



J-series™ Services Router

Configuration Guide

Release 7.3

Juniper Networks, Inc.

1194 North Mathilda Avenue
Sunnyvale, California 94089
USA

408-745-2000

www.juniper.net

Part Number: 530-013469-01, Revision 2

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 7.3
Copyright © 2005, Juniper Networks, Inc.
All rights reserved. Printed in USA.

Writing: 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
7 July 2005—Revision 2.

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. Other Juniper documentation for the Software (such as product purchase documents, documents accompanying the product, the Software user manual(s), Juniper's website for the Software, or messages displayed by the Software) may specify limits to Customer's use of the Software. Such limits may restrict use to a maximum number of seats, concurrent users, sessions, subscribers, nodes, or transactions, or require the purchase of separate licenses to use particular features, functionalities, or capabilities, or provide 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, 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; (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, or capability without first purchasing the applicable license(s) and obtaining a valid key from Juniper, even if such feature, function, 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.

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.** If the Software is distributed on physical media (such as CD), Juniper warrants for 90 days from delivery that the media on which the Software is delivered will be free of defects in material and workmanship under normal use. This limited warranty extends only to the Customer. Except as may be expressly provided in separate documentation from Juniper, no other warranties apply to the Software, and the Software is otherwise provided AS IS. Customer assumes all risks arising from use of the Software. Customer's sole remedy and Juniper's entire liability under this limited warranty is that Juniper, at its option, will repair or replace the media containing the Software, or provide a refund, provided that Customer makes a proper warranty claim to Juniper, in writing, within the warranty period. Nothing in this Agreement shall give rise to any obligation to support the Software. Any such support shall be governed by a separate, written agreement. To the maximum extent permitted by law, Juniper shall not be liable for any liability for 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 HEREIN OR IN SEPARATE DOCUMENTATION PROVIDED FROM JUNIPER AND 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.

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 you may contain encryption or other capabilities restricting your 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. **Miscellaneous.** This Agreement shall be governed by the laws of the State of California without reference to its conflicts of laws principles. 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.

If you have any questions about this agreement, contact Juniper Networks at the following address:

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089
USA
Attn: Contracts Administrator

Abbreviated Table of Contents

About This Guide	xxi
-------------------------	-----

Part 1 **Using the Configuration Interfaces**

Chapter 1	Using Services Router Configuration Tools	3
------------------	--	----------

Part 2 **Configuring Router Interfaces**

Chapter 2	Interfaces Overview	41
Chapter 3	Configuring Network Interfaces	101
Chapter 4	Configuring Point-to-Point Protocol over Ethernet	143
Chapter 5	Configuring ISDN	159

Part 3 **Configuring Routing Protocols**

Chapter 6	Routing Overview	193
Chapter 7	Configuring Static Routes	223
Chapter 8	Configuring a RIP Network	235
Chapter 9	Configuring an OSPF Network	251
Chapter 10	Configuring BGP Sessions	273

Part 4 Configuring Private Communications over Public Networks with MPLS

Chapter 11	Multiprotocol Label Switching Overview	293
Chapter 12	Configuring Signaling Protocols for Traffic Engineering	309
Chapter 13	Configuring Virtual Private Networks	321
Chapter 14	Configuring CLNS VPNs	345
Chapter 15	Configuring IPSec for Secure Packet Exchange	357

Part 5 Managing Multicast Transmissions

Chapter 16	Multicast Overview	375
Chapter 17	Configuring a Multicast Network	385

Part 6 Configuring Routing Policy, Firewall Filters, and Class of Service

Chapter 18	Policy, Firewall Filter, and Class-of-Service Overview	397
Chapter 19	Configuring Routing Policies	423
Chapter 20	Configuring Firewall Filters and NAT	437
Chapter 21	Configuring Class of Service with DiffServ	477

Part 7 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 Services Router Configuration Tools	3
Configuration Tools Terms	3
Configuration Tools Overview	4
Editing and Committing a Configuration.....	4
J-Web Configuration Options.....	5
CLI Configuration Commands	5
Filtering Configuration Command Output	6
Before You Begin.....	7
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	44
Media Types	45
Network Interface Naming	45
J-series Interface Naming Conventions	46
Understanding CLI Output for J-series Interfaces	48
Data Link Layer Overview.....	49
Physical Addressing.....	49
Network Topology.....	49
Error Notification	49
Frame Sequencing	49
Flow Control.....	50
Data Link Sublayers.....	50
MAC Addressing.....	50
Ethernet Interface Overview	51
Ethernet Access Control and Transmission	51
Collisions and Detection.....	52
Collision Detection	52
Backoff Algorithm	52
Collision Domains and LAN Segments	53
Repeaters	53
Bridges and Switches	53
Broadcast Domains	54
Ethernet Frames	54
T1 and E1 Interfaces Overview	55
T1 Overview.....	56

E1 Overview	56
T1 and E1 Signals	56
Encoding	57
AMI Encoding	57
B8ZS and HDB3 Encoding	57
T1 and E1 Framing	58
Superframe (D4) Framing for T1	58
Extended Superframe (ESF) Framing for T1	58
T1 and E1 Loopback Signals	59
T3 and E3 Interfaces Overview	59
Multiplexing DS1 Signals	59
DS2 Bit Stuffing	60
DS3 Framing	60
M13 Asynchronous Framing	61
C-Bit Parity Framing	62
Serial Interface Overview	64
Serial Transmissions	65
Signal Polarity	66
Serial Clocking Modes	66
Serial Interface Transmit Clock Inversion	67
DTE Clock Rate Reduction	67
Serial Line Protocols	67
EIA-530	68
RS-232	68
RS-422/449	69
V.35	69
X.21	70
ADSL Interface Overview	70
ADSL Systems	71
ADSL2 and ADSL2+	72
Asynchronous Transfer Mode	72
ISDN Interface Overview	73
ISDN Channels	73
ISDN Interfaces	73
Typical ISDN Network	73
NT Devices and S and T Interfaces	74
U Interface	75
ISDN Call Setup	75
Layer 2 ISDN Connection Initialization	75
Layer 3 ISDN Session Establishment	75
Interface Physical Properties	76
Bit Error Rate Testing	77
Interface Clocking	77
Data Stream Clocking	78
Explicit Clocking Signal Transmission	78
Frame Check Sequences	79
Cyclic Redundancy Checks and Checksums	79
Two-Dimensional Parity	79
Physical Encapsulation on an Interface	80
Frame Relay	80
Virtual Circuits	81
Switched and Permanent Virtual Circuits	81
Data-Link Connection Identifiers	81

Congestion Control and Discard Eligibility	81
Point-to-Point Protocol	82
Link Control Protocol	82
CHAP Authentication	83
Network Control Protocols	83
Magic Numbers	84
CSU/DSU Devices	84
Point-to-Point Protocol over Ethernet	85
PPPoE Discovery	85
PPPoE Sessions	86
High-Level Data Link Control	86
HDLC Stations	86
HDLC Operational Modes	87
Interface Logical Properties	87
Protocol Families	88
Common Protocol Suites	88
Other Protocol Suites	89
IPv4 Addressing	89
IPv4 Classful Addressing	89
IPv4 Dotted Decimal Notation	90
IPv4 Subnetting	90
IPv4 Variable-Length Subnet Masks	91
IPv6 Addressing	92
IPv6 Address Representation	92
IPv6 Address Types	93
IPv6 Address Scope	93
IPv6 Address Structure	93
Virtual LANs	94
Special Interfaces	95
Discard Interface	97
Loopback Interface	98
Management Interface	98
Services Interfaces	99
MLPPP and MLFR	99
MLFR Frame Relay Forum	99
CRTP	99

Chapter 3 Configuring Network Interfaces 101

Before You Begin	101
Configuring Network Interfaces with Quick Configuration	102
Configuring an E1 Interface with Quick Configuration	103
Configuring an E3 Interface with Quick Configuration	106
Configuring a Fast Ethernet Interface with Quick Configuration	111
Configuring a T1 Interface with Quick Configuration	113
Configuring a T3 Interface with Quick Configuration	117
Configuring a Serial Interface with Quick Configuration	120
Configuring Network Interfaces with a Configuration Editor	124
Adding a Network Interface with a Configuration Editor	124
Adding an ATM-over-ADSL Network Interface with a Configuration Editor	126

	Configuring CHAP on the ATM-over-ADSL Interface (Optional)	131
	Configuring Compressed Real-Time Transport Protocol (CRTP)	133
	Deleting a Network Interface with a Configuration Editor	134
	Verifying Interface Configuration	135
	Verifying the Link State of All Interfaces	135
	Verifying Interface Properties	136
	Verifying ADSL Interface Properties	137
	Displaying a PPPoA Configuration for an ATM-over-ADSL Interface	141
Chapter 4	Configuring Point-to-Point Protocol over Ethernet	143
	PPPoE Terms	143
	PPPoE Overview	144
	PPPoE Interfaces	145
	Fast Ethernet Interface	145
	ATM-over-ADSL Interface	145
	PPPoE Stages	146
	PPPoE Discovery Stage	146
	PPPoE Session Stage	146
	Optional CHAP Authentication	147
	Before You Begin	147
	Configuring PPPoE with a Configuration Editor	147
	Setting the Appropriate Encapsulation on the Interface (Required)	147
	Configuring PPPoE Encapsulation on an Ethernet Interface	148
	Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface ..	149
	Configuring a PPPoE Interface (Required)	150
	Configuring CHAP (Optional)	152
	Verifying a PPPoE Configuration	153
	Displaying a PPPoE Configuration for an ATM-over-ADSL Interface	154
	Verifying PPPoE Interfaces	155
	Verifying PPPoE Sessions	156
	Verifying the PPPoE Version	157
	Verifying PPPoE Statistics	157
Chapter 5	Configuring ISDN	159
	ISDN Terms	159
	ISDN Overview	160
	ISDN Interfaces	161
	Before You Begin	162
	Configuring ISDN Interfaces with a Configuration Editor	162
	Adding an ISDN Interface (Required)	162
	Configuring a Dialer Interface (Required)	165
	Enabling an ISDN Interface as a Secondary Connection (Optional)	168
	Configuring Dial-on-Demand Connectivity (Optional)	169
	Configuring a Dialer Filter	169
	Applying the Dial-on-Demand Dialer Filter to the Dialer Interface ..	170
	Configuring Bandwidth-on-Demand (Optional)	171
	Configuring a Dialer Interface for Bandwidth-on-Demand	171
	Configuring an ISDN Interface for Bandwidth-on-Demand	173
	Configuring Dial-on-Demand Routing (Optional)	174
	Configuring the Dial-on-Demand Dialer Filter	174

Applying the Dial-on-Demand Dialer Filter to the Dialer Interface...	175
Configuring Dialer Watch (Optional)	176
Adding a Dialer Watch Interface on the Services Router	176
Configuring the ISDN Interface for Dialer Watch	179
Configuring Dial-on-Demand Routing with OSPF Support (Optional)	180
Configuring Dialer Profiles (Optional)	181
Verifying the ISDN Configuration	182
Displaying the ISDN Status	183
Verifying an ISDN Interface	183
Checking B-Channel Statistics	184
Checking D-Channel Interface Statistics	186
Verifying Dialer Interface Configuration	187

Part 3

Configuring Routing Protocols

Chapter 6

Routing Overview 193

Routing Terms	193
Routing Overview	197
Networks and Subnetworks	198
Autonomous Systems	198
Interior and Exterior Gateway Protocols	198
Routing Tables	199
Forwarding Tables	199
Dynamic and Static Routing	200
Route Advertisements	201
Route Aggregation	201
RIP Overview	203
Distance-Vector Routing Protocols	203
Maximizing Hop Count	204
RIP Packets	205
Split Horizon and Poison Reverse Efficiency Techniques	205
Limitations of Unidirectional Connectivity	206
OSPF Overview	207
Link-State Advertisements	208
Role of the Designated Router	208
Path Cost Metrics	209
Areas and Area Border Routers	209
Role of the Backbone Area	210
Stub Areas and Not-So-Stubby Areas	211
BGP Overview	212
Point-to-Point Connections	213
BGP Messages for Session Establishment	214
BGP Messages for Session Maintenance	214
IBGP and EBGP	214
Route Selection	215
Local Preference	216
AS Path	217
Origin	217

	Multiple Exit Discriminator.....	218
	Scaling BGP for Large Networks	218
	Route Reflectors—for Added Hierarchy	219
	Confederations—for Subdivision	221
Chapter 7	Configuring Static Routes	223
	Static Routing Overview.....	223
	Static Route Preferences.....	224
	Qualified Next Hops	224
	Control of Static Routes	224
	Route Retention	225
	Readvertisement Prevention	225
	Forced Rejection of Passive Route Traffic	225
	Default Properties.....	225
	Before You Begin.....	226
	Configuring Static Routes with Quick Configuration	226
	Configuring Static Routes with a Configuration Editor.....	228
	Configuring a Basic Set of Static Routes (Required)	228
	Controlling Static Route Selection (Optional)	229
	Controlling Static Routes in the Routing and Forwarding Tables (Optional).....	231
	Defining Default Behavior for All Static Routes (Optional).....	232
	Verifying the Static Route Configuration	233
	Displaying the Routing Table.....	233
Chapter 8	Configuring a RIP Network	235
	RIP Overview.....	235
	RIP Traffic Control with Metrics.....	235
	Authentication.....	236
	Before You Begin.....	236
	Configuring a RIP Network with Quick Configuration	236
	Configuring a RIP Network with a Configuration Editor.....	239
	Configuring a Basic RIP Network (Required).....	239
	Controlling Traffic in a RIP Network (Optional).....	242
	Controlling Traffic with the Incoming Metric.....	242
	Controlling Traffic with the Outgoing Metric.....	243
	Enabling Authentication for RIP Exchanges (Optional)	245
	Enabling Authentication with Plain-Text Passwords.....	245
	Enabling Authentication with MD5 Authentication	246
	Verifying the RIP Configuration.....	247
	Verifying the RIP-Enabled Interfaces	247
	Verifying the Exchange of RIP Messages.....	247
	Verifying Reachability of All Hosts in the RIP Network	248
Chapter 9	Configuring an OSPF Network	251
	OSPF Overview	251
	Enabling OSPF	251
	OSPF Areas.....	252
	Path Cost Metrics	252

OSPF Dial-on-Demand Circuits	252
Before You Begin.....	253
Configuring an OSPF Network with Quick Configuration	253
Configuring an OSPF Network with a Configuration Editor	255
Configuring the Router Identifier (Required).....	256
Configuring a Single-Area OSPF Network (Required)	256
Configuring a Multiarea OSPF Network (Optional)	258
Creating the Backbone Area.....	259
Creating Additional OSPF Areas.....	259
Configuring Area Border Routers	260
Configuring Stub and Not-So-Stubby Areas (Optional)	261
Tuning an OSPF Network for Efficient Operation	263
Controlling Route Selection in the Forwarding Table.....	264
Controlling the Cost of Individual Network Segments	264
Enabling Authentication for OSPF Exchanges	265
Controlling Designated Router Election	267
Verifying an OSPF Configuration	268
Verifying OSPF-Enabled Interfaces	268
Verifying OSPF Neighbors	269
Verifying the Number of OSPF Routes	270
Verifying Reachability of All Hosts in an OSPF Network.....	271

