADSL encapsulation 157
 - serial interface 124
 - T1 interface 117
 - T3 interface 121
 - long buildout *See* line buildout
 - longer route list match type 462
 - loop clocking mode 67
 - loopback address
 - for PE routers in VPNs 353
 - internal, lo0.16385 97
 - loopback interface
 - applying stateless firewall filters to (configuration editor) 516
 - functions 99
 - NET on for IS-IS 290
 - loopback signals, E1 and T1 60
 - loose hops, RSVP 324
 - loss priority, CoS 452
 - ls-0/0/0
 - adding CRTP 140
 - interface description 97
 - LSAs (link-state advertisements)
 - description 218
 - three-way handshake 218
 - lsi interface 97
 - LSPs (label-switched paths)
 - bandwidth 338
 - description 318
 - disabling CSPF 338
 - dynamic LSPs 321
 - for RSVP in a VPN 350
 - keepalive interval for LDP link 335
 - label operations 320
 - label switching 318
 - labels 320
 - LDP 322
 - LDP-signaled LSPs 333
 - LSR types 319
 - overview 331
 - PHP 321
 - RSVP 322
 - RSVP-signaled LSPs 335
 - static LSPs 321
 - verifying LDP-signaled LSPs 338
 - verifying RSVP-signaled LSPs 341
 - LSPs (link-state PDUs)
 - CSNPs 225
 - overview 224
 - PSNPs 225
 - LSRs (label-switching routers) 319
 - lt-0/0/0 interface 97
- ## M
- M13 frame format 62
 - MAC (media access control) *See* MAC addresses
 - MAC addresses
 - as IS-IS system identifiers 288
 - EUI-64 addresses 51
 - MAC-48 address format 51
 - overview 51
 - physical addressing 50
 - MAC-48 addresses 51
 - magic numbers, PPP 85
 - management interfaces
 - disabling PIM on 409
 - overview 99
 - managing files *See* file management
 - manuals
 - comments on xxvi
 - mapping, CoS forwarding classes to schedulers 553
 - match conditions
 - routing policy 436
 - routing policy, summary of 436
 - stateful firewall filters and NAT 443
 - stateless firewall filters 445
 - stateless firewall filters, summary of 446
 - Match Destination tab, stateless firewall filters 491
 - Match Interface tab, stateless firewall filters 493
 - Match Network tab, stateless firewall filters 494
 - Match Packet and Network tab, stateless firewall filters 494
 - Match Source or Destination tab, stateless firewall filters 492
 - Match Source tab, stateless firewall filters 491
 - match types 462
 - maximum hop count, RIP 214
 - maximum transmission unit *See* MTU
 - MED (multiple exit discriminator)
 - description 230

role in route selection	228	T1	117
media access control <i>See</i> MAC addresses		T3	121, 124
media types supported	46	mtun interface	97
merging a configuration file	35	multiarea network, OSPF	272
example	37	multicast	
messages, LDP	322	administrative scoping	401
metrics <i>See</i> path cost metrics		architecture	398
MF classifier	532	Auto-RP	402
MLFR (Multilink Frame Relay)	100	BSR	402
MLFR FRF.15 and FRF.16	100	downstream interface	398
mlfr-end-to-end protocol family	90	DVMRP	401
mlfr-uni-nni protocol family	90	forwarding state notation	399
MLPPP (Multilink Point-to-Point Protocol)	100	*,G notation	399
mlppp protocol family	90	IGMP <i>See</i> IGMP	
MPLS (Multiprotocol Label Switching)	326	IP address ranges	399
dynamic LSPs	321	MSDP	403
label operations	320	network elements	399
label switching	318	overview	395
labels	320	PGM	403
Layer 2 VPNs and Layer 2 circuits	349	PIM dense mode <i>See</i> PIM	
LDP	322	PIM source-specific multicast (SSM)	402
LSP for RSVP in a VPN	350	PIM sparse mode <i>See</i> PIM	
LSPs	318	preparation	406
LSR types	319	preventing routing loops	400
overview	315	protocols	401
PHP	321	reverse-path forwarding (RPF)	400
RSVP	322	routing modes <i>See</i> multicast routing modes	
static LSPs	321	S,G notation	399
traffic engineering <i>See</i> MPLS traffic engineering		SAP and SDP <i>See</i> SAP; SDP	
verifying	338	session announcements	406
<i>See also</i> VPNs		shortest-path tree (SPT)	401
MPLS protocol family	89	static RP	408
MTU value for PPPoE	160	<i>See also</i> RP	
MPLS traffic engineering		subnetwork leaves and branches	398
LDP signaling	332	upstream interface	398
LDP-signaled LSPs	333	verifying IGMP versions	412
overview	331	verifying PIM mode and interface configuration ..	413
requirements	332	verifying PIM RPF routing table	414
RSVP signaling	332	verifying PIM RPs	413
RSVP-signaled LSPs	335	verifying SAP and SDP configuration	412
signaling protocols overview	322	multicast IPv6 addresses	94
verifying LDP neighbors	338	multicast routing modes	
verifying LDP sessions	339	dense mode	400
verifying LDP-signaled LSPs	340	dense mode, caution for use	400
verifying RSVP neighbors	341	sparse mode	400
verifying RSVP sessions	341	Multicast Source Discovery Protocol	403
verifying RSVP-signaled LSPs	342	multifield classifier	532
verifying traffic forwarding over LDP-signaled		Multilink Frame Relay (MLFR)	100
LSPs	340	Multilink Frame Relay Forum	100
MSDP (Multicast Source Discovery Protocol)	403	Multilink Point-to-Point Protocol (MLPPP)	100
mt-0/0/0 interface	97	multilink services	
MTU (maximum transmission unit)		CRTP	100
E1	107	MLFR	100
E3	110	MLFR FRF.15 and FRF.16	100
Fast Ethernet	115	MLPPP	100

multiple exit discriminator *See* MED
 multiple push label operation 321
 Multiprotocol Label Switching *See* MPLS