Chapter 10

Configuring BGP Sessions 273

BGP Overview.....	273
BGP Peering Sessions.....	273
IBGP Full Mesh Requirement.....	274
Route Reflectors and Clusters	274
BGP Confederations	274
Before You Begin.....	275
Configuring BGP Sessions with Quick Configuration	275
Configuring BGP Sessions with a Configuration Editor	277
Configuring a Point-to-Point Peering Session (Required).....	277
Configuring BGP Within a Network (Required)	280
Configuring a Route Reflector (Optional)	281
Configuring BGP Confederations (Optional)	284
Verifying a BGP Configuration	286
Verifying BGP Neighbors	286
Verifying BGP Groups.....	287
Verifying BGP Summary Information	288
Verifying Reachability of All Peers in a BGP Network	289

Part 4

Configuring Private Communications over Public Networks with MPLS

Chapter 11

Multiprotocol Label Switching Overview 293

MPLS and VPN Terms	293
--------------------------	-----

MPLS Overview	295
Label Switching	296
Label-Switched Paths	296
Label-Switching Routers	297
Labels	298
Label Operations	298
Penultimate Hop Popping	299
LSP Establishment	299
Static LSPs	299
Dynamic LSPs	299
Signaling Protocols Overview	300
Label Distribution Protocol	300
LDP Operation	300
LDP Messages	300
Resource Reservation Protocol	300
RSVP Fundamentals	301
Bandwidth Reservation Requirement	301
Explicit Route Objects	301
Constrained Shortest Path First	303
Link Coloring	303
VPN Overview	304
VPN Components	304
VPN Routing Requirements	305
VPN Routing Information	306
VRF Instances	306
Route Distinguishers	306
Route Targets to Control the VRF Table	307
Types of VPNs	307
Layer 2 VPNs	307
Layer 2 Circuits	307
Layer 3 VPNs	307

Chapter 12

Configuring Signaling Protocols for Traffic Engineering 309

Signaling Protocol Overview	309
LDP Signaling Protocol	310
RSVP Signaling Protocol	310
Before You Begin	310
Configuring LDP and RSVP with a Configuration Editor	311
Configuring LDP-Signaled LSPs	311
Configuring RSVP-Signaled LSPs	313
Verifying an MPLS Configuration	316
Verifying an LDP-Signaled LSP	316
Verifying LDP Neighbors	316
Verifying LDP Sessions	317
Verifying the Presence of LDP-Signaled LSPs	318
Verifying Traffic Forwarding over the LDP-Signaled LSP	318
Verifying an RSVP-Signaled LSP	319
Verifying RSVP Neighbors	319
Verifying RSVP Sessions	319
Verifying the Presence of RSVP-Signaled LSPs	320

Chapter 13	Configuring Virtual Private Networks	321
	VPN Configuration Overview	321
	Sample VPN Topology	322
	Basic Layer 2 VPN Configuration	322
	Basic Layer 2 Circuit Configuration	323
	Basic Layer 3 VPN Configuration	323
	Before You Begin.....	324
	Configuring VPNs with a Configuration Editor	324
	Configuring Interfaces Participating in a VPN	325
	Configuring Protocols Used by a VPN	327
	Configuring MPLS for VPNs	327
	Configuring a BGP Session	329
	Configuring Routing Options for VPNs	330
	Configuring an IGP and a Signaling Protocol.....	331
	Configuring LDP for Signaling.....	331
	Configuring RSVP for Signaling	333
	Configuring a Layer 2 Circuit	334
	Configuring a VPN Routing Instance	335
	Configuring a VPN Routing Policy	337
	Configuring a Routing Policy for Layer 2 VPNs	338
	Configuring a Routing Policy for Layer 3 VPNs	341
	Verifying a VPN Configuration	342
	Pinging a Layer 2 VPN.....	343
	Pinging a Layer 3 VPN.....	343
	Pinging a Layer 2 Circuit	343
Chapter 14	Configuring CLNS VPNs	345
	CLNS Terms	345
	CLNS Overview	346
	Before You Begin.....	347
	Configuring CLNS with a Configuration Editor	347
	Configuring a VPN Routing Instance (Required).....	348
	Configuring ES-IS	349
	Configuring IS-IS for CLNS	350
	Configuring CLNS Static Routes.....	352
	Configuring BGP for CLNS.....	353
	Verifying CLNS VPN Configuration	354
	Displaying CLNS VPN Configuration	354
Chapter 15	Configuring IPSec for Secure Packet Exchange	357
	IPSec Tunnel Overview.....	357
	Security Associations	358
	Translating Outgoing Traffic.....	358
	Before You Begin.....	358
	Configuring an IPSec Tunnel with Quick Configuration	358
	Configuring an IPSec Tunnel with a Configuration Editor	360
	Configuring IPSec Services Interfaces	361

Configuring IPSec Service Sets	362
Configuring an IPSec Stateful Firewall Filter Rule	366
Configuring a NAT Pool	368
Verifying the IPSec Tunnel Configuration	370
Verifying IPSec Tunnel Statistics	371

Part 5

Managing Multicast Transmissions

Chapter 16

Multicast Overview 375

Multicast Terms	375
Multicast Architecture	378
Upstream and Downstream Interfaces	378
Subnetwork Leaves and Branches	378
Multicast IP Address Ranges	379
Notation for Multicast Forwarding States	379
Dense and Sparse Routing Modes	380
Strategies for Preventing Routing Loops	380
Reverse-Path Forwarding for Loop Prevention	380
Shortest-Path Tree for Loop Prevention	381
Administrative Scoping for Loop Prevention	381
Multicast Protocol Building Blocks	381

Chapter 17

Configuring a Multicast Network 385

Before You Begin	386
Configuring a Multicast Network with a Configuration Editor	386
Configuring SAP and SDP (Optional)	386
Configuring IGMP (Required)	387
Configuring the PIM Static RP (Optional)	388
Configuring a PIM RPF Routing Table (Optional)	390
Verifying a Multicast Configuration	392
Verifying SAP and SDP Addresses and Ports	392
Verifying the IGMP Version	392
Verifying the PIM Mode and Interface Configuration	393
Verifying the PIM RP Configuration	393
Verifying the RPF Routing Table Configuration	394

Part 6

Configuring Routing Policy, Firewall Filters, and Class of Service

Chapter 18

Policy, Firewall Filter, and Class-of-Service Overview 397

Policy, Firewall Filter, and CoS Terms	397
Routing Policy Overview	399
Routing Policy Components	399
Routing Policy Terms	400

Routing Policy Match Conditions	400
Routing Policy Actions	402
Default and Final Actions	404
Applying Routing Policies	404
Firewall Filter Overview	404
Stateful and Stateless Firewall Filters	405
Process for Configuring a Stateful Firewall Filter and NAT	406
Summary of Stateful Firewall Filter and NAT Match Conditions and Actions	406
Planning a Stateless Firewall Filter	408
Stateless Firewall Filter Match Conditions, Actions, and Action Modifiers	409
Class-of-Service Overview	413
Benefits of DiffServ CoS	413
DSCPs and Forwarding Service Classes	414
JUNOS CoS Functions	415
How Forwarding Classes and Schedulers Work	417
Default Forwarding Class Queue Assignments	417
Default Scheduler Settings	418
Default Behavior Aggregate (BA) Classifiers	419
DSCP Rewrites	420
Sample BA Classification	420

Chapter 19**Configuring Routing Policies****423**

Before You Begin	424
Configuring a Routing Policy with a Configuration Editor	424
Configuring the Policy Name (Required)	425
Configuring a Policy Term (Required)	425
Rejecting Known Invalid Routes (Optional)	426
Injecting OSPF Routes into the BGP Routing Table (Optional)	428
Grouping Source and Destination Prefixes in a Forwarding Class (Optional)	430
Configuring a Policy to Prepend the AS Path (Optional)	431
Configuring Damping Parameters (Optional)	433

Chapter 20**Configuring Firewall Filters and NAT****437**

Before You Begin	438
Configuring a Stateful Firewall Filter with Quick Configuration	438
Configuring a Stateful Firewall Filter with a Configuration Editor	442
Configuring a Stateless Firewall Filter with a Configuration Editor	448
Stateless Firewall Filter Strategies	449
Strategy for a Typical Stateless Firewall Filter	449
Strategy for Handling Packet Fragments	449
Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources	450
Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods	453
Configuring a Routing Engine Firewall Filter to Handle Fragments	459

Applying a Stateless Firewall Filter to an Interface	464
Verifying Firewall Filter Configuration	465
Displaying Firewall Filter Configurations	465
Verifying a Stateful Firewall Filter	470
Displaying Firewall Filter Logs	471
Displaying Firewall Filter Statistics	472
Verifying a Services, Protocols, and Trusted Sources Firewall Filter	473
Verifying a TCP and ICMP Flood Firewall Filter	474
Verifying a Firewall Filter That Handles Fragments	475

Chapter 21

Configuring Class of Service with DiffServ 477

Before You Begin	478
Configuring CoS with DiffServ with a Configuration Editor	478
Configuring a Policer for a Firewall Filter (Required)	479
Configuring and Applying a Firewall Filter for a Multifield Classifier (Required)	480
Assigning Forwarding Classes to Output Queues (Required)	484
Configuring and Applying Rewrite Rules (Required)	485
Configuring and Applying Behavior Aggregate Classifiers (Required) ...	490
Configuring RED Drop Profiles for Assured Forwarding Congestion Control (Required)	494
Configuring Schedulers (Optional)	496
Configuring and Applying Scheduler Maps (Optional)	500
Configuring and Applying Virtual Channels (Optional)	503
Configuring and Applying Adaptive Shaping (Optional)	507
Verifying a DiffServ Configuration	508
Verifying Multicast Session Announcements	509
Verifying an Adaptive Shaper Configuration	509
Verifying a Virtual Channel Configuration	510
Verifying a Virtual Channel Group Configuration	510

Part 7

Index

Index	513
-------------	-----

About This Guide

This preface provides the following guidelines for using this manual and related Juniper Networks, Inc., technical documents:

- 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 7.3 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-series Configuration Tools”	<i>JUNOS System Basics Configuration Guide</i>
“Interfaces Overview”	■ <i>JUNOS Network Interfaces and Class of Service 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 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>
“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 and Class of Service 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 Network Interfaces and 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 and Diagnosing a Services Router”	<ul style="list-style-type: none"> ■ <i>JUNOS System Basics and Services Command Reference</i> ■ <i>JUNOS Interfaces Command Reference</i> ■ <i>JUNOS Routing Protocols and Policies Command Reference</i>
“Monitoring 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 Services Router Configuration Tools on page 3

Chapter 1

Using Services Router Configuration Tools

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

This chapter contains the following topics:

- Configuration Tools Terms on page 3
- Configuration Tools Overview on page 4
- Before You Begin on page 7
- Using J-Web Quick Configuration on page 7
- Using the J-Web Configuration Editor on page 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.

Table 6: Configuration Tools Terms (continued)

Term	Definition
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.
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	
commit	Commit the set of configuration changes in the candidate configuration to take operational effect.
load	Load a configuration from an ASCII configuration file or from terminal input.
rollback	Return to a previously committed configuration.
save	Save the configuration to an ASCII file.
Modifying the Configuration and Its Statements	
activate	Activate a previously deactivated statement or identifier.
annotate	Add a comment to a statement.
copy	Copy and add a statement to the configuration.
deactivate	Deactivate a statement or identifier.
delete	Delete a statement or identifier from the configuration.
insert	Insert an identifier into an existing hierarchy.
rename	Rename an existing statement or identifier.
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

The screenshot displays the Juniper J-Web interface for a ROUTER - J4300. The top navigation bar includes the Juniper Networks logo, the router model, and the user 'regress' logged in. Below the navigation bar are tabs for Monitor, Configuration, Diagnose, and Manage. The left sidebar shows a tree view under 'Quick Configuration' with options: Set Up, SSL, Interfaces, Users, SNMP, Routing, Firewall/NAT, IPSec Tunnels, Realtime Performance Monitoring, View and Edit, History, and Rescue. The main content area shows the 'Quick Configuration' page with a 'Summary' section and a 'Router Configuration' section. The 'Router Configuration' section lists four options: Set Up, SSL, Interfaces, and Users, each with a brief description of its function.

Juniper NETWORKS

ROUTER - J4300

Logged in as: **regress**

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

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

[Configuration > Quick Configuration > Summary](#)

Quick Configuration

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

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.

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

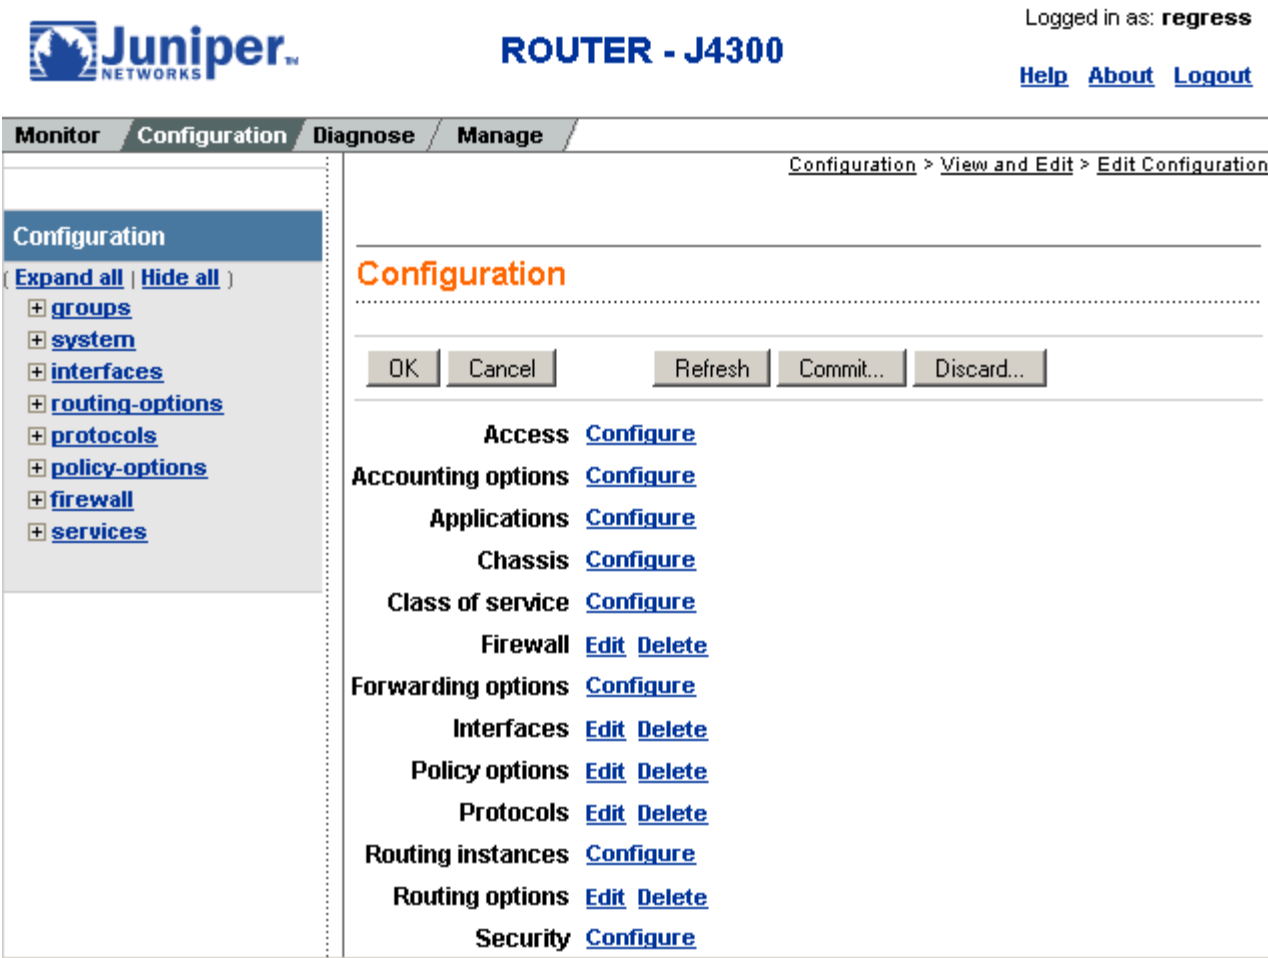
Use the J-Web configuration editor's clickable interface to perform the following configuration tasks on a Services Router:

- 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.

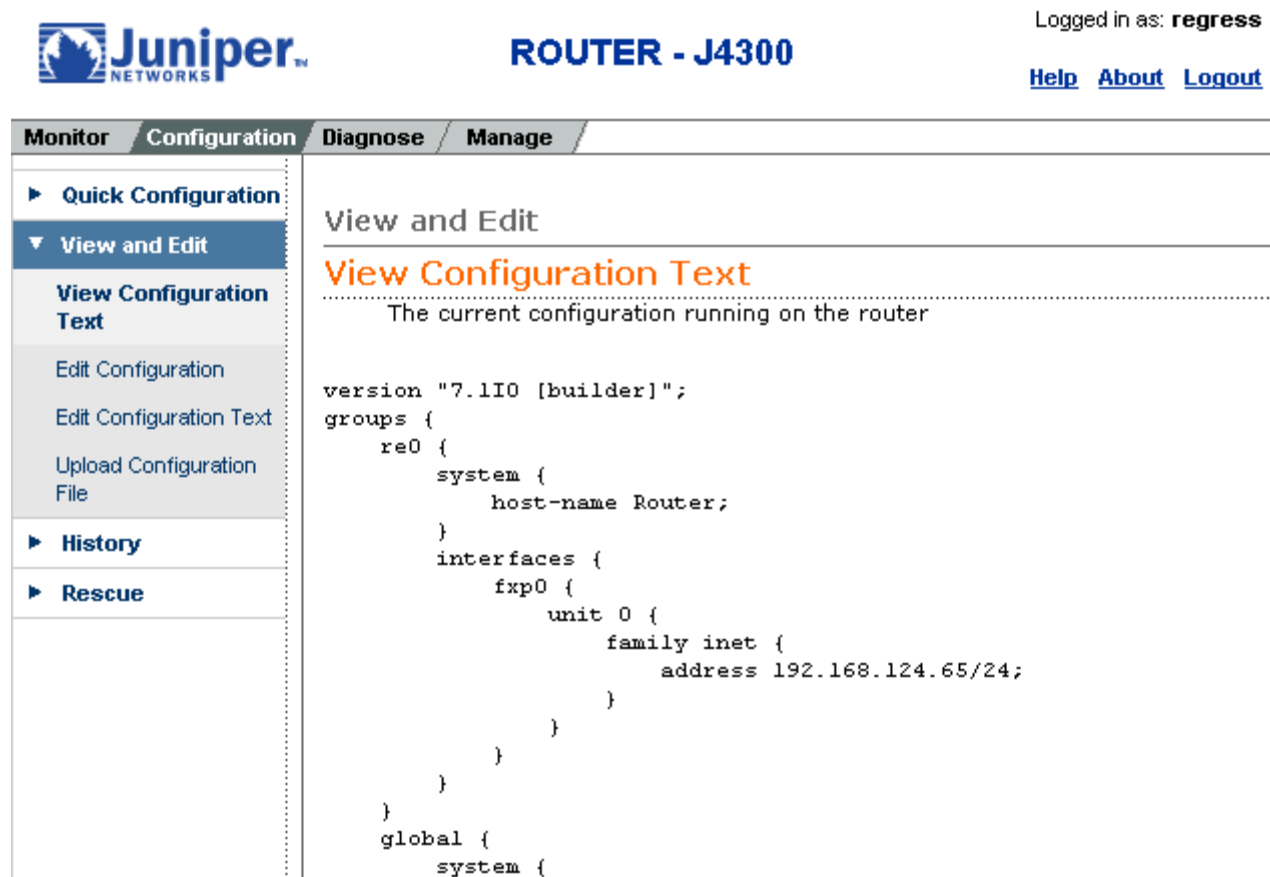
3. To display all the edits applied to the running configuration, click **Refresh**.

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

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

Edit Configuration Text

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.