N

n-selectors, in IS-IS NET addresses 288
 names, of network interfaces 48
 NAPT (Network Address Port Translation) 441
 NAT (Network Address Translation)
 actions 443
 applying to an interface (configuration editor) ... 485
 configuration editor 480, 482
 description 440
 enabling (Quick Configuration) 479
 match conditions 443
 pools for IPSec tunnels (configuration editor) ... 389
 preparation 476
 Quick Configuration 476
 sample rules 481
 verifying 522
 NCPs (Network Control Protocols) 84
 neighbors *See* adjacencies, IS-IS; BGP peers; DLSw
 peers; OSPF neighbors; RIP neighbors
 NETs (network entity titles)
 n-selectors 288
 on an Ethernet interface 290
 on the loopback interface 290
 parts 223
 system identifier 224
 Network Address Port Translation (NAPT) 441
 Network Address Translation *See* NAT
 Network Control Protocols (NCPs) 84
 network entity titles *See* NETs
 network interfaces
 adding 126
 ATM-over-ADSL configuration 128
 ATM-over-ADSL interfaces 71
 ATM-over-SHDSL configuration 135
 ATM-over-SHDSL interfaces 74
 clocking 78
 deleting 142
 DS3 configuration 119
 E1 configuration 105
 E1 interfaces 56
 E3 configuration 108
 E3 interfaces 60
 enabling PIM on 409
 enabling RIP on 252
 Ethernet interfaces 52
 Fast Ethernet configuration 113
 FCS 80
 G.SHDSL interfaces 74
 IPv4 addressing 90
 IPv6 addressing 93
 ISDN interfaces 74

logical properties 88
 media types 46
 multicast, upstream and downstream 398
 names 48
 naming conventions 47
 output, understanding 49
 physical encapsulation 81
 See also encapsulation types
 physical properties 77
 preparation 103
 protocol families 89
 sample name 49
 serial configuration 122
 serial interfaces 65
 supported 45
 T1 configuration 115
 T1 interfaces 56
 T3 configuration 119
 T3 interfaces 60
 verifying ATM-over-ADSL properties 145
 verifying link states 143
 verifying PIM on 413
 verifying properties 144
 verifying RIP message exchange 261
 verifying RIP on 261
 VLANs 95
 VPN configuration 347
 network layer reachability information (NLRI), BGP, for
 CLNS 375
 network service access points *See* NSAPs
 networks 344
 description 208
 designated router *See* designated router, OSPF
 IPv4 subnets 91
 path cost metrics *See* path cost metrics
 PPPoE session on an ATM-over-ADSL loop 154
 PPPoE session on an Ethernet loop 153
 sample BGP AS path 230
 sample BGP confederation 307
 sample BGP confederations 235
 sample BGP external and internal links 302
 sample BGP local preference use 229
 sample BGP MED use 231
 sample BGP peer network 300
 sample BGP peer session 226
 sample BGP route reflector (one cluster) ... 232, 304
 sample BGP route reflectors (cluster of clusters) .. 234
 sample BGP route reflectors (multiple clusters) .. 233
 sample distance-vector routing 214
 sample DLSw topology 419
 sample LSP topology 319
 sample multiarea OSPF routing 220
 sample OSPF backbone area 221
 sample OSPF multiarea network 272
 sample OSPF network with stubs and NSSAs ... 222

sample OSPF single-area network	271	Open Systems Interconnection (OSI) networks, CLNS	
sample OSPF stub areas and NSSAs	276	VPNs	367
sample OSPF topology	283	operational mode, entering during configuration	34
sample poison reverse routing	216	origin, of BGP route	230
sample RIP network with incoming metric	256	orlonger route list match type	462
sample RIP network with outgoing metric	258	OSI (Open Systems Interconnection) networks, CLNS	
sample RIP topology	253	VPNs	367
sample route advertisement	211	OSPF (Open Shortest Path First)	
sample route aggregation	212	and LDP for VPNs	354
sample routing topology	209	and RSVP for VPNs	355
sample RSVP topology	324	area border routers <i>See</i> area border routers	
sample split horizon routing	215	area type (Quick Configuration)	269
sample static route, preferred path	244	areas	219, 266
sample stub network for static routes	242	<i>See also</i> area border routers; backbone area;	
sample unidirectional routing	217	NSSAs; stub areas	
sample VPN topology	344	authenticating exchanges (OSPFv2 only)	279
static routing	210	backbone area <i>See</i> backbone area	
trusted	440	controlling designated router election	280
untrusted	440	controlling route cost	278
<i>See also</i> VPNs		designated router <i>See</i> designated router, OSPF	
next hop		designating OSPF interfaces (configuration	
address for static routes	241	editor)	272, 274
defining for static routes	243	designating OSPF interfaces (Quick	
qualified, defining for static routes	245	Configuration)	269
qualified, for static routes	238	dial-on-demand routing support, ISDN	190
service set, for IPsec tunnels	384	enabling (Quick Configuration)	268
next term, filter action	498	enabling, description	265
NLRI (network layer reachability information), BGP, for		ensuring efficient operation	277
CLNS	375	injecting OSPF routes into BGP	464
no filter action	497	ISDN dial-on-demand routing support	190
non-UR-2 operating mode	131	LSAs	218
Normal Response Mode, HDLC	88	multiarea network (configuration editor)	272
not-so-stubby areas <i>See</i> NSSAs		NSSAs <i>See</i> NSSAs	
notice icons	xxiii	overview	217, 265
NRM, HDLC	88	path cost metrics <i>See</i> path cost metrics	
NSAPs (network service access points)		Quick Configuration	267
NSAP addresses for IS-IS routers	287	requirements	267
overview	368	route preferences	278
samples	374	router ID (configuration editor)	270
NSSAs (not-so-stubby areas)		router ID (Quick Configuration)	268
area ID (configuration editor)	274	sample multiarea network	272
area ID (Quick Configuration)	268	sample network topology	283
area type (Quick Configuration)	269	sample NSSAs	276
creating (configuration editor)	275	sample single-area network	271
description	221	sample stub areas	276
example	222	single-area network (configuration editor)	270
sample topology	276	stub areas <i>See</i> stub areas	
NT1 devices	75	supported versions	218
numeric range match conditions	446	three-way handshake	218
O		tuning an OSPF network	277
OK button		verifying host reachability	284
J-Web configuration editor	11	verifying neighbors	282
Quick Configuration	8	verifying RIP-enabled interfaces	281
Open Shortest Path First protocol <i>See</i> OSPF		verifying routes	283

OSPF interfaces	
enabling	269
enabling (configuration editor)	272, 274
enabling, for area border routers	275
verifying	281
OSPF neighbors, verifying	282
OSPF page	267
field summary	268
outbound router, in an LSP	319
outgoing metric (RIP)	
description	250
modifying	259
output filters, assigning to interfaces	501
output queues	
assigning forwarding classes	536
sample assignments	536
overriding a configuration file	35
example	36