Configuration

```
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:

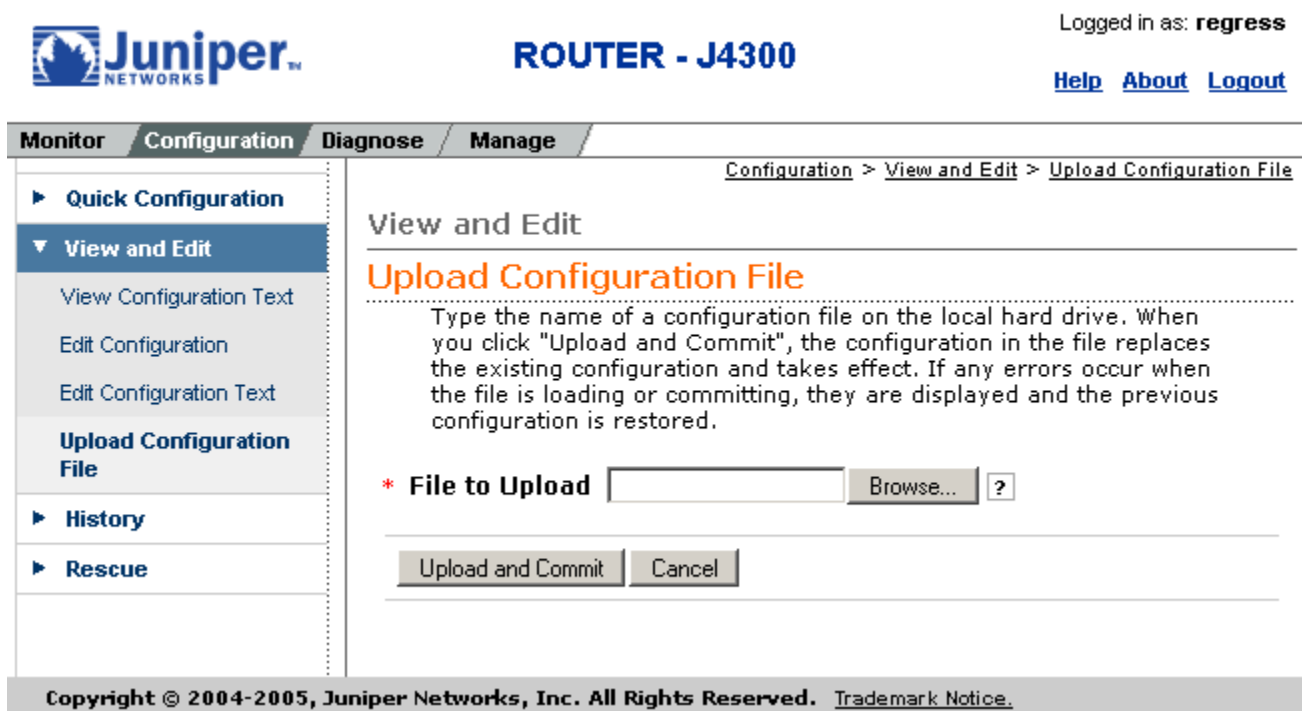
1. Select **Configuration > View and Edit > Upload Configuration File**.

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-ds4JDbisz6A0hcbs2goG"; ## 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 7.

To enter and exit configuration mode:

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

Select the form of the **configure** command (see Table 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.


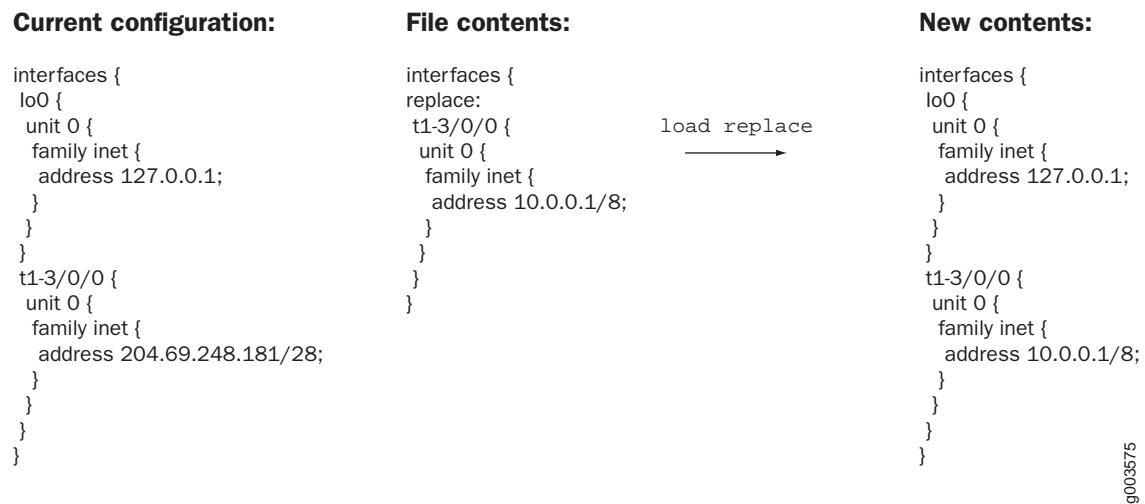
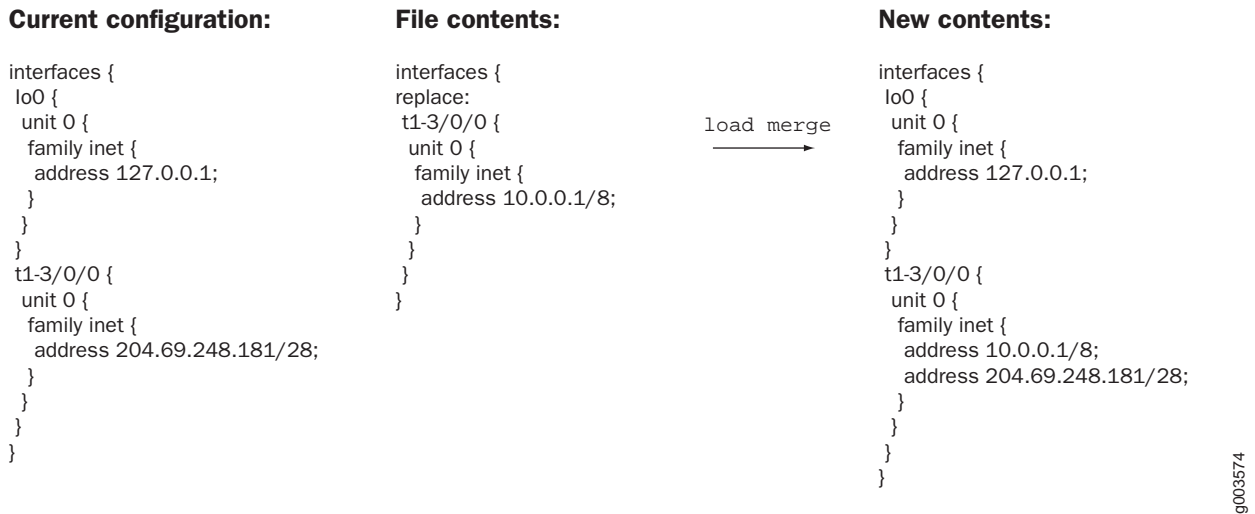
Figure 8: Loading a Configuration with the Override Operation**Figure 9: Loading a Configuration with the Replace Operation**

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 101
- Configuring Point-to-Point Protocol over Ethernet on page 143
- Configuring ISDN on page 159

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 and Class of Service Configuration Guide*, the *JUNOS Services Interfaces Configuration Guide*, and the *JUNOS Interfaces Command Reference*.

- Interfaces Terms on page 42
- Network Interfaces on page 44
- Data Link Layer Overview on page 49
- Ethernet Interface Overview on page 51
- T1 and E1 Interfaces Overview on page 55
- T3 and E3 Interfaces Overview on page 59
- Serial Interface Overview on page 64
- ADSL Interface Overview on page 70
- ISDN Interface Overview on page 73
- Interface Physical Properties on page 76
- Physical Encapsulation on an Interface on page 80
- Interface Logical Properties on page 87
- Special Interfaces on page 95

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	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.992.1 Standard	International Telecommunications 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.
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).

Table 17: Network Interfaces Terms (continued)

Term	Definition
Multilink Point-to-Point Protocol (MLPPP)	Protocol that allows you to bundle multiple Point-to-Point Protocol (PPP) links into a single logical unit. MLPPP improves bandwidth efficiency and fault tolerance and reduces latency.
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
Point-to-Point Protocol (PPP)	Link-layer protocol that provides multiprotocol encapsulation. PPP is used for link-layer and network-layer configuration.
serial interface	<p>Physical LAN interface for transmitting data between computing devices. A Services Router has two types of serial interfaces:</p> <ul style="list-style-type: none"> ■ Asynchronous serial interface—Console port, with speeds up to 110.5 Kbps. The console port supports an RS-232 (EIA-232) standard serial cable with a 25-pin (DB-25) connector. ■ Synchronous serial interface—Port that transmits packets to and from, for example, a T1 device or microwave link, at speeds up to 8 Mbps. You cannot use this serial interface to connect a console. Services Router synchronous serial interfaces support RS-232 (EIA-232), RS-422/449 (EIA-449), RS-530 (EIA-530), V.35, and X.21 cable types. For details, see “Serial Line Protocols” on page 67. <p>For cable details, see the <i>J-series Services Router Getting Started Guide</i>.</p>
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 45
- Network Interface Naming on page 45

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.

- 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 46
- Understanding CLI Output for J-series Interfaces on page 48

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"> ■ <i>at</i>—ATM-over-ADSL WAN interface ■ <i>bc</i>—Bearer channel on an ISDN BRI ■ <i>br</i>—Basic Rate Interface for establishing ISDN connections ■ <i>dc</i>—Delta channel on an ISDN BRI ■ <i>dl</i>—Dialer interface for initiating ISDN connections ■ <i>e1</i>—E1 WAN interface ■ <i>e3</i>—E3 WAN interface ■ <i>fe</i>—Fast Ethernet LAN interface ■ <i>se</i>—Serial interface (either RS-232, RS-422/499, RS-530, V.35, or X.21) ■ <i>t1</i>—T1 (also called DS1) WAN interface ■ <i>t3</i>—T3 (also called DS3) WAN interface <p>Special interface identifiers: (See “Special Interfaces” on page 95.)</p> <ul style="list-style-type: none"> ■ <i>dsc</i> ■ <i>gr, gre</i> ■ <i>ip, ipip</i> ■ <i>lo</i> ■ <i>ls</i> ■ <i>lsi</i> ■ <i>mtun</i> ■ <i>pd, pimd</i> ■ <i>pe, pime</i> ■ <i>sp</i> ■ <i>tap</i>
<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 98.)</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 87.</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 51
- Collisions and Detection on page 52
- Collision Domains and LAN Segments on page 53
- Broadcast Domains on page 54
- Ethernet Frames on page 54

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.

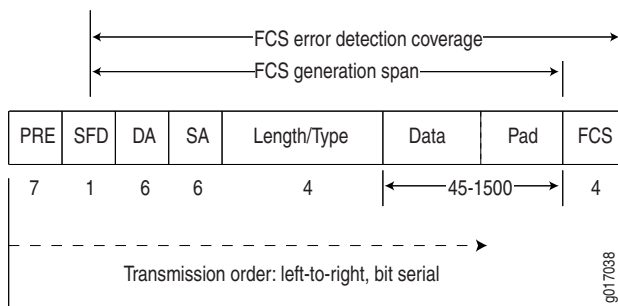
Broadcast Domains

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 56
- E1 Overview on page 56
- T1 and E1 Signals on page 56
- Encoding on page 57
- T1 and E1 Framing on page 58
- T1 and E1 Loopback Signals on page 59

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 57.

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

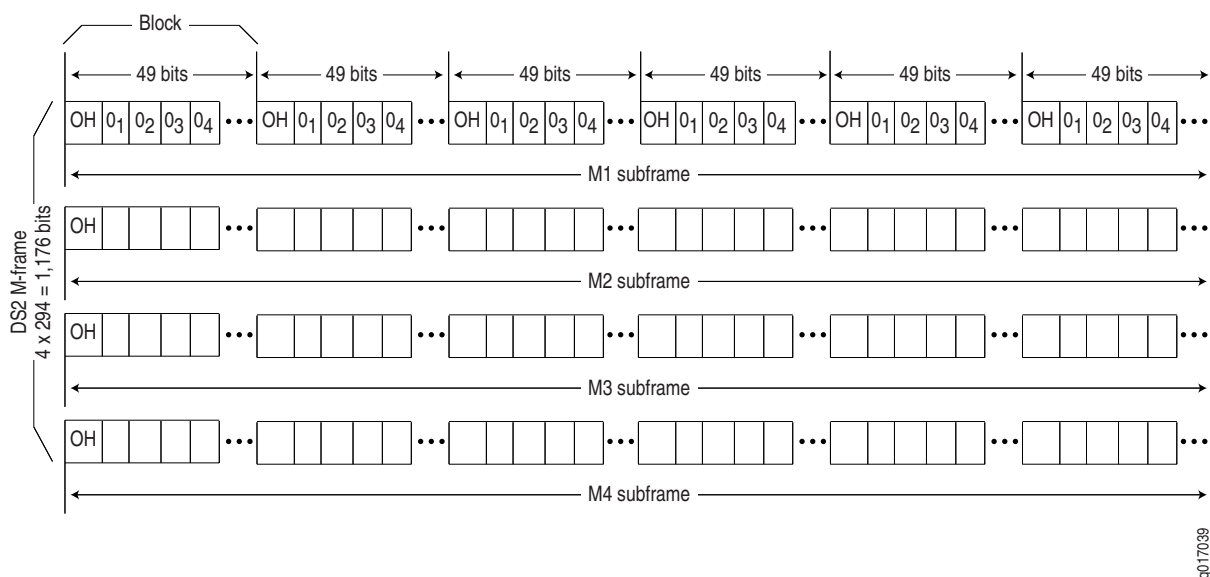
T3 is a high-speed data-transmission medium formed by multiplexing 28 DS1 signals into seven separate DS2 signals, and combining the DS2 signals into a single DS3 signal. T3 links operate at 43.736 Mbps. T3 is also called DS3.

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.

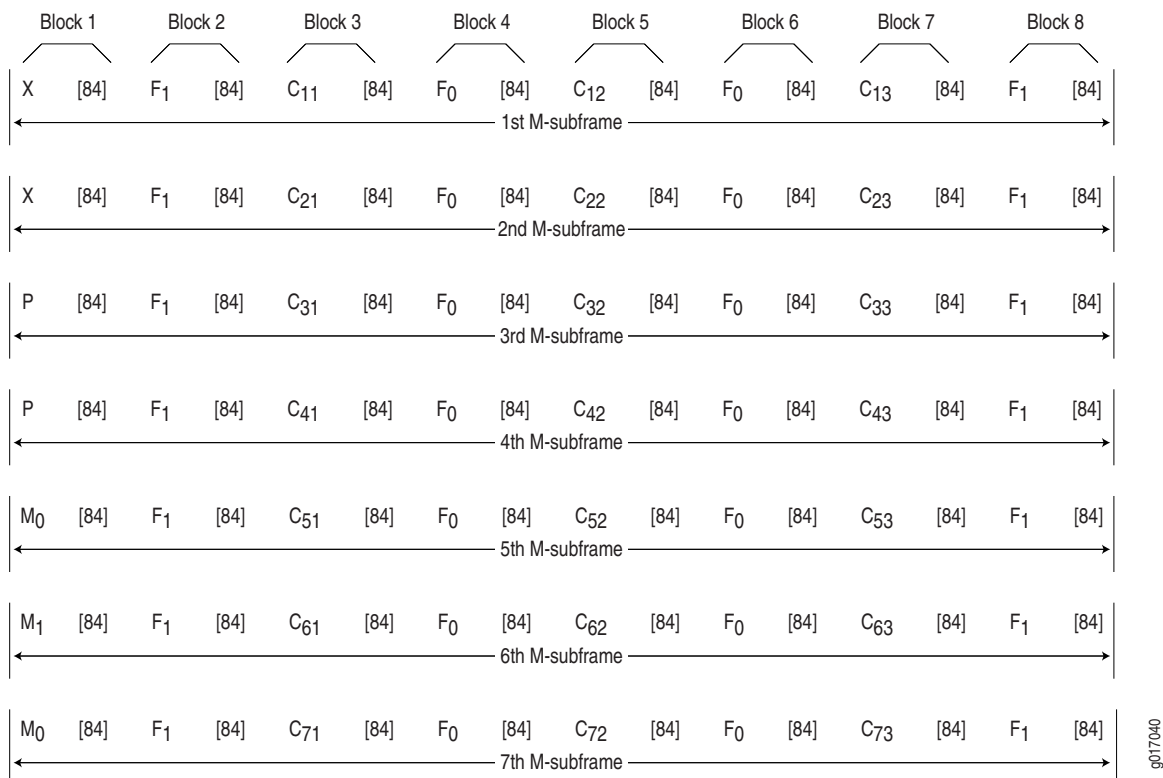
DS3 Framing

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 68

- RS-232 on page 68
- RS-422/449 on page 69
- V.35 on page 69
- X.21 on page 70

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 interfaces can be configured with the properties associated with traditional ATM interfaces, including virtual circuit and path information and ATM encapsulation.

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.

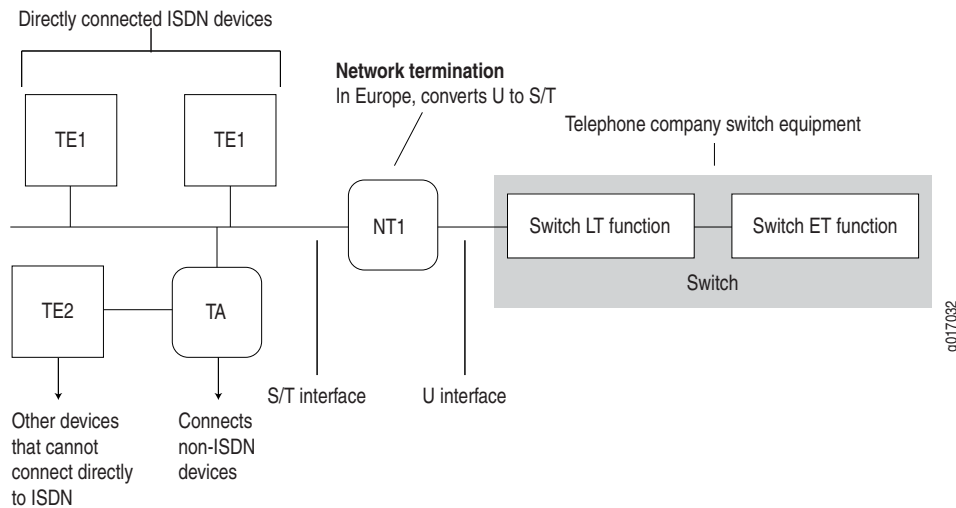
ISDN Interfaces

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 77.
bert-period	Bit error rate test (BERT) time period over which bit errors are sampled. See “Bit Error Rate Testing” on page 77.
chap	Challenge Handshake Authentication Protocol (CHAP). Specifying chap enables CHAP authentication on the interface. See “CHAP Authentication” on page 83.
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 77.
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 80.
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 79.
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 80
- Point-to-Point Protocol on page 82
- Point-to-Point Protocol over Ethernet on page 85
- High-Level Data Link Control on page 86

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 88
- IPv4 Addressing on page 89
- IPv6 Addressing on page 92
- Virtual LANs on page 94

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 97.
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 98.
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 98.
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 99.</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 143.</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 and Class of Service Configuration Guide*. To configure PPPoE interfaces, see “Configuring Point-to-Point Protocol over Ethernet” on page 143. To configure ISDN interfaces, see “Configuring ISDN” on page 159.

- Before You Begin on page 101
- Configuring Network Interfaces with Quick Configuration on page 102
- Configuring Network Interfaces with a Configuration Editor on page 124
- Verifying Interface Configuration on page 135

Before You Begin

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

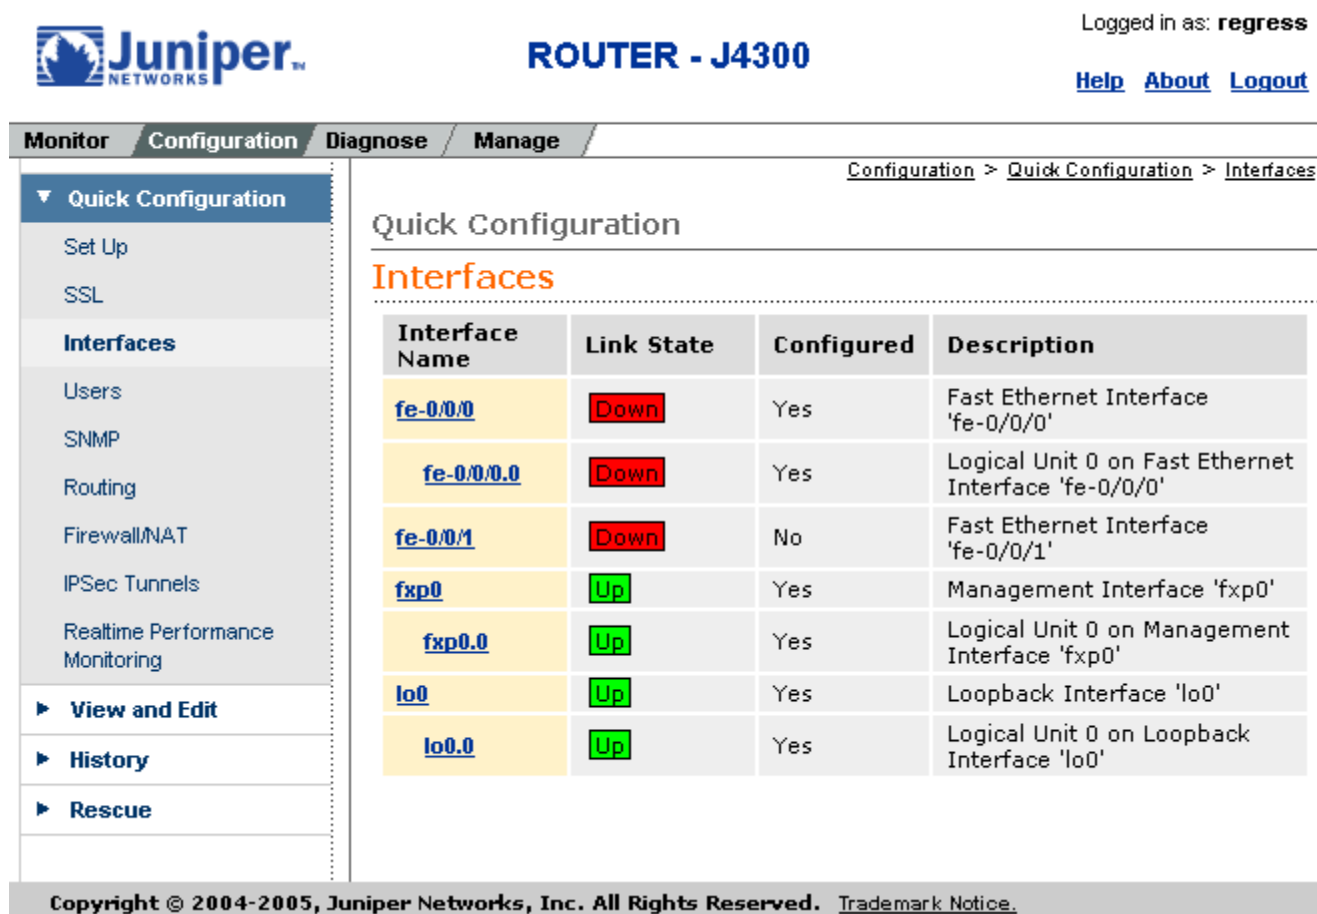
- Install Services Router hardware. For more information, see the *J-series Services Router Getting Started Guide*.
- Establish basic connectivity. For more information, see the *J-series Services Router Getting Started Guide*.
- If you do not already have a basic understanding of physical and logical 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



Juniper NETWORKS ROUTER - J4300

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

Monitor Configuration Diagnose Manage

[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 103
- Configuring an E3 Interface with Quick Configuration on page 106
- Configuring a Fast Ethernet Interface with Quick Configuration on page 111
- Configuring a T1 Interface with Quick Configuration on page 113
- Configuring a T3 Interface with Quick Configuration on page 117
- Configuring a Serial Interface with Quick Configuration on page 120

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. At the top, the Juniper logo and 'ROUTER - J6300' are displayed. The user is logged in as 'regress'. Navigation links for 'Help', 'About', and 'Logout' are present. The main menu includes 'Monitor', 'Configuration', 'Diagnose', and 'Manage'. The 'Configuration' tab is active, showing a breadcrumb trail: 'Configuration > Quick Configuration > Interfaces'. The left sidebar lists various configuration options, with 'Interfaces' selected. The main content area is titled 'Quick Configuration' and 'Interfaces', showing the 'Physical Interface: 'e1-1/0/0''. It includes sections for 'Logical Interfaces' (with an 'Add...' button), 'Physical Interface Description' (with a text input field), and 'Encapsulation' (with a dropdown menu). The 'CHAP Local Identity' section includes a 'Use System Host Name' checkbox and input fields for 'Local Name', 'CHAP Peer Identity', and 'CHAP Secret'.

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 135.

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 in dotted decimal notation. 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)		

Table 24: E1 Quick Configuration Summary (continued)

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

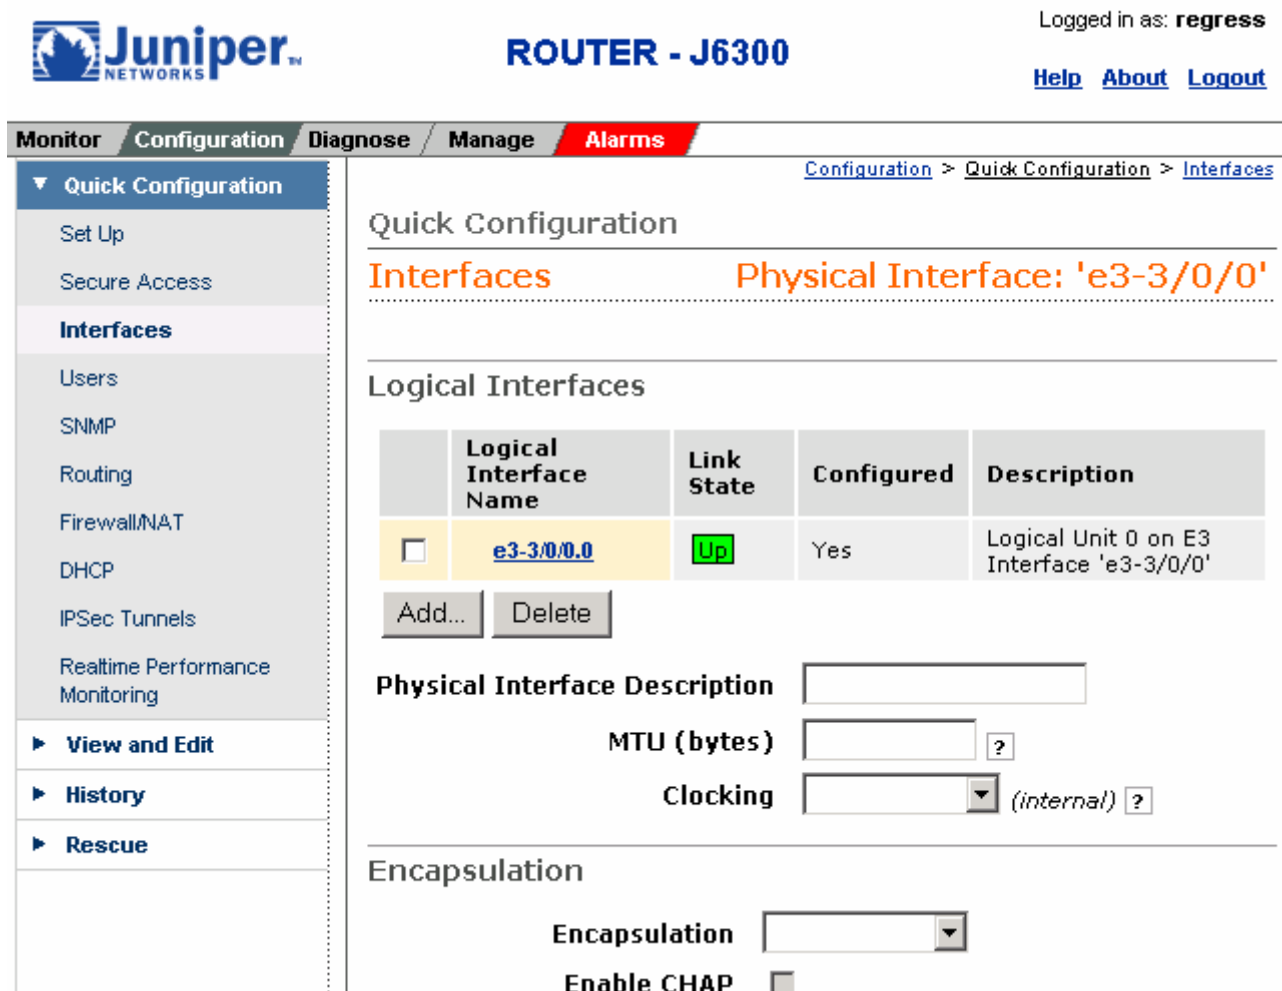
Configuring an E3 Interface with Quick Configuration

To configure properties on an E3 interface:

1. From the Quick Configuration page, as shown in Figure 22, 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**

[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 135.

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 in dotted decimal notation. 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)		

Table 25: E3 Quick Configuration Summary (continued)

Field	Function	Your Action
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.
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 breadcrumb trail is 'Configuration > Quick Configuration > Interfaces'. The left sidebar shows a tree view with 'Quick Configuration' expanded, containing options like 'Set Up', 'SSL', 'Interfaces', '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 displays a table of 'Logical Interfaces' with columns for 'Logical Interface Name', 'Link State', 'Configured', and 'Description'. One interface, 'fe-0/0/0.0', is listed 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. The footer contains copyright information for Juniper Networks, Inc. (2004-2005).

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

Router - J4300

Monitor Configuration Diagnose Manage

Configuration > Quick Configuration > Interfaces

Quick Configuration

Interfaces Physical Interface: 'fe-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'

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 135.

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 in dotted decimal notation. 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 135.

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 in dotted decimal notation. 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)		

Table 27: T1 Quick Configuration Summary (continued)

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

Table 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**

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

Quick Configuration

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

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

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**

2. Enter information into the Quick Configuration page, as described in Table 28.
3. Click one of the following buttons:
 - To apply the configuration and stay on the Quick Configuration page, click **Apply**.
 - To apply the configuration and return to the main configuration page, click **OK**.
 - To cancel your entries and return to the main page, click **Cancel**.
4. To verify that the T3 interface is configured correctly, see “Verifying Interface Configuration” on page 135.

Table 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 in dotted decimal notation. 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)		

Table 28: T3 Quick Configuration Summary (continued)

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

Configuring 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 the Juniper logo, the router model 'ROUTER - J6300', and the user 'regress'. Below this is a secondary navigation bar with 'Monitor', 'Configuration', 'Diagnose', and 'Manage'. The 'Configuration' tab is active, and the breadcrumb trail is 'Configuration > Quick Configuration > Interfaces'. The left sidebar contains a tree view with 'Quick Configuration' expanded, showing options like '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''. 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. The 'Enable CHAP' checkbox is unchecked. Under 'CHAP Local Identity', the 'Use System Host Name' checkbox is checked. The 'Local Name' field is empty. The 'CHAP Peer Identity' and 'CHAP Secret' fields are marked with red asterisks and are empty.

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 135.

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"> 1. Type one or more IPv4 addresses and prefixes in dotted decimal notation. For example: 10.10.10.10/24 2. Click Add. 3. Click OK.
Physical Interface Description	(Optional) Adds supplementary information about the physical serial interface.	Type a text description of the serial interface to more clearly identify it in monitoring displays.
MTU (bytes)	Specifies the maximum transmission unit size for a serial interface.	Type a value between 256 and 9192 bytes. The default MTU for serial interfaces is 1504 .
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.

Table 29: Serial Quick Configuration Summary (continued)

Field	Function	Your Action
Serial Options		
Clocking Mode	<p>Specifies the clock source to determine the timing on serial interfaces.</p> <p>If you use an externally timed clocking mode—dce or loop—long cables might introduce a phase shift of DTE-transmitted clock and data. At high speeds, this phase shift might cause errors.</p> <p>Inverting the transmit clock corrects the phase shift, thereby reducing error rates. By default, the transmit clock is not inverted. To invert the transmit clock, do either of the following:</p> <ul style="list-style-type: none"> ■ In the J-Web configuration editor, set the Transmit clock value to invert on the Interfaces > interface-name > Serial options page. ■ In the CLI configuration editor, include the transmit-clock invert statement at the [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	Specifies the line speed in kilohertz or megahertz for serial interfaces that use the DTE clocking mode.	<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 102. 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 124
- Adding an ATM-over-ADSL Network Interface with a Configuration Editor on page 126
- Configuring CHAP on the ATM-over-ADSL Interface (Optional) on page 131
- Configuring Compressed Real-Time Transport Protocol (CRTP) on page 133
- Deleting a Network Interface with a Configuration Editor on page 134

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

Adding a Network Interface with a Configuration Editor

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 135.

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 45. 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 and Class of Service Configuration Guide*.

You configure the underlying ADSL as an ATM interface, with an interface name of `at-pim/0/port`. Multiple encapsulation types are supported on both the physical and logical ATM-over-ADSL interface.

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 131.
 - 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 143.

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 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</code> <code>oam-liveness up-count 200</code> <code>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</code> <code>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 135.)		
<ul style="list-style-type: none"> ■ atm-ppp-llc—AAL5 logical link control (LLC) encapsulation. ■ atm-ppp-vc-mux—Use AAL5 multiplex encapsulation. 		
Other encapsulation types supported on the ATM-over-ADSL interfaces:		
<ul style="list-style-type: none"> ■ ppp-over-ether-over-atm-llc—Use PPP over Ethernet over ATM LLC encapsulation. When you use this encapsulation type, you cannot configure the interface address. Instead you configure the interface address on the PPP interface. ■ atm-mlppp-llc—Use ATM Multilink PPP over AAL5 LLC encapsulation. ■ 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 menu, select vci. In the Vci field, type <code>35</code>. Next to Oam liveness, click Configure. In the Down count box, type <code>200</code>. In the Up count box, type <code>200</code>. Click OK. Next to Oam period box, click Configure. From the Oam period choices list, select Oam period. In the Oam period box, type <code>100</code>. 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 <code>0</code> and <code>4089</code>—for example, <code>35</code>, with VCI's <code>0</code> through <code>31</code> reserved. Operation, Maintenance, and Administration (OAM) F5 loopback cell thresholds on ATM virtual circuits. The range is between <code>1</code> and <code>255</code>, and the default is <code>5</code> 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, <code>200</code>. Up count—Number of consecutive OAM loopback cells an ATM virtual interface must receive to be identified as operational—for example, <code>200</code>. OAM period—Interval, in seconds, at which OAM cells are transmitted on ATM virtual circuits—for example, <code>100</code>. The range is between <code>1</code> and <code>900</code> 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 and Class of Service 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 135.

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>

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:

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. If you are finished configuring the router, commit the configuration.

Table 33: 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 33: 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 and Class of Service 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 34.



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 34: 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"> In the configuration editor hierarchy, select Configuration > View and Edit > Edit Configuration. 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"> Click Discard. 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 135
- Verifying Interface Properties on page 136
- Verifying ADSL Interface Properties on page 137
- Displaying a PPPoA Configuration for an ATM-over-ADSL Interface on page 141

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:
- 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 143
- PPPoE Overview on page 144
- Before You Begin on page 147
- Configuring PPPoE with a Configuration Editor on page 147
- Verifying a PPPoE Configuration on page 153

## PPPoE Terms

---

Before configuring PPPoE on a Services Router, become familiar with the terms defined in Table 35.

**Table 35: 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 145

- PPPoE Stages on page 146
- Optional CHAP Authentication on page 147

## PPPoE Interfaces

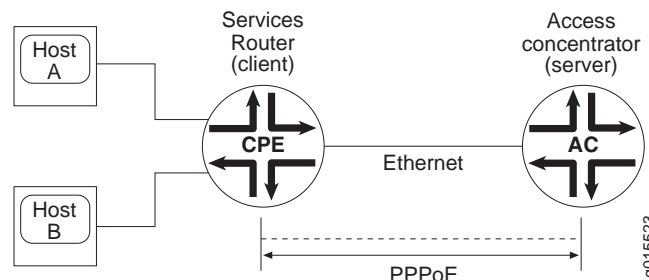
The PPPoE interface to the access concentrator can be either a Fast Ethernet interface on any Services Router or an ATM-over-ADSL interface on a J4300 or J6300 Services Router. The PPPoE configuration is the same for both interfaces. The only difference is the encapsulation for the underlying interface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

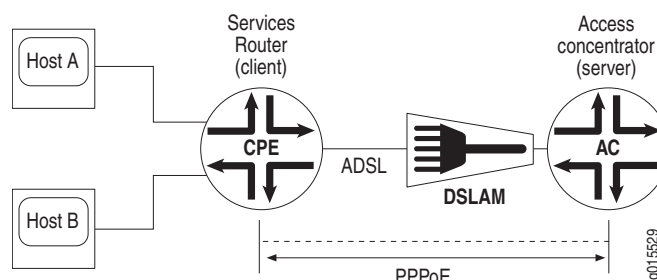
The Services Router encapsulates each PPP frame in an Ethernet frame and transports the frames over an Ethernet loop. Figure 29 shows a typical PPPoE session between a Services Router and an access concentrator on the Ethernet loop.

**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 the 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. This 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 and Class of Service 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 101.

For a PPPoE over ATM interface, see “Adding an ATM-over-ADSL Network Interface with a Configuration Editor” on page 126.

## 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 147
- Configuring a PPPoE Interface (Required) on page 150
- Configuring CHAP (Optional) on page 152

### 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 the 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 148
- Configuring PPPoE Encapsulation on an ATM-over-ADSL Interface on page 149

## **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 36.
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 150.
  - To enable authentication on the interface, see “Configuring CHAP (Optional)” on page 152.
  - To check the configuration, see “Verifying a PPPoE Configuration” on page 153.

**Table 36: Configuring PPPoE Encapsulation on an Ethernet 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                                          |
| 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> to apply your entries to the configuration.</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 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 150.
  - To enable authentication on the interface, see “Configuring CHAP (Optional)” on page 152.
  - To check the configuration, see “Verifying a PPPoE Configuration” on page 153.

**Table 37: Configuring PPPoE Encapsulation on an ATM-over-ADSL 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 |

**Table 37: 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 to apply your entries to the configuration.</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> to apply your entries to the configuration.</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 38.
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 152.
  - To check the configuration, see “Verifying a PPPoE Configuration” on page 153.



**Table 38: 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>d10.0</b> .                                                                                      | <ol style="list-style-type: none"> <li>Next to Backup options, click <b>Configure</b>.</li> <li>In the Interface box, type <b>d10.0</b>.</li> <li>Click <b>OK</b>.</li> </ol>                                                                                                                                                    | Enter<br><br>set backup-options interface d10.0                                                                                                                                                                                                                                                |
| 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 the following:<br/><br/>For the logical Ethernet interface, type <b>fe-0/0/1.0</b>.<br/><br/>For the logical ATM interface type, <b>at-2/0/0.0</b>.</li> </ol>                                   | Enter one of the following:<br><br><ul style="list-style-type: none"> <li>For the logical Ethernet interface, type <b>set pppoe-options underlying-interface fe-0/0/1.0</b>.</li> <li>For the logical ATM interface, type <b>set pppoe-options underlying-interface at-2/0/0.0</b>.</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> to apply your entries to the configuration.</li> </ol>                                                                                                                                                      | Enter<br><br>set pppoe-options service-name video@ispl.com                                                                                                                                                                                                                                     |
| Configure the maximum transmission unit (MTU) of the protocol—for example, <b>1492</b> .                                                                                                     | <ol style="list-style-type: none"> <li>In the Inet box, select <b>Yes</b> and click <b>Configure</b>.</li> <li>In the Mtu box, type <b>1492</b>.</li> </ol>                                                                                                                                                                      | Enter<br><br>up<br><br>set pp0 mtu 1492                                                                                                                                                                                                                                                        |

**Table 38: Configuring a PPPoE Interface (continued)**

| Task                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | CLI Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the PPPoE interface address in one of the following ways: <ul style="list-style-type: none"> <li>Assign source and destination addresses—for example, 192.168.1.1/32 and 192.168.1.2.</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 address:</p> <ol style="list-style-type: none"> <li>Next to Address, click <b>Add new entry</b>.</li> <li>In the Source box, type 192.168.1.1/32.</li> <li>In the Destination box, type 192.168.1.2.</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>Enter <b>up</b>, then do one of the following:</p> <ul style="list-style-type: none"> <li>To assign the source and destination address, type <b>set pp0.0 family inet address 192.168.1.1/32 destination 192.168.1.2</b>.</li> <li>To derive the source address and assign the destination address, type <b>set pp0.0 family inet unnumbered-address lo0.0 destination 192.168.1.2</b>.</li> <li>To obtain an IP address from the remote end, type <b>set pp0.0 family inet negotiate-address</b>.</li> </ul> |
| Disable the sending of keepalives on a logical interface—for example, no-keepalives.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <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><b>set pp0.0 no-keepalives</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

### Configuring CHAP (Optional)

To configure CHAP on the PPPoE interface:

- Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- Perform the configuration tasks described in Table 39.
- If you are finished configuring the router, commit the configuration.
- To check the configuration, see “Verifying a PPPoE Configuration” on page 153.

**Table 39: Configuring CHAP**

| <b>Task</b>                                                                                                                                               | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                             | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                          |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 <b>A-ppp-client</b>.</li> <li>3. Next to Client, click <b>Add new entry</b>.</li> <li>4. In the Name box, type <b>client1</b>.</li> <li>5. In the Chap secret box, type <b>my-secret</b>.</li> <li>6. Click <b>OK</b> until you return to the Configuration page.</li> </ol> |                                                                                                                                                                                                                                                                          |
| Navigate to the <b>pp0</b> unit <b>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 <b>A-ppp-client</b>.</li> </ol>                                                                                                                                                                                          | <p>Enter</p> <p>set ppp-options chap access-profile A-ppp-client</p>                                                                                                                                                                                                     |
| Specify a unique hostname to be used in CHAP challenge and response packets—for example, <b>A-fe-0/0/1.0</b> or <b>A-at-2/0/0.0</b> .                     | <p>In the Local name box, type one of the following:</p> <ul style="list-style-type: none"> <li>■ For the Ethernet interface, type <b>A-fe-0/0/1.0</b>.</li> <li>■ For the ATM interface, type <b>A-at-2/0/0.0</b>.</li> </ul>                                                                                                                                                                                | <p>Enter one of the following:</p> <ul style="list-style-type: none"> <li>■ For the Ethernet interface, type <b>set ppp-options chap local-name A-fe-0/0/1.0</b>.</li> <li>■ For the ATM interface, type <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>set ppp-options chap passive</p>                                                                                                                                                                                                                         |

## Verifying a PPPoE Configuration

To verify PPPoE configuration perform the following tasks:

- Displaying a PPPoE Configuration for an ATM-over-ADSL Interface on page 154
- Verifying PPPoE Interfaces on page 155

- Verifying PPPoE Sessions on page 156
- Verifying the PPPoE Version on page 157
- Verifying PPPoE Statistics on page 157

### ***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, Clocking: 1
 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@isp1.com, Configured AC name: isp1.com,
  Session AC name: belur, AC MAC address: 00:90:1a:40:f6:4e,
  Auto-reconnect timeout: 1 seconds,
  Underlying interface: 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 159
- ISDN Overview on page 160
- Before You Begin on page 162
- Configuring ISDN Interfaces with a Configuration Editor on page 162
- Verifying the ISDN Configuration on page 182

ISDN Terms

Before configuring ISDN, become familiar with the terms defined in Table 40.

Table 40: 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 40: ISDN Terminology (continued)

Term	Definition