P

P routers	<i>See</i> provider routers
packet encapsulation	
E1 interfaces	107
E3 interfaces	110
Layer 2 circuits	348
Layer 2 VPNs	348
serial interfaces	124
T1 interfaces	117
T3 interfaces	121
packet filters	<i>See</i> stateful firewall filters; stateless
firewall filters	
packet fragments, matching with a filter	494
packet loss priority, setting with a filter	499
packets	
applying CoS scheduling rules	555
handling packet fragments	502
handling packet fragments (configuration editor)	513
ICMP, matching with a filter	495
length, matching with a filter	497
PADI	86
PADO	87
PADR	87
PADS	87
PADT	87
PPPoE discovery	86, 154
RIP, description	215
TCP, matching with a filter	494
PADI packets	86
PADO packets	87
PADR packets	87
PADS packets	87
PADT packets	87
parentheses, in syntax descriptions	xxiv
partial sequence number PDU (PSNP)	225
passive routes, rejection, in static routing	239
password	
for OSPFv2 authentication	280
for RIPv2 authentication	259
patching a configuration file	35
path cost metrics	
for OSPF routes, description	219, 266
for OSPF routes, modifying	278
for RIP routes, description	249
for RIP routes, modifying	256
path selection	
IS-IS	224
RSVP for MPLS	<i>See</i> traffic engineering database
path-vector protocol	<i>See</i> BGP
pd-0/0/0 interface	98
PDU (protocol data units)	
CSNPs	225
hello PDUs	224
LSPs	224
overview	224
PSNPs	225
PE (provider edge) routers	344
description	327
ES-IS for a CLNS island	371
route distinguishers	357
verifying Layer 2 circuit connections	365
verifying Layer 2 circuit interfaces	365
verifying Layer 2 VPN connections	365
verifying Layer 2 VPN interfaces	365
verifying Layer 3 VPN connections	365
VPN task overview	346
VPN topology	344
<i>See also</i> VPNs	
pe-0/0/0 interface	98
peer routers	<i>See</i> adjacencies, IS-IS; BGP peers; DLSw
peers;	
peering sessions	<i>See</i> BGP peers; BGP sessions
penultimate hop popping (PHP)	321
penultimate router, in an LSP	319
permanent routes, adding	237
permanent virtual circuits (PVCs)	82
PGM (Pragmatic General Multicast)	403
PHP (penultimate hop popping)	321
Physical Interface Module	<i>See</i> PIMs
physical interface properties	
BERT	78
encapsulation	81
FCS	80
interface clocking	78
key properties	77
PIC (PIM on a Services Router)	<i>See</i> PIMs
PIM (Protocol Independent Multicast)	
dense mode	402
disabling on the network management interface	408

- RPF routing table group 410
 - source-specific multicast (SSM) 402
 - sparse mode 402
 - static RP router 408
 - supported versions 405
 - verifying the mode 413
 - verifying the RP 413
- pimd interface 98
- pime interface 98
- PIMs (Physical Interface Modules)
 - G.SHDSL 135
 - See also* G.SHDSL PIMs
 - output about, understanding 49
 - PIM number, always 0 48
 - PIM slot number 48
- ping
 - verifying link states 143
 - VPN connection 364
- ping command 527
 - explanation 527
- Ping Host page, output for BGP 311
- ping mpls l2circuit interface command 365
- ping mpls l2circuit virtual-circuit command 365
- ping mpls l2vpn instance 365
- ping mpls l2vpn interface command 365
- ping mpls l3vpn command 365
- ping trusted-nw-trusted-host 522
 - explanation 523
- ping untrusted-nw-untrusted-host command 522
 - explanation 523
- plesiochronous networks 79
- Point-to-Point Protocol *See* PPP
- Point-to-Point Protocol over ATM *See* PPPoA
- Point-to-Point Protocol over Ethernet *See* PPPoE
- poison reverse technique 215
- polarity, signal 67
- policers
 - description 453
 - for firewall filter 531
 - for stateless firewall filters 508
- policy *See* routing policies
- pop label operation 320
- ports
 - DS1 *See* T1 ports
 - DS3 *See* T3 ports
 - E1 *See* E1 ports
 - E3 *See* E3 ports
 - interfaces overview 41
 - See also* ATM-over-ADSL interfaces;
 - ATM-over-SHDSL interfaces; ISDN
 - interfaces; loopback interface;
 - management interfaces; network
 - interfaces; services interfaces; special
 - interfaces
 - number in interface name 49
 - T1 *See* T1 ports
 - T3 *See* T3 ports
- pp0
 - creating 158
 - enabling CHAP 161
 - information about 164
 - interface description 98
 - logical Ethernet interface on 159
- PPP encapsulation
 - CHAP authentication 84
 - CSU/DSU devices 85
 - LCP connection process 83
 - magic numbers 85
 - NCPs 84
 - overview 83
- PPP over ATM *See* PPPoA
- PPP over ATM AAL5 multiplex encapsulation ... 132, 139
- 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 133
 - logical encapsulation 132
 - logical encapsulation (ATM-over-ADSL) 132
 - logical encapsulation (ATM-over-SHDSL) 139
 - physical encapsulation (ATM-over-ADSL) 131
 - physical encapsulation (ATM-over-SHDSL) 137
 - verifying ATM-over-ADSL configuration 149
- PPPoE (Point-to-Point Protocol over Ethernet) 153
 - address assignment 160
 - CHAP 155, 161
 - client and server 152
 - creating the pp0 interface 158
 - discovery packets 86, 154
 - encapsulation on an Ethernet interface 86, 156
 - interfaces 153
 - MTU values 160
 - overview 152
 - preparation 155
 - sample topology 153
 - service type 159
 - session limit 87, 154
 - session overview 87, 154
 - session reconnection time 159
 - verifying interfaces 164
 - verifying sessions 165
 - verifying statistics 166
 - verifying version information 166
 - See also* PPPoE over ATM-over-ADSL
- PPPoE Active Discovery Initiation (PADI) packets 86
- PPPoE Active Discovery Offer (PADO) packets 87
- PPPoE Active Discovery Request (PADR) packets 87
- PPPoE Active Discovery Session-Confirmation (PADS)
 - packets 87
- PPPoE Active Discovery Termination (PADT) packets ... 87

- PPPoE encapsulation *See* PPPoE
 - PPPoE over ATM LLC encapsulation 132, 139
 - PPPoE over ATM-over-ADSL 153
 - CHAP 161
 - creating the pp0 interface 158
 - encapsulation 157
 - preparation 155
 - sample topology 153
 - verifying configuration 162–163
 - See also* PPPoE
 - PPPoEoA *See* PPPoE over ATM-over-ADSL
 - Pragmatic General Multicast 403
 - precedence
 - matching with a filter 496
 - ToS value for DSLw 426
 - preferences
 - for OSPF routes 278
 - for static routes 238
 - setting for static routes 245
 - prefix-length-range match type 463
 - primary stations, HDLC 87
 - priority of a packet, setting with a filter 499
 - propagation, suppressing 470
 - properties, verifying
 - for ATM-over-ADSL network interfaces 145
 - for network interfaces 144
 - protocol data units *See* PDUs
 - protocol families
 - ccc 90
 - common protocol suites 89
 - inet 89
 - inet6 89
 - ISO 89
 - mlfr-end-to-end 90
 - mlfr-uni-nni 90
 - mlppp 90
 - MPLS 89
 - overview 89
 - tcc 90
 - tnp 90
 - vpls 90
 - Protocol Independent Multicast *See* PIM
 - protocols
 - Auto-RP 402
 - BGP *See* BGP
 - CRTP 140
 - distance vector *See* RIP
 - DVMRP 401
 - EGPs 208
 - EIA-530 69
 - IGMP *See* IGMP
 - IGPs 208
 - IPSec *See* IPSec
 - IPv4, matching with a filter 494
 - IPv6, matching with a filter 495
 - IS-IS *See* IS-IS
 - LDP *See* LDP
 - MPLS *See* MPLS
 - MSDP 403
 - multicast *See* multicast
 - NAT *See* NAT
 - OSPF *See* OSPF
 - overview 203
 - path vector *See* BGP
 - PGM 403
 - PIM dense mode *See* PIM
 - PIM source-specific multicast (SSM) 402
 - PIM sparse mode *See* PIM
 - PPPoE *See* PPPoE
 - RIP *See* RIP
 - RS-232 69
 - RS-422/449 70
 - RSVP *See* RSVP
 - SAP and SDP *See* SAP; SDP
 - serial 68
 - SSP for DSLw 419
 - V.35 70
 - X.21 71
 - provider edge routers *See* PE routers
 - provider routers 344
 - description 327
 - VPN task overview 346
 - VPN topology 344
 - See also* VPNs
 - PSNP (partial sequence number PDU) 225
 - push label operation 320
 - PVCs (permanent virtual circuits) 82
- ## Q
- queuing rules, CoS 555
 - Quick Configuration
 - BGP page 297
 - buttons 8
 - E1 Interfaces page 106
 - E3 Interfaces page 109
 - Fast Ethernet Interfaces page 114
 - Firewall Filters initial page 487
 - Firewall Filters interface assignment page 500
 - Firewall Filters match conditions and actions
 - page 488
 - Firewall/NAT application page 478
 - Firewall/NAT main page 477
 - Interfaces page 104
 - IPSec Tunnels page 381
 - network interfaces 104
 - OSPF page 267
 - overview 7
 - RIP page 251
 - serial Interfaces page 123
 - Static Routes page 240

Summary page	7
T1 Interfaces page	116
T3 (DS3) Interfaces page	120

R

radio buttons

Delete Configuration Below This Point	11
Discard All Changes	11
Discard Changes Below This Point	11
RADIUS authentication, of PPP sessions	155
random early detection <i>See</i> RED drop profiles	
reachability	
verifying for BGP peers	311
verifying for DSLw peers	429
reactivate command	29
RED (random early detection) drop profiles	547
samples	546
redistributing routes	465
Refresh button	11
reject, filter action	498
rejecting invalid routes	464
relative option	35
release notes, URL	xxi
Remote Authentication Dial-In User Service (RADIUS)	
authentication, of PPP sessions	155
remote router, DSLw	423
<i>See also</i> DSLw peers	
remote tunnel endpoint, IPSec	382
rename command	27
renaming configuration identifiers	27
repeaters, on LAN segments	54
replacing a configuration file	35
example	36
request system configuration rescue delete	
command	33
request system configuration rescue save command	33
rescue configuration	
deleting (CLI configuration editor)	33
deleting (J-Web)	21–22
disabling CONFIG button for	33
loading with the CONFIG button	21, 33
setting (CLI configuration editor)	33
setting (J-Web)	21
viewing (CLI configuration editor)	33
viewing (J-Web)	21–22
reservation <i>See</i> RSVP	
reset button, for return to factory configuration	
<i>See</i> CONFIG button	
Resource Reservation Protocol <i>See</i> RSVP	
reverse-path forwarding <i>See</i> RPF	
rewrite rules	
description	453
replacing DSCPs	538
sample rules	537
when applied	456

RIB *See* routing table

RIP (Routing Information Protocol)

authentication (RIPv2 only)	250
authentication (RIPv2 only), configuring	259
basic network (configuration editor)	253
designating RIP interfaces	252
distance vector protocol	213
efficiency techniques	215
enabling (Quick Configuration)	252
maximum hop count	214
overview	213, 249
packets	215
path cost metrics <i>See</i> path cost metrics	
poison reverse technique	215
Quick Configuration	250
requirements	250
routing policy (configuration editor)	253
sample network with incoming metric	256
sample network with outgoing metric	258
sample topology	253
split horizon technique	215
supported versions	213
traffic control with metrics <i>See</i> path cost metrics	
traffic control with metrics, configuring	256
unidirectional limitations	216
verifying host reachability	262
verifying RIP message exchange	261
verifying RIP-enabled interfaces	261
RIP neighbors, verifying	261
RIP page	251
field summary	252
rollback ? command	33
rollback command	32
rollback rescue command	32
rolling back a configuration file	
during configuration (CLI configuration editor)	32
during configuration (J-Web)	21
route advertisements	
AS path in	229
BGP, update messages	227
description	211
external, EBGP	227
internal, IBGP	227
LSAs	218
stub areas and NSSAs, to control	221
route aggregation	211
route distinguishers	
description	328
formats for	357
route injection	464
route list match types	462
route manipulation actions, routing policies	439
route redistribution	464
route reflectors <i>See</i> BGP route reflectors	