dialer profile	Set of characteristics configured for the ISDN dialer interface. Dialer profiles allow the configuration of physical interfaces to be separated from the logical configuration of dialer interfaces required for ISDN connectivity. This feature also allows physical and logical interfaces to be bound together dynamically on a per-connection basis.
dialer watch	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.
floating static route	Routes with an administrative distance greater than the administrative distance of dynamic routes.
integrated services digital network (ISDN)	Digital communication service provided by telecommunication service providers. It is an all-digital dialup (on-demand) service that carries voice, data, and video transmissions over telephone lines.
terminal endpoint identifier (TEI)	Number that identifies a terminal endpoint, an ISDN-capable device attached to an ISDN network through an ISDN interface on the Services Router. The TEI is a number between 0 and 127. The range 0–63 are used for static TEI assignment; 64–126 are used for dynamic assignment; and 127 is used for group assignment.
service profile identifier (SPID)	Number that specifies the services available to you on the service provider switch and defines the feature set ordered when the ISDN service is provisioned.

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 /O/ port*.

Each B-channel is identified by *bc-pim /O/ 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 /O/ 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 101.

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 162
- Configuring a Dialer Interface (Required) on page 165
- Enabling an ISDN Interface as a Secondary Connection (Optional) on page 168
- Configuring Dial-on-Demand Connectivity (Optional) on page 169
- Configuring Bandwidth-on-Demand (Optional) on page 171
- Configuring Dial-on-Demand Routing (Optional) on page 174
- Configuring Dialer Watch (Optional) on page 176
- Configuring Dial-on-Demand Routing with OSPF Support (Optional) on page 180
- Configuring Dialer Profiles (Optional) on page 181

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 41.
3. When you are finished configuring the interface, go to “Configuring a Dialer Interface (Required)” on page 165.
4. To verify that the network interface is configured correctly, see “Verifying the ISDN Configuration” on page 182.

Table 41: Adding an ISDN 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 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"> 1. Next to Interface, click Add new entry. 2. In the Interface name box, type the name of the new interface, br-1/0/3. 3. Click OK. 	
Configure dialer options. <ul style="list-style-type: none"> ■ Name the dialer pool—for example, ISDN-dialer-group. ■ Set the dialer pool priority—for example, 25. 	<ol style="list-style-type: none"> 1. In the Encapsulation column, next to the new interface, click Edit. 2. Next to Dialer options, select Yes, and then click Edit. 3. Next to Pool, click Add new entry. 4. In the Pool identifier box, type isdn-dialer-group. 5. In the Priority field, type 25. 6. Click OK. 	<p>From the [edit interfaces br-1/0/3] hierarchy, enter</p> <p>set dialer-options pool ISDN-dialer-group priority 25</p>
Dialer pool priority has a range from 0 to 255 with 0 designating lowest-priority interfaces and 255 designating the highest-priority interfaces.		

Table 41: Adding an ISDN Interface (continued)

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

Table 41: Adding an ISDN Interface (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure Q.931 timers. Q.931 is a Layer 3 protocol for the setup and termination of connections. The default value for each timer is 10 seconds, but can be configured between 1 and 65536 seconds—for example, 15 .	<ol style="list-style-type: none"> 1. In the T306 field, type 15. 2. In the T310 field, type 15. 	<ol style="list-style-type: none"> 1. Enter <code>set isdn-options t306 15</code> 2. Enter <code>set isdn-options t310 15</code>
Configure when the TEI negotiates with the ISDN provider. <ul style="list-style-type: none"> ■ first-call—Activation does not occur until a call is sent. ■ power-up—Activation occurs when the Services Router is powered on. This is the default value. 	<ol style="list-style-type: none"> 1. From the Tei option list, select power-up. 2. Click OK to return to the Interfaces page. 	To initiate activation at power-up, enter <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 42.
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 168.
 - “Configuring Dial-on-Demand Connectivity (Optional)” on page 169.
 - “Configuring Bandwidth-on-Demand (Optional)” on page 171.
 - “Configuring Dial-on-Demand Routing (Optional)” on page 174.
 - “Configuring Dialer Watch (Optional)” on page 176.
 - “Configuring Dial-on-Demand Routing with OSPF Support (Optional)” on page 180.

- “Configuring Dialer Profiles (Optional)” on page 181.
5. To verify that the network interface is configured correctly, see “Verifying the ISDN Configuration” on page 182.

Table 42: Adding a Dialer Interface to a Services Router

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Configure or Edit. 	From the top of the configuration hierarchy, enter edit interfaces
Create the new interface—for example, d10 . Adding a description can differentiate between different dialer interfaces—for example, backup .	<ol style="list-style-type: none"> 1. Next to Interface, click Add new entry. 2. In the Interface name box, type d10. 3. In the Description box, type backup. 4. Click OK. 	Create and name the interface: <ol style="list-style-type: none"> 1. edit interfaces d10 2. set interfaces d10 description backup
Configure encapsulation options—for example, cisco-hdlc . <ul style="list-style-type: none"> ■ cisco-hdlc—Cisco-compatible High-Level Data Link Control is a group of protocols for transmitting data between network points. ■ ppp—Point-to-Point Protocol is a protocol used for communication between two computers using a serial interface. 	<ol style="list-style-type: none"> 1. In the Encapsulation column, next to the new interface, click Edit. 2. From the Encapsulation list, select cisco-hdlc. 	Enter set encapsulation cisco-hdlc
Enter a hold-time value in milliseconds—for example, 60 . 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"> 1. In the Hold time section, type 60 in the Down box. 2. In the Up box, type 60. 	<ol style="list-style-type: none"> 1. Enter set hold-time down 60 2. Enter set hold-time up 60

Table 42: Adding a Dialer Interface to a Services Router (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the logical unit—for example, 0. NOTE: 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"> Next to Unit, click Add new entry. In the Interface unit number box, type 0. Next to Dialer options, select Yes, and then click Configure. 	Enter set unit 0
Configure dialer options.	<ol style="list-style-type: none"> In the Activation delay box, type 60. In the Deactivation delay box, type 30. In the Idle timeout box, type 30. In the Initial route check box, type 30. In the Pool box, type 3. 	<ol style="list-style-type: none"> Enter edit dialer-options Enter set activation-delay 30 Enter set deactivation-delay 30 Enter set idle-timeout 30 initial-route-check 30 pool 3 Enter
<ul style="list-style-type: none"> ■ Activation delay—Time to wait before activating the backup interface once the primary interface is down—for example, 30. Default value is 0 seconds with a maximum value of 60 seconds. Use only for dialer backup and dialer watch. ■ Deactivation delay—Time to wait before deactivating the backup interface once the primary interface is up—for example, 30. Default value is 0 seconds with a maximum value of 60 seconds. Use only for dialer backup and dialer watch. ■ Idle timeout—Time a connection is idle before disconnecting—for example, 30. Default value is 120 seconds with a range from 0 to 4294967295. ■ Initial route check—Time to wait before checking if the primary interface is up—for example, 30. Default value is 120 seconds with a range of 1 to 300 seconds. ■ Load interval—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. ■ Load threshold—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. ■ Pool—Name of a group of ISDN interfaces configured to use the dialer interface—for example, 3. 		
Configure the remote destination to call—for example, 5551212.	In the Dial string field, type 5551212.	Enter set dial-string 5551212

Table 42: Adding a Dialer Interface to a Services Router (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
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"> 1. Next to Watch list, click Add new entry. 2. In the Prefix field, type 192.1.1.0/24. 3. Click OK until you return to the Unit page. 	Enter set watch-list 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"> 1. Select Inet under Family, and click Edit. 2. Next to Address, click Add new entry. 3. In the Source box, type 172.20.10.2. 4. In the Destination box, type 172.20.10.1. 	<ol style="list-style-type: none"> 1. Enter set family inet address 172.20.10.2 2. Enter set family inet address destination 172.20.10.1

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 43.
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 182.

Table 43: Configuring an Interface for ISDN Backup

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	<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"> 1. In the Interface name column, click the interface name. 1. Under Unit, in the Nested Configuration column, click Edit. 	
Configure the backup interface—for instance, d10.0.	<ol style="list-style-type: none"> 1. Next to Backup options, click Configure. 2. In the Interface box, type d10.0. 3. Click OK until you return to the Interfaces page. 	<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 169
- Applying the Dial-on-Demand Dialer Filter to the Dialer Interface on page 170

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 44.
3. Go on to “Applying the Dial-on-Demand Dialer Filter to the Dialer Interface” on page 170.

Table 44: Configuring the Dialer Filter for Interesting Packets

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Firewall level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Configuration > View and Edit > Edit Configuration. 2. Next to Firewall, click Edit. 	<p>From the top of the configuration hierarchy, enter</p> <pre>edit firewall</pre>
Configure the dialer filter name—for example, int-packet .	<ol style="list-style-type: none"> 1. Next to Inet, click Edit. 2. Next to Dialer filter, click Add new entry. 3. In the Filter name box, type int-packet. 	<ol style="list-style-type: none"> 1. Enter <pre>edit family inet</pre> 2. Then enter <pre>edit dialer-filter int-packet</pre>
<p>Configure the firewall rule name—for example, term1.</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 from and then statements.</p>	<ol style="list-style-type: none"> 1. Next to Term, click Add new entry. 2. In the Rule name field, type term1. 3. Next to From, click Configure. 4. From the Protocol choice list, select Protocol. 5. Next to Protocol, click Add new entry. 6. From the Value keyword list, select icmp. 7. Click OK twice to return to the Term page. 	<p>Enter</p> <pre>set term term1 from protocol icmp</pre>
Configure the then part of the dialer filter.	<ol style="list-style-type: none"> 1. Next to Then, click Configure. 2. From the Designation list, select Note. 	<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 45.
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 182.

Table 45: Applying the Dialer Filter to the Dialer Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	<p>From the top of the configuration editor hierarchy, enter</p> <pre>edit interfaces d10 unit 0</pre>
Select the dialer interface to apply the filter—for example, d10.0	<ol style="list-style-type: none"> 1. In the Interface name column, click d10.0. 2. Under Unit, in the Nested Configuration column, click Edit. 	
Select the dialer filter and apply it to the dialer interface.	<ol style="list-style-type: none"> 1. In the Family section, next to Inet, click Edit. 2. Next to Filter, click Configure. 3. In the Dialer box, type int-packet, the dialer-filter configured in “Configuring a Dialer Filter” on page 169, as the dialer-filter. 4. Click OK. 	<ol style="list-style-type: none"> 1. Enter edit family inet filter 2. Enter set dialer int-packet

Configuring 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 46.
3. Go on to “Configuring an ISDN Interface for Bandwidth-on-Demand” on page 173.

Table 46: Configuring a Dialer Interface for Bandwidth-on-Demand

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<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. 3. Next to the dialer interface name, click Edit. 	<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"> 1. Select multilink-ppp as the encapsulation type. 	<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"> ■ Dial string—Telephone number for the interface to dial that establishes ISDN connectivity—for example, 4085551515. ■ Idle timeout—Time a connection is idle before disconnecting—for example, 300. Default value is 120 seconds with a range from 0 to 4294967295. ■ Load interval—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. ■ Load threshold—Percentage of load on all links—for example 90. Default value is 100 with a range from 0 to 100. ■ Pool—Name of a group of ISDN interfaces configured to use the dialer interface—for example, 3. 	<ol style="list-style-type: none"> 1. In the Unit section, click Dialer options under Encapsulation. 2. Next to Dial string, click Add new entry. 3. In the Value box, type 4085551515 and click OK. 4. In the Idle timeout box, type 300. 5. In the Load interval box, type 90. 6. In the Load threshold box, type 95. 7. In the Pool box, type bw-pool. 8. Click OK. 	<ol style="list-style-type: none"> 1. Enter edit unit 0 2. Enter edit dialer-options 3. Enter set dial-string 4085551515 4. Enter set idle-timeout 300 5. Enter set load-interval 90 6. Enter set load-threshold 95 7. Enter set pool bw-pool
<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>F max period 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"> 1. Next to Compression, select Yes, and then Configure. 2. Select Rtp, and then Configure. 3. In the F max period box, type 100. 4. Next to Queues, click Add new entry. 5. From the Value list, select q3. Then click OK and OK again. 	<ol style="list-style-type: none"> 1. From the edit interfaces dl hierarchy, enter edit unit 0 2. Enter set compression rtp f-max-period 500 queues q3

Table 46: Configuring a Dialer Interface for Bandwidth-on-Demand (continued)

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

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 47.
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 182.

Table 47: Configuring an ISDN Interface for Bandwidth-on-Demand

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<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. 3. Next to the ISDN interface name, click Edit. 	<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 46.</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 br interfaces together when configuring bandwidth-on-demand with a total of eight B-channels providing connectivity.</p>	<ol style="list-style-type: none"> 1. Next to the interface name, click Dialer options. 2. Next to Pool, click Add new entry. 3. In the Pool identifier box, type the name of the dialer pool—for example, bw-pool. 4. Click OK. 	<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 48.