- route selection
 - BGP process 228
 - BGP, determining by AS path 229
 - BGP, determining by local preference 228
 - BGP, determining by MED metric 230
 - BGP, lowest origin value preferred 230
 - static routes, defining 243
- route targets, VPN 329
 - in a routing instance 357
- route-flap damping 470
 - parameters 470
- router *See* Services Router
- routing 203
 - advertisements 211
 - aggregation 211
 - BGP *See* BGP
 - configuring PPPoE 151
 - configuring VPNs 343
 - DLSw *See* DLSw
 - dynamic 210
 - filtering and classifying routes 433
 - filtering routes with policies 459
 - filtering traffic through a firewall 475
 - forwarding tables 209
 - from one source to many destinations 405
 - IBM networking *See* DLSw
 - in multiple ASs with BGP 295
 - in one AS with OSPF 265
 - in one AS with RIP 249
 - IS-IS *See* IS-IS
 - MPLS for VPNs 315
 - MPLS traffic engineering 331
 - multicast *See* multicast
 - neighbors *See* BGP peers; OSPF neighbors; RIP neighbors
 - OSPF *See* OSPF
 - overriding default packet forwarding with CoS .. 529
 - protecting local IP addresses with NAT 475
 - protocol overview 203
 - RIP *See* RIP
 - RIP statistics 261
 - routing tables 209
 - static *See* static routing
 - through IPSec tunnels 379
 - VPNs 343
 - See also* protocols; routing policies; routing solutions
- Routing Engine
 - handling packet fragments for (configuration editor) 511
 - protecting against DoS attacks (configuration editor) 506
 - protecting against untrusted services and protocols (configuration editor) 502
- routing information base *See* routing table
- Routing Information Protocol *See* RIP
- routing instance
 - filter action, setting 498
 - for CLNS static routes (with IS-IS) 370
 - for CLNS static routes (without IS-IS) 374
 - VPN configuration 357
 - VPN route target 357
 - VRF instances 328
 - VRF table 357
- routing policies
 - actions 438
 - applying 440
 - BGP export, for CLNS 372
 - BGP routing policy (configuration editor) 303
 - components 435
 - configuration tasks 460
 - default actions 440
 - export statement 440
 - final actions 440
 - forwarding class with source and destination ... 466
 - grouping source and destination prefixes 466
 - import statement 440
 - injecting routes from one protocol into another .. 464
 - Layer 2 VPN export policy 361
 - Layer 2 VPN import policy 360
 - Layer 3 VPNs 363
 - making BGP routes less preferable 467
 - match conditions 436
 - overview 435
 - policy name 461
 - preparation 460
 - prepending AS paths 467
 - reducing update messages with flap damping ... 470
 - rejecting invalid routes 462
 - RIP routing policy (configuration editor) 253
 - route redistribution 464
 - route-flap damping 470
 - terms 436
 - terms, creating 461
 - VPN configuration 359
- routing protocols *See* protocols
- routing solutions
 - BGP confederations, for scaling problems 306
 - BGP route reflectors, for scaling problems 303
 - BGP scaling techniques 231
 - controlling designated router election 280
 - controlling OSPF route cost 278
 - controlling OSPF route selection 278
 - controlling RIP traffic with the incoming metric .. 256
 - controlling RIP traffic with the outgoing metric .. 257
 - CoS with DiffServ 449, 529
 - designated router, to reduce flooding 218
 - directing BGP traffic by local preference 228
 - filtering unwanted services and protocols 502
 - firewall filters and NAT 440, 475

handling packet fragments	502
handling packet fragments (configuration editor)	511
making BGP routes less preferable	467
MPLS traffic engineering	331
multicast administrative scoping	401
multicast reverse-path forwarding (RPF)	400
multicast shortest-path tree (SPT)	401
NSSAs, to control route advertisement	221
path cost metrics, for packet flow control <i>See</i> path cost metrics	
point-to-point sessions over Ethernet	151
poison reverse, for traffic reduction	215
preventing multicast routing loops	400
protecting against DoS attacks	506
reducing update messages with flap damping	470
rejecting invalid routes	462
routing policies	435, 459
securing OSPF routing (OSPFv2 only)	279
split horizon, for traffic reduction	215
static route control techniques	238
stub areas, to control route advertisement	221
VPNs	343
routing table	
controlling static routes in	238, 245
description	209
displaying static routes in	247
RPF group, for multicast	410
sample distance-vector routing	214
updates, limitations in RIP	216
verifying for RPF	414
verifying LDP-signaled LSPs	340
verifying OSPF routes	283
verifying RSVP-signaled LSPs	342
RP (rendezvous point)	
static	408
verifying	413
RPF (reverse-path forwarding)	
description	400
routing table group	410
verifying the routing table	414
RS-232	69
RS-422/449	70
RS-530	69
RSVP (Resource Reservation Protocol)	
and OSPF for VPNs	355
bandwidth reservation	323
CSPF	325
disabling CSPF	338
EROS	323
fundamentals	323
link coloring	325
overview	332
requirements	332
RSVP-signaled LSPs	335
verifying LSPs	342
verifying neighbors	341
verifying sessions	341
verifying the routing table on the entry router	342
RSVP neighbors, verifying	341
RSVP-signaled LSP <i>See</i> RSVP	
run command	34
S	
S,G notation, for multicast forwarding states	399
S/T interfaces	
overview	75
PIMs	171
samples	149, 162–163
CLNS VPN configuration	376
firewall filter configurations	518
PPPoA for ATM-over-ADSL configuration	149
PPPoE over ATM-over-ADSL configuration	163
PPPoE over Ethernet configuration	162
sample DSLw topology	419
<i>See also</i> networks; topology	
sampling traffic on an interface, with a filter	499
SAP (Session Announcement Protocol)	
description	403
session announcements	406
verifying	412
saving, configuration files	37
scaling BGP <i>See</i> BGP confederations; BGP route reflectors	
schedulers	
assigning resources	549
default settings	454
description	453
mapping to forwarding classes	553
sample mappings	553
sample schedulers	549
scheduling a commit	31
scope, IPv6 addresses	
global unicast	94
link-local unicast	94
multicast types	94
site-local unicast	94
scoping, administrative	401
SDP (Session Discovery Protocol)	
description	403
session announcements	406
verifying	412
secondary stations, HDLC	88
secret, CHAP <i>See</i> CHAP, local identity	
security	
IPSec tunnels	379
MD5 authentication for OSPF	280
MD5 authentication for RIPv2	260
password authentication for OSPFv2	280
password authentication for RIPv2	259

- security association *See* IPSec security associations
- serial interfaces 65
 - clocking modes 67
 - connection process 66
 - DTE default clock rate reduction 68
 - EIA-530 69
 - inverting the transmit clock 68
 - line protocols 68
 - RS-232 69
 - RS-422/449 70
 - signal polarity 67
 - transmission signals 66
 - V.35 70
 - X.21 71
 - See also* serial ports
- serial numbers, in MAC addresses 51
- serial ports 65
 - CHAP 124
 - clock rate 125
 - clocking 125
 - clocking, inverting the transmit clock 125
 - configuring 122
 - encapsulation type 124
 - line speed 125
 - logical interfaces 124
 - See also* serial interfaces
- service classes, corresponding DSCPs 450
- service sets
 - for IPSec tunnels 384
 - for NAT rules 485
 - for stateful firewall filters 485
- service types, naming for PPPoE 159
- services interfaces
 - applying a NAT rule to (configuration editor) 485
 - applying a stateful firewall filter to (configuration editor) 485
 - CRTCP 100
 - for IPSec tunnels 383
 - MLFR 100
 - MLFR FRF.15 and FRF.16 100
 - MLPPP 100
 - overview 100
- Services Router
 - as a PPPoE client 152
 - BGP routing 295
 - CLNS VPNs 367
 - configuration tools 3
 - CoS overview 449
 - CoS with DiffServ 529
 - CPE, with PPPoE 151
 - See also* PPPoE
 - DLSw 417
 - firewall filter overview 440
 - firewall filters 475
 - IBM networking 417
 - interfaces overview 41
 - IPSec tunnels 379
 - IS-IS protocol 287
 - ISDN connections 169
 - MPLS for VPNs overview 315
 - MPLS traffic engineering 331
 - multicast 405
 - multicast overview 395
 - NAT 475
 - network interfaces 103
 - OSPF routing 265
 - PPPoE 151
 - RIP routing 249
 - routing policies 459
 - routing policy overview 435
 - routing protocols overview 203
 - static routing 237
 - VPNs 343
- Session Announcement Protocol *See* SAP; SDP
- sessions
 - announcements, multicast 406
 - BGP session establishment 226
 - BGP session maintenance 227
 - ISDN session establishment 76
 - LDP, verifying 339
 - limit on PPPoE sessions 154
 - PPPoE 87, 154
 - PPPoE, reconnection time 159
 - RSVP, verifying 341
- SHDSL ports *See* ATM-over-SHDSL interfaces
- shortest path first algorithm 217
- shortest-path tree 401
- show access command 149
- show bgp group command 309
 - explanation 310
- show bgp neighbor command 308
 - explanation 309
- show bgp summary command 310
 - explanation 311
- show chassis hardware command 49
- show class-of-service adaptive-shaper command 561
- show class-of-service interface command 561
- show class-of-service virtual-channel command 562
- show class-of-service virtual-channel-group
 - command 562
- show cli history command 34
- show command 25, 376
- show dlsw capabilities command 427
- show dlsw circuits command 427
- show dlsw peers command 428
- show dlsw peers detail command 428
 - explanation 429
- show dlsw reachability command 429
- show firewall command 518
- show firewall filter protect-RE command 524

show firewall log command	523	show sap listen command	412, 561
explanation	524	explanation	412, 561
show igmp interface command	412	show services command	518
explanation	413	show services ipsec-vpn ipsec statistics command ...	391
show interfaces bc-0/0/4 extensive command	194	explanation	391
show interfaces br-6/0/0 extensive command	193	show statement dlsw circuits detail command	428
show interfaces command	149, 162–163	show system reboot command	34
show interfaces dc-0/0/4 extensive command	196	signaling protocols	331
show interfaces detail command	144	overview	322
show interfaces dl0 extensive command	197	VPNs	353
show interfaces extensive command	145	<i>See also</i> LDP; MPLS traffic engineering; RSVP	
explanation, for ATM-over-ADSL interfaces	148	signals	
explanation, for ISDN interfaces	194, 196–197	DS1	57
show interfaces fe-3/0/0 command	427	E1 loopback (control)	60
show interfaces lo0 command	517	explicit clocking signal transmission	79
show interfaces ppo command	164	multiplexing DS1 into DS2 signal	60
show isdn status command	193	serial polarity	67
show isis adjacency brief command	292	serial transmission	66
show isis adjacency extensive command	293	T1 loopback (control)	60
explanation	294	V.35	70
show isis interface brief command	291	X.21	71
show isis interface detail command	291	single-area network, OSPF	270
explanation	292	site-local unicast IPv6 addresses	94
show ldp neighbor command	338	SNA forwarding <i>See</i> DLSw	
explanation	339	source-specific multicast	402
show ldp session detail command	339	sp-0/0/0	
explanation	340	for IPSec tunnels (configuration editor)	383
show multicast rpf command	414	interface description	98
explanation	414	no stateful firewall filters	442
show ospf interface command	281	sparse mode <i>See</i> multicast routing modes	
explanation	282	special interfaces	
show ospf neighbor command	282	CRTP	100
explanation	282	dsc interface	98
show ospf route command	283	IPv4 addressing	90
results	284	IPv6 addressing	93
show pim interface command	413	logical properties	88
explanation	413	loopback interface	99
show pim rps command	413	management interface	99
explanation	414	MLFR	100
show ppoe interfaces command	165	MLFR FRF.15 and FRF.16	100
show ppoe statistics command	166	MLPPP	100
show ppoe version command	166	names	48
show rip neighbor command	261	naming conventions	47
explanation	261	output, understanding	49
show rip statistics command	261	overview	96
show route summary command	525, 527	physical properties	77
explanation	526, 528	protocol families	89
show route table inet.3 command	340, 342	services interfaces	100
explanation	340, 342	summary	96
show route terse command	247	SPF (shortest path first) algorithm	217
explanation	248	split horizon technique	215
show rsvp neighbor command	341	SPT (shortest-path tree)	401
explanation	341	ssh command	525
show rsvp session detail command	342	explanation	526
explanation	342	SSP (Switch-to-Switch Protocol) for DLSw	419