- Go on to “Applying the Dial-on-Demand Dialer Filter to the Dialer Interface” on page 175.

Table 48: Configuring a Dialer Filter for Interesting Packets and Dial-on-Demand Routing

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

- Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
- Perform the configuration tasks described in Table 49.
- 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 182.

Table 49: Applying the Dialer Filter to the Dialer Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Configuration > View and Edit > Edit Configuration. 2. Next to Interfaces, click Edit. 	From the top of the CLI configuration hierarchy, enter edit interfaces
Select the dialer interface to apply the dialer filter—for example, d10.	<ol style="list-style-type: none"> 1. In the Interface name column, click d10. 2. Under Unit, in the Nested Configuration column, click Edit. 	Enter edit interfaces d10 unit 0
Select the dialer filter, ddr-packet , and apply it to the dialer interface.	<ol style="list-style-type: none"> 1. In the Family section, next to Inet, click Edit. 2. Next to Filter, click Configure. 3. In the Dialer box, type ddr-packet. 4. Click OK. 	Enter 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 50.
3. Go on to “Configuring the ISDN Interface for Dialer Watch” on page 179.

Table 50: Adding a Dialer Watch Interface

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

Table 50: 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 edit dialer-options
<ul style="list-style-type: none"> ■ Activation delay—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. 	2. In the Deactivation delay box, type 60.	2. Enter set activation-delay 60
	3. In the Dialer string box, type 18005555555.	3. Enter set deactivation-delay 60
	4. In the Idle timeout box, type 30.	4. Enter set dial-string 18005555555
<ul style="list-style-type: none"> ■ Deactivation delay—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. 	5. In the Initial route check box, type 30.	5. Enter set dialer-options idle-timeout 30
	6. In the Pool box, type dw-group.	6. Enter set initial-route-check 30
<ul style="list-style-type: none"> ■ Dialer string—Telephone number for the interface to dial that establishes ISDN connectivity—for example, 8005555555. 		7. Enter set pool dw-group
<ul style="list-style-type: none"> ■ Idle timeout—Time a connection is idle before disconnecting—for example, 30. Default value is 120 seconds with a range from 0 to 4294967295. 		
<ul style="list-style-type: none"> ■ Initial route check—Time to wait before checking if the primary interface is up—for example, 30. Default value is 120 seconds with a range of 1 to 300 seconds. 		
<ul style="list-style-type: none"> ■ Pool—Name of a group of ISDN interfaces configured to use the dialer interface—for example, dw-group. 		

Table 50: Adding a Dialer Watch Interface (continued)

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

- Navigate to the top of the interfaces configuration hierarchy in either the J-Web or CLI configuration editor.
- Perform the configuration tasks described in Table 51.
- If you are finished configuring the router, commit the configuration.
- To verify that the network interface is configured correctly, see “Verifying the ISDN Configuration” on page 182.

Table 51: Configuring an ISDN Interface for Dialer Watch

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. 3. Next to the ISDN interface name, click Edit. 	<p>From the top of the configuration hierarchy, enter</p> <p>edit interfaces br-1/0/3 dialer options pool dw-group</p>
Configure dialer watch options for each ISDN interface participating in the dialer watch feature.	<ol style="list-style-type: none"> 1. Next to the interface name, click Dialer options. 2. Next to Pool, click Add new entry. 	
Each ISDN interface must have the same pool identifier to participate in dialer watch. Therefore, the dialer watch interface configured in Table 50 is used when configuring the ISDN interface.	<ol style="list-style-type: none"> 3. In the Pool identifier field, type dw-group. 4. Click OK. 	

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 251.

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 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 182.

Table 52: Configuring OSPF Demand Circuits

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Protocols level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Configuration > View and Edit > Edit Configuration. 2. Next to Protocols, click Edit. 3. Next to the OSPF, click Edit. 	<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"> 1. Next to the Area id, click Edit. 2. Next to Interfaces, click Add new entry. 3. In the Interface box, type br-5/0/0. 4. Select Demand circuit. 5. Click OK. 	<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 53, 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 53.
3. When you are finished configuring the router, commit the configuration.

Table 53: Dialer Profile Configuration

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. 3. Under Interface name, click dl0. 	<p>From the top of the configuration editor hierarchy, enter</p> <p>edit interfaces dl0 unit 0</p>

Table 53: Dialer Profile Configuration (continued)

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

Verifying the ISDN Configuration

To verify an ISDN configuration, perform the following tasks:

- Displaying the ISDN Status on page 183
- Verifying an ISDN Interface on page 183
- Checking B-Channel Statistics on page 184
- Checking D-Channel Interface Statistics on page 186

- Verifying Dialer Interface Configuration on page 187

Displaying the ISDN Status

Purpose	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.
Action	From the operational mode in the CLI, enter <code>show isdn status</code> .
Sample Output	<pre>user@host> 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>
What It Means	<p>The output shows a summary of interface information. Verify the following information:</p> <ul style="list-style-type: none"> ■ 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

Purpose	Verify that the ISDN interface is correctly configured.
Action	From the CLI, enter the <code>show interfaces extensive</code> command.
Sample Output	<pre>user@host> 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)
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 193
- Configuring Static Routes on page 223
- Configuring a RIP Network on page 235
- Configuring an OSPF Network on page 251
- Configuring BGP Sessions on page 273

Chapter 6

Routing Overview

At its most fundamental level, routing is the process of delivering a message across a network or networks. This task is divided into two primary components: the exchange of routing information to accurately forward packets from source to destination and the packet-forwarding process.

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 193
- Routing Overview on page 197
- RIP Overview on page 203
- OSPF Overview on page 207
- BGP Overview on page 212

Routing Terms

To understand routing, become familiar with the terms defined in Table 54 .

Table 54: Routing Terms

Term	Definition
adjacency	Portion of the local routing information that pertains to the reachability of a single neighbor over a single circuit or interface.
area	Administrative group of OSPF networks within an autonomous system (AS) that operates independently from other areas in the AS. Multiple areas within an AS reduce the amount of link-state advertisement (LSA) traffic on the network and the size of topology databases.
area border router (ABR)	In OSPF, a router having interfaces in multiple areas of an autonomous system (AS) so that it can link the areas to each other. An area border router maintains a separate topological database for each area it is connected to and shares topology information between areas.
AS path	In BGP, the list of autonomous system (ASs) that a packet must traverse to reach a given set of destinations within a single AS.
autonomous system (AS)	Network or collection of routers under a single administrative authority.
backbone area	In OSPF, the central area in an autonomous system (AS) to which all other areas are connected by area border routers (ABRs). The backbone area always has the area ID 0.0.0.0.
bidirectional connectivity	Ability of directly connected devices to communicate with each other over the same link.
Border Gateway Protocol (BGP)	Exterior gateway protocol used to exchange routing information among routers in different autonomous systems.
broadcast	Operation of sending network traffic from one network node to all other network nodes.
cluster	In BGP, a set of routers that have been grouped together. A cluster consists of one system that acts as a route reflector, along with any number of client peers. The client peers receive their route information only from the route reflector system. Routers in a cluster do not need to be fully meshed.
confederation	In BGP, a group of autonomous systems (ASs) that appears to external ASs to be a single AS.
confederation sequence	Ordered set of autonomous systems (ASs) for a confederation. The closest AS in the path is first in the sequence.
convergence	After a topology change, the time all the routers in a network take to receive the information and update their routing tables.
cost	Unitless number assigned to a path between neighbors, based on throughput, round-trip time, and reliability. The sum of path costs between source and destination hosts determines the overall path cost. OSPF uses the lowest cost to determine the best path.
designated router (DR)	In OSPF, a node designated to process link-state advertisements (LSAs) and distribute topology updates for an autonomous system (AS).
distance vector	Number of hops to a routing destination.
dynamic routing	Routing method that enables the route of a message through a network to change as network conditions change. Compare <i>static routing</i> .
exterior gateway protocol (EGP)	Protocol that exchanges routing information between autonomous systems (ASs). BGP is an EGP. Compare <i>interior gateway protocol (IGP)</i> .
external BGP (EBGP)	BGP configuration in which sessions are established between routers in different autonomous systems (ASs).
external peer	In BGP, a peer that resides in a different autonomous system (AS) from the Services Router.
external route	Route to an area outside the network.

Table 54: Routing Terms (continued)

Term	Definition
flooding	Technique by which a router forwards traffic to every node attached to the router, except the node from which the traffic arrived. Flooding is a simple but sometimes inefficient way to distribute routing information quickly to every node in a network. RIP and OSPF are flooding protocols, but BGP is not.
forwarding table	JUNOS software forwarding information base (FIB). The JUNOS routing protocol process installs active routes from its routing tables into the Routing Engine forwarding table. The kernel copies this forwarding table into the Packet Forwarding Engine, which is responsible for determining which interface transmits the packets.
full mesh	Network in which devices are organized in a mesh topology, with each node connected to every other network node.
gateway router	Node on a network that serves as an entrance to another network.
global AS	Global autonomous system (AS). An AS consisting of multiple subautonomous systems (sub-ASs).
handshake	Process of exchanging signaling information between two communications devices to establish the method and transmission speed of a connection.
hello packet	In OSPF, a packet sent periodically by a router to first establish and then maintain network adjacency, and to discover neighbor routers.
hold time	Maximum number of seconds allowed to elapse between the time a BGP system receives successive keepalive or update messages from a peer.
hop	Trip a data packet takes from one router to another in the network. The number of routers through which a packet passes to get from its source to its destination is known as the hop count. In general, the best route is the one with the shortest hop count.
Intermediate System-to-Intermediate System (IS-IS)	Link-state, interior gateway routing protocol for IP networks that also uses the shortest-path-first (SPF) algorithm to determine routes.
interior gateway protocol (IGP)	Protocol that exchanges routing information within autonomous systems (ASs). IS-IS, OSPF, and RIP are IGPs. Compare <i>exterior gateway protocol (EGP)</i> .
Internal BGP (IBGP)	BGP configuration in which sessions are established between routers in the same autonomous systems (ASs).
internal peer	In BGP, a peer that resides in the same autonomous system (AS) as the Services Router.
keepalive message	Periodic message sent by one BGP peer to another to verify that the session between them is still active.
latency	Delay.
link-state advertisement (LSA)	Messages that announce the presence of OSPF-enabled interfaces to adjacent OSPF interfaces (neighbors). The exchange of LSAs establishes bidirectional connectivity between neighbors.
local preference	Optional BGP path attribute carried in internal BGP update packets that indicates the degree of preference for an external route.
mesh	Network topology in which devices are organized in a manageable, segmented manner with many, often redundant, interconnections between network nodes. See also <i>full mesh</i> .
metric	Numerical value that determines how quickly a packet can reach its destination. See also <i>cost</i> .
multiple exit discriminator (MED)	Optional BGP path attribute consisting of a metric value that is used to determine the exit point to a destination when all other factors in determining the exit point are equal.
neighbor	Adjacent router interface. A node can directly route packets to its neighbors only. See also <i>peer</i> .

Table 54: Routing Terms (continued)

Term	Definition
network	Series of nodes interconnected by communication paths.
network diameter	Maximum hop count in a network.
network topology	Arrangement of nodes and connections in a network.
node	Connection point that operates as a redistribution point or an end point in a network, recognizing data transmissions and either forwarding or processing them.
notification message	Message sent between BGP peers to inform the receiving peer that the sending peer is terminating the session because an error occurred, and explaining the error.
not-so-stubby area (NSSA)	In OSPF, a type of stub area in which external route advertisements can be flooded.
open message	Message sent between BGP peers to establish communication.
Open Shortest Path First protocol (OSPF)	A link-state interior gateway protocol (IGP) that makes routing decisions based on the shortest-path-first (SPF) algorithm (also referred to as the Dijkstra algorithm).
origin	Value assigned to a BGP route to indicate whether the first router to advertise the route learned it from an external, internal, or unknown source.
path-vector protocol	Protocol that uses the path between autonomous systems (ASs) to select the best route, rather than the shortest distance or the characteristics of the route (link state). BGP is a path-vector protocol. In contrast, RIP is a distance-vector protocol, and OSPF and IS-IS are link-state protocols.
peer	Immediately adjacent router with which a protocol relationship has been established. See also <i>neighbor</i> .