- stateful firewall filters
 - actions 443
 - applying to an interface (configuration editor) ... 485
 - automatic discard rule..... 441
 - configuration editor 480, 482
 - configuration overview 442
 - description 441
 - do not apply to sp-0/0/0..... 442
 - enabling (Quick Configuration) 479
 - IPSec tunnels, rules for (configuration editor) 387
 - junos-algs-outbound default group..... 442
 - match conditions 443
 - preparation..... 476
 - Quick Configuration 476
 - sample rules 481
 - untrusted network..... 442
 - verifying..... 522
 - verifying actions 525
- stateless firewall filters
 - action modifiers (Quick Configuration) 498
 - Action tab (Quick Configuration) 497
 - actions and action modifiers..... 448
 - adding (Quick Configuration) 490
 - applying to an interface (configuration editor) ... 516
 - assigning to interfaces (Quick Configuration) 499
 - automatic discard rule..... 441, 444
 - bit-field logical operators 448
 - description 441
 - destination matching (Quick Configuration) 491
 - filter actions (Quick Configuration) 497
 - See also* actions
 - handling packet fragments..... 502
 - handling packet fragments (configuration editor)..... 511
 - input filters, interface assignment (Quick Configuration) 501
 - interface matching (Quick Configuration) 493
 - IPv4 filters (Quick Configuration) 487
 - IPv6 filters (Quick Configuration) 487
 - match conditions 445
 - Match Destination tab (Quick Configuration) 491
 - Match Interface tab (Quick Configuration) 493
 - Match Match Network tab (Quick Configuration) 494
 - Match Match Packet and Network tab (Quick Configuration) 494
 - Match Source or Destination tab (Quick Configuration) 492
 - Match Source tab (Quick Configuration)..... 491
 - network matching (Quick Configuration)..... 494
 - output filters, interface assignment (Quick Configuration) 501
 - packet matching (Quick Configuration) 494
 - planning..... 444, 502
 - policers for 508
 - preparation..... 476
 - protecting the Routing Engine against ICMP floods (configuration editor) 506
 - protecting the Routing Engine against TCP floods (configuration editor) 506
 - protecting the Routing Engine against untrusted protocols (configuration editor)..... 502
 - protecting the Routing Engine against untrusted services (configuration editor)..... 502
 - Quick Configuration 486
 - sample terms, to filter fragments..... 512
 - sample terms, to filter services and protocols.... 503
 - sample terms, to protect against DoS attacks 507
 - source matching (Quick Configuration)..... 491
 - source or destination matching (Quick Configuration) 492
 - summary (Quick Configuration) 489
 - terms, adding (Quick Configuration)..... 490
 - typical, planning..... 502
- statements
 - adding or modifying..... 26
 - copying..... 27
 - deactivating..... 29
 - deleting..... 26
 - replacing 35
- static LSPs 321
- static routes
 - CLNS VPNs (with IS-IS) 370
 - CLNS VPNs (without IS-IS)..... 374
 - configuring basic routes (configuration editor) ... 242
 - controlling..... 238
 - controlling in routing and forwarding tables 245
 - default properties..... 239
 - default properties, setting 246
 - defining route selection 243
 - preferences 238
 - preventing readvertisement..... 239
 - qualified next hops 238
 - Quick Configuration 240
 - rejecting passive traffic..... 239
 - requirements..... 240
 - route retention 239
 - sample preferred path..... 244
 - sample stub network 242
 - verifying..... 247
- Static Routes page 240
 - field summary 241
- static routing
 - default gateway 241
 - description 210
 - overview 237
 - See also* static routes
- static RP router 408
 - See also* RP

statistics	
ATM-over-ADSL interfaces	149
firewall filters	524
IPSec tunnels	391
PPPoE	166
RIP	261
status command	22
status, link states, verifying	143
strict hops, RSVP	324
stub areas	
area ID (configuration editor)	274
area ID (Quick Configuration)	268
area type (Quick Configuration)	269
controlling OSPF route cost	279
creating (configuration editor)	275
description	221
example	222
sample topology	276
sub-ASs, BGP	234
subautonomous systems, BGP	234
subnet masks	92
subnets <i>See</i> subnetworks	
subnetworks	
description	208
IPv4 subnets	91
multicast leaves and branches	398
route aggregation	212
Summary Quick Configuration page	7
superframe framing	59
support, technical <i>See</i> technical support	
SVCs (switched virtual circuits)	82
swap and push label operation	321
swap label operation	320
switch types, ISDN	174
Switch-to-Switch Protocol (SSP) for DLSw	419
switched virtual circuits (SVCs)	82
switches, on LAN segments	54
symmetric high-speed digital subscriber line (SHDSL)	
<i>See</i> ATM-over-SHDLS interfaces	
synchronous networks	78
syntax conventions	xxiii
system clock <i>See</i> clocking	
system identifier, IS-IS	
all zeros not supported	288
formats, MAC or IP address	288
identifier-to-hostname mapping	288
overview	224
system log, of packet information	498
Systems Network Architecture (SNA) forwarding	
<i>See</i> DLSw	
T	
T1 interfaces	56
AMI encoding	58
B8ZS encoding	58
D4 framing	59
data stream	57
encoding	58
ESF framing	59
framing	59
loopback	60
overview	57
signals	57
superframe framing	59
<i>See also</i> T1 ports	
T1 ports	56
adding CRTTP	140
cable length	119
CHAP	117
clocking	117
configuring	115
data inversion	118
encapsulation type	117
fractional, channel number	49
frame checksum	119
framing	118
logical interfaces	117
MTU	117
overview	57
time slots	118
<i>See also</i> T1 interfaces	
T3 interfaces	60
bit stuffing	61
data stream	60
DS3 framing	61
multiplexing on	61
overview	60
<i>See also</i> T3 ports	
T3 ports	60
C-bit parity	122
cable length	122
CHAP	121
clocking	121
configuring	119
encapsulation type	121
frame checksum	122
framing	122
logical interfaces	121
MTU	121, 124
overview	60
<i>See also</i> T3 interfaces	
tap interface	98
tcc protocol family	90
TCP packets, matching with a filter	494
TCP policers	508
technical support	
contacting JTAC	xxvi
TED <i>See</i> traffic engineering database	
telnet command	526
explanation	527

- terminology
 - CLNS 367
 - configuration 3
 - CoS 433
 - DLSw 417
 - firewall filters 433
 - interfaces 42
 - ISDN 169
 - MPLS 315
 - multicast 395
 - ports 42
 - PPPoE 151
 - routing policies 433
 - routing protocols 203
 - VPNs 315
- terms
 - firewall filter, for multifold classifier 532
 - in a routing policy 436
 - in a routing policy, creating 461
 - stateless firewall filters, adding (Quick Configuration) 490
- three-way handshake 218
- through route list match type 463
- time slots
 - E1 108
 - number in interface name 49
 - T1 118
- time-to-live (TTL) value, matching with an IPv4 filter 496
- tnp protocol family 90
- to statement, routing policy match conditions 436
- top command 25
- topology
 - data link layer 50
 - IPv4 subnets 92
 - PPPoE session on an ATM-over-ADSL loop 154
 - PPPoE session on an Ethernet loop 153
 - sample ATM-over-ADSL 73
 - sample BGP AS path 230
 - sample BGP confederation 307
 - sample BGP confederations 235
 - sample BGP external and internal links 302
 - sample BGP local preference use 229
 - sample BGP MED use 231
 - sample BGP peer network 300
 - sample BGP peer session 226
 - sample BGP route reflector (one cluster) ... 232, 304
 - sample BGP route reflectors (cluster of clusters) .. 234
 - sample BGP route reflectors (multiple clusters) .. 233
 - sample distance-vector routing 214
 - sample DLSw topology 419
 - sample Frame Relay network 81
 - sample ISDN network 75
 - sample LAN 95
 - sample LSP network 319
 - sample multiarea OSPF routing 220
 - sample OSPF backbone area 221
 - sample OSPF multiarea network 272
 - sample OSPF network 283
 - sample OSPF network with stubs and NSSAs 222
 - sample OSPF single-area network 271
 - sample OSPF stub areas and NSSAs 276
 - sample poison reverse routing 216
 - sample RIP network 253
 - sample RIP network with incoming metric 256
 - sample RIP network with outgoing metric 258
 - sample route advertisement 211
 - sample route aggregation 212
 - sample router network 209
 - sample RSVP-signaled LSP 324
 - sample split horizon routing 215
 - sample static route 210
 - sample static route, preferred path 244
 - sample stub network for static routes 242
 - sample unidirectional routing 217
 - sample VLAN 96
 - sample VPN 344
- topology database, OSPF 265
- ToS (type of service), precedence for DLSw packets .. 424
- Traceroute page
 - results for OSPF 285
 - results for RIP 263
- traceroute source bypass-routing gateway
 - command 340
 - explanation 341
- traffic
 - controlling with incoming RIP metric 256
 - controlling with outgoing RIP metric 257
 - outgoing, securing 380
 - sampling on an interface, with a filter 499
- traffic engineering *See* MPLS traffic engineering
- traffic engineering database
 - CSPF constraints on path selection 325
 - CSPF rules for path selection 325
 - link coloring for CSPF path selection 325
- transit interfaces
 - LDP-signaled LSPs for 333
 - RSVP-signaled LSPs for 335
- transit routers, in an LSP 319
- transmit clock source *See* clocking
- trusted networks, firewall filter protection 440
- TTL (time-to-live) value, matching with an IPv4 filter 496
- tunnels, through a public network *See* IPSec tunnels; VPNs
- two-dimensional parity 80
- two-wire mode (2 ports), SHDSL *See* ATM-over-SHDSL interfaces
- type of service (ToS), precedence for DLSw packets .. 424
- types of interfaces 48

U

U interface	
overview	76
PIMs	171
unicast IPv6 addresses	94
untrusted networks, firewall filter actions on	440
up command	24
uploading a configuration file	14
upstream interfaces	398
<i>See also</i> multicast	
upto route list match type	463
UR-2 operating mode	131
URLs	
release notes	xxi

V

V.35	70
variable-length subnet masks (VLSMs)	92
VCI (virtual channel identifier)	
ATM-over-ADSL interfaces	133
ATM-over-SHDSL interfaces	140
PPPoE over ATM-over-ADSL interfaces	158
verification	
adaptive shaping	561
ATM-over-ADSL interface properties	145
B-channels	194
BGP configuration	310
BGP groups	309
BGP peer reachability	311
BGP peers (neighbors)	308
CLNS VPNs	376
configuration syntax	30
D-channel	196
dialer interfaces	197
DLSw capabilities	427
DLSw circuit state	427
DLSw circuit state (detail)	428
DLSw LLC2 properties	427
DLSw peers	428
DLSw peers (detail)	428
DLSw reachability	429
firewall filter actions	525
firewall filter flood protection	526
firewall filter handles fragments	527
firewall filter operation	523
firewall filters	517
firewall statistics	524
IGMP version	412
IPSec tunnel operation	391
IS-IS adjacencies	292
IS-IS adjacencies (detail)	293
IS-IS interface configuration	291
IS-IS interface configuration (detail)	291
IS-IS neighbors	292
IS-IS neighbors (detail)	293

ISDN interfaces	193
ISDN status	193
LDP neighbors	338
LDP sessions	339
LDP-signaled LSP	340
MPLS traffic engineering	338
multicast SAP and SDP	412
multicast session announcements	561
network interfaces	143
OSPF host reachability	284
OSPF neighbors	282
OSPF routes	283
OSPF-enabled interfaces	281
PIM mode and interface configuration	413
PIM RP address	413
PIM RPF routing table	414
PPPoA for ATM-over-ADSL configuration	149
PPPoE interfaces	164
PPPoE over ATM-over-ADSL configuration	162–163
PPPoE sessions	165
PPPoE statistics	166
PPPoE version	166
RIP host reachability	262
RIP message exchange	261
RIP-enabled interfaces	261
RSVP neighbors	341
RSVP sessions	341
RSVP-signaled LSP	342
stateful firewall filters	522
static routes in the routing table	247
traffic forwarding over LDP-signaled LSPs	340
virtual channel	562
virtual channel group	562
VPNs	364
version	
OSPF, supported	218
PPPoE, verifying	166
RIP, supported	213
View Configuration Text page	13
virtual channel identifier <i>See</i> VCI	
virtual channels	
applying CoS rules to logical interfaces	555
filter action modifier, setting	498
virtual circuit ID, for Layer 2 circuits	356
virtual circuits	
DLCIs	82
overview	82
PVCs	82
SVCs	82
virtual LANs <i>See</i> VLANs	
virtual link, through the backbone area	220
virtual path identifier (VPI), PPPoE over ATM-over-ADSL	
interfaces	158
virtual private networks <i>See</i> VPNs	

- VLANs (virtual LANs)
 - LAN comparison 95
 - overview 95
 - topology 96
 - VLSMs (variable-length subnet masks) 92
 - VPI, PPPoE over ATM-over-ADSL interfaces 158
 - vpls protocol family 90
 - VPN routing and forwarding (VRF) instances 328
 - VPN routing and forwarding table *See* VRF table
 - VPNs (virtual private networks) 343
 - AS number 352
 - basic Layer 2 circuit description 345
 - basic Layer 2 VPN description 344
 - basic Layer 3 VPN description 345
 - BGP 351
 - CLNS *See* CLNS
 - components 326
 - configuration overview 343
 - configuration task overview 346
 - IGPs 353
 - Layer 2 circuit configuration 356
 - LSP for RSVP 350
 - MPLS 349
 - overview 315, 326
 - participating interfaces 347
 - preparation 346
 - protocols for 349
 - route distinguishers 328, 357
 - route target 357
 - route targets 329
 - routing information 328
 - routing instance 357
 - routing instance, for CLNS static routes (with IS-IS) 370
 - routing instance, for CLNS static routes (without IS-IS) 374
 - routing policies 359
 - routing requirements 327
 - sample topology 344
 - signaling protocols 353
 - tunneling process 327
 - types 329
 - verifying connectivity 364
 - VRF instances 328
 - VRF table *See* VRF table
 - See also* Layer 2 circuits; Layer 2 VPNs: Layer 3 VPNs; MPLS
 - VRF (VPN routing and forwarding) table 357
 - route targets 329
 - VRF instances 328
 - VRF instances 328
- X**
- X.21 71