peering	The practice of exchanging Internet traffic with directly connected peers according to commercial and contractual agreements.
point of presence (POP)	Access point to the Internet, having a unique IP address, where telecommunications equipment is located. POPs usually belong to Internet service providers (ISPs) or telephone companies.
poison reverse	An efficiency technique in a RIP network. By setting the number of hops to an unavailable router to 16 hops or more, a router informs all the other routers in the network. Because RIP allows only up to 15 hops to another router, this technique reduces RIP updates and helps defeat large routing loops. See also <i>split horizon</i> .
propagation	Process of translating and forwarding route information discovered by one routing protocol in the update messages of another routing protocol. Route propagation is also called route redistribution.
reachability	In BGP, the feasibility of a route.
round-robin	Scheduling algorithm in which items have the same priority and are handled in a fixed cyclic order.
route advertisement	Distribution of routing information at specified intervals throughout a network, to establish adjacencies with neighbors and communicate usable routes to active destinations. See also <i>link-state advertisement (LSA)</i> .
route aggregation	Combining groups of routes with common addresses into a single entry in the routing table, to decrease routing table size and the number of route advertisements sent by a router.
route reflection	In BGP, configuring a group of routers into a cluster and having one system act as a route reflector, redistributing routes from outside the cluster to all routers in the cluster. Routers in a cluster do not need to be fully meshed.
Routing Information Protocol (RIP)	Distance-vector routing protocol that keeps a database of routing information gathered from periodic broadcasts by each router in a network.

Table 54: Routing Terms (continued)

Term	Definition
routing table	Table stored on a router that keeps track of all possible paths (routes) between sources and destinations in a network and, in some cases, metrics associated with the routes.
split horizon	An efficiency technique in a RIP network. A router reduces the number of RIP updates in the network by not retransmitting a route advertisement out the interface through which it was received. Split-horizon updates also help prevent routing loops. See also <i>poison reverse</i> .
static routing	Routing method in which routes are manually entered in the routing table and do not change unless you explicitly update them. Unlike dynamic routes, which must be imported into the routing table each time a host comes online, static routes are available immediately. Static routes are generally preferred over other types of routes. Compare <i>dynamic routing</i> .
stub area	In OSPF, an area through which or into which autonomous system (AS) external route advertisements are not flooded.
subautonomous system (sub-AS)	Autonomous system (AS) members of a BGP confederation.
subnetwork	Subdivision of a network, which functions exactly like a network except that it has a more specific address and subnet mask (destination prefix).
three-way handshake	Process by which two routers synchronize protocols and establish a bidirectional connection.
topology database	Map of connections between the nodes in a network. The topology database is stored in each node.
triggered update	In a network that uses RIP, a routing update that is automatically sent whenever routing information changes.
virtual link	In OSPF, a link you create between two area border routers (ABRs) that have an interface to a common nonbackbone area, to connect a third area to the backbone area. One of the area border routers must be directly connected to the backbone area.

Routing Overview

Routing is the transmission of data packets from a source to a destination address. For packets to be correctly forwarded to the appropriate host address, the host must have a unique numeric identifier or IP address. The unique IP address of the destination host forms entries in the routing table. These entries are primarily responsible for determining the path that a packet traverses when transmitted from source to destination.

This overview contains the following topics:

- Networks and Subnetworks on page 198
- Autonomous Systems on page 198
- Interior and Exterior Gateway Protocols on page 198
- Routing Tables on page 199
- Forwarding Tables on page 199

- Dynamic and Static Routing on page 200
- Route Advertisements on page 201
- Route Aggregation on page 201

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:

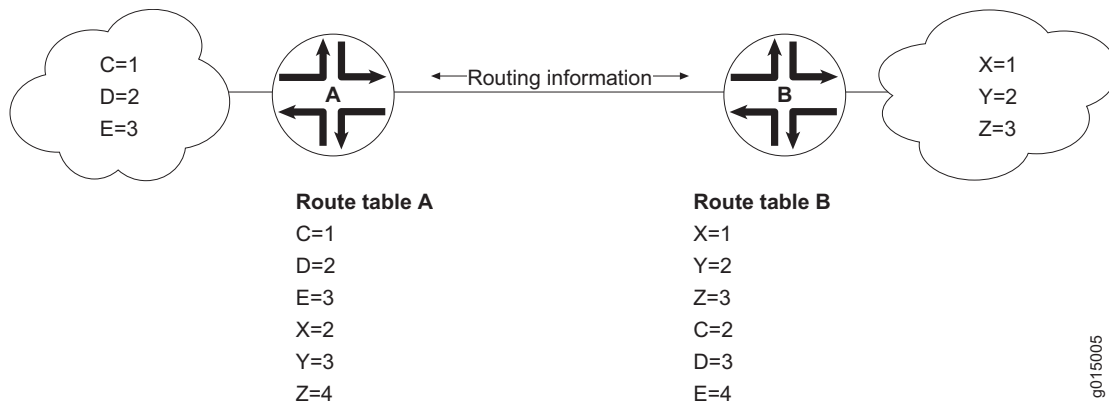
- Distance-Vector Routing Protocols on page 203
- Maximizing Hop Count on page 204
- RIP Packets on page 205
- Split Horizon and Poison Reverse Efficiency Techniques on page 205
- Limitations of Unidirectional Connectivity on page 206



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 208.

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 208
- Role of the Designated Router on page 208
- Path Cost Metrics on page 209
- Areas and Area Border Routers on page 209
- Role of the Backbone Area on page 210
- Stub Areas and Not-So-Stubby Areas on page 211




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.

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 and OSPF, 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 213
- BGP Messages for Session Establishment on page 214
- BGP Messages for Session Maintenance on page 214
- IBGP and EBGP on page 214
- Route Selection on page 215
- Local Preference on page 216
- AS Path on page 217
- Origin on page 217
- Multiple Exit Discriminator on page 218
- Scaling BGP for Large Networks on page 218

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 EBGp

BGP uses two primary modes of information exchange, internal BGP (IBGP) and external BGP (EBGP), to communicate with internal and external peers, respectively.

Peer ASs establish links through an external peer BGP session. As a result, all route advertisement between the external peers takes place by means of the EBGP mode of information exchange. To propagate the routes through the AS and advertise them to internal peers, BGP uses IBGP. To advertise the routes to a different peer AS, BGP again uses EBGP.

To avoid routing loops, IBGP does not advertise routes learned from an internal BGP peer to other internal BGP peers. For this reason, BGP cannot propagate routes throughout an AS by passing them from one router to another. Instead, BGP requires that all internal peers be fully meshed so that any route advertised by one router is advertised to all peers within the AS.

As a network grows, the full mesh requirement becomes increasingly difficult to manage. In a network with 1000 routers, the addition of a single router requires that all the routers in the network be modified to account for the new addition. To combat these scaling problems, BGP uses route reflection and BGP confederations.

For information about route reflection, see “Scaling BGP for Large Networks” on page 218. For information about routing confederations, see “Scaling BGP for Large Networks” on page 218.

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 216.)
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 217.)
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 217.)
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 218.)


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 219

- Confederations—for Subdivision on page 221

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).

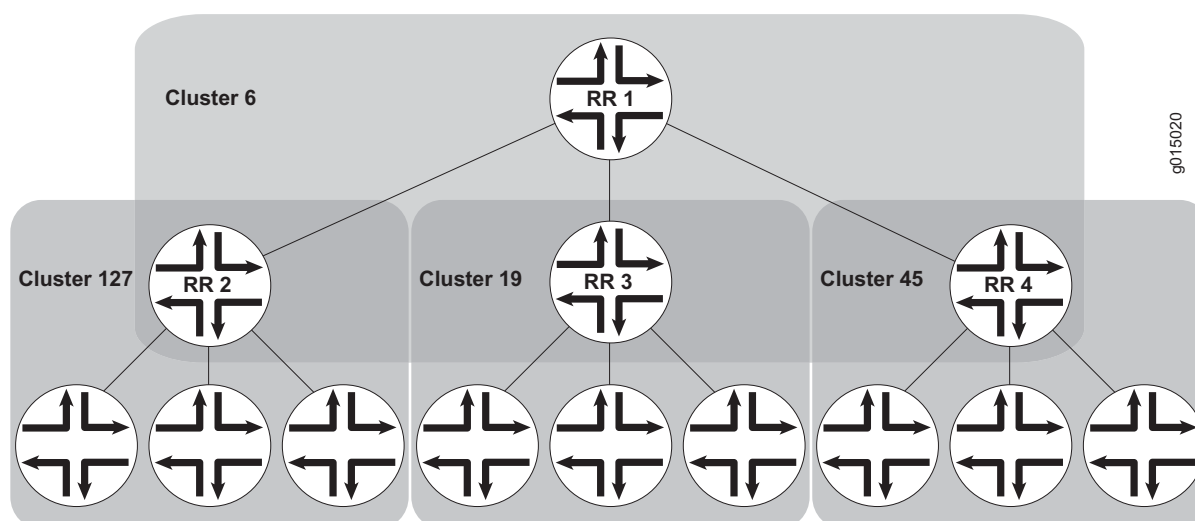
Figure 49: Hierarchical Route Reflection (Clusters of Clusters)

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 223
- Before You Begin on page 226
- Configuring Static Routes with Quick Configuration on page 226
- Configuring Static Routes with a Configuration Editor on page 228
- Verifying the Static Route Configuration on page 233

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 224
- Qualified Next Hops on page 224
- Control of Static Routes on page 224
- Default Properties on page 225

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 225.
- **no-readvertise**—Prevents the route from being readvertised to other routing protocols. For more information, see “Readvertisement Prevention” on page 225.
- **passive**—Rejects traffic destined for the route. For more information, see “Forced Rejection of Passive Route Traffic” on page 225.

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 101.

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

Logged in as: **regress**

Router - J4300

[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

Default Route

Default Route

Static Routes

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

Add... **Delete**

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 55.
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 233.

Table 55: Static Routing Quick Configuration Summary

Field	Function	Your Action
Default Route		
Default Route	Specifies the default gateway for the router.	Type the 32-bit IP address of the Services Router's default route in dotted decimal notation.
Static Routes		
Static Route Address (required)	Specifies the static route to add to the routing table.	<ol style="list-style-type: none"> 1. On the main static routing Quick Configuration page, click Add. 2. In the Static Route Address box, type the 32-bit IP address of the static route in dotted decimal notation.
Next-Hop Addresses	Specifies the next-hop address or addresses to be used when routing traffic to the static route.	<ol style="list-style-type: none"> 1. In the Add box, type the 32-bit IP address of the next-hop host. 2. Click Add. 3. Add more next-hop addresses as necessary. <p>NOTE: If a route has multiple next-hop addresses, traffic is routed across each address in round-robin fashion.</p> <ol style="list-style-type: none"> 4. When you have finished adding next-hop addresses, click OK.

Configuring Static Routes with a Configuration Editor

To configure static routes on the Services Router, you must perform the following tasks marked *(Required)*.


- Configuring a Basic Set of Static Routes (Required) on page 228
- Controlling Static Route Selection (Optional) on page 229
- Controlling Static Routes in the Routing and Forwarding Tables (Optional) on page 231
- Defining Default Behavior for All Static Routes (Optional) on page 232

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

Configuring a Basic Set of Static Routes (Required)

Customer routes that are connected to stub networks are often configured as static routes. Figure 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 56.


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 229.
 - 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 231.
 - To define default properties for static routes, see “Defining Default Behavior for All Static Routes (Optional)” on page 232.
 - To check the configuration, see “Verifying the Static Route Configuration” on page 233.

Table 56: Configuring Basic Static Routes

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

Controlling Static Route Selection (Optional)

When multiple next hops exist for a single static route (see Figure 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 57.
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 231.
 - To define default properties for static routes, see “Defining Default Behavior for All Static Routes (Optional)” on page 232.
 - To check the configuration, see “Verifying the Static Route Configuration” on page 233.

Table 57: Controlling Static Route Selection

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Static level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Static .	From the top of the configuration hierarchy, enter edit routing-options static
Add the static route 192.168.47.5/32 , and define the next-hop address 10.10.10.10 .	<ol style="list-style-type: none"> 1. In the Route field, click Add new entry. 2. In the Destination box, type 192.168.47.5/32. 3. From the Next hop list, select Next hop. 4. In the Next hop field, click Add new entry. 5. In the Value box, type 10.10.10.10. 6. Click OK. 	Define the static route and set the next-hop address: set route 192.168.47.5 next-hop 10.10.10.10
Set the preference for the 10.10.10.10 next hop to 7 .	<ol style="list-style-type: none"> 1. Under Preference, in the Metric value box, enter 7. 2. Click OK. 	Set the preference to 7: set route 192.168.47.5 next-hop 10.10.10.10 preference 7
Define the qualified next-hop address 10.10.10.7 .	<ol style="list-style-type: none"> 1. In the Qualified next hop field, click Add new entry. 2. In the Nexthop field, enter 10.10.10.7. 3. Click OK. 	Set the qualified-next-hop address: set route 192.168.47.5 qualified-next-hop 10.10.10.7
Set the preference for the 10.10.10.7 qualified next hop to 6 .	<ol style="list-style-type: none"> 1. Under Preference, in the Metric value box, enter 6. 2. Click OK. 	Set the preference to 6: set route 192.168.47.5 qualified-next-hop 10.10.10.7 preference 6

Controlling Static Routes in the Routing and Forwarding Tables (Optional)

Static routes have a number of attributes that define how they are inserted and maintained in the routing and forwarding tables. To customize this behavior for the static route **192.168.47.5/32**, perform these steps:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 58.
3. If you are finished configuring the router, commit the configuration.
4. Go on to one of the following procedures:

- To define default properties for static routes, see “Defining Default Behavior for All Static Routes (Optional)” on page 232.
- To check the configuration, see “Verifying the Static Route Configuration” on page 233.

Table 58: Controlling Static Routes in the Routing and Forwarding Tables

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the 192.168.47.5/32 level in the configuration hierarchy.	In the configuration editor hierarchy, select Routing options > Static , then click 192.168.47.5/32 in the Destination field.	From the top of the configuration hierarchy, enter edit routing-options static route 192.168.47.5/32
Specify that the route is to be retained in the forwarding table after the routing process shuts down. By default, static routes are not retained.	1. Next to Retain, select the Yes check box. 2. Click OK .	Set the retain attribute: set retain
Specify that the static route is not to be readvertised. By default, static routes are eligible to be readvertised.	1. Next to Readvertise, select the No check box. 2. Click OK .	Set the no-readvertise attribute: set no-readvertise
Specify that the static route is to be included in the routing table whether the route is active or not. By default, passive routes are not included in the routing table.	1. From the Passive flag list, select Passive . 2. Click OK .	Set the passive attribute: set passive

Defining Default Behavior for All Static Routes (Optional)

Attributes that define static route behavior can be configured either at the individual route level or as a default behavior that applies to all static routes. In the case of conflicting configuration, the configuration at the individual route level overrides static route defaults. To configure static route defaults, perform these steps:

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 59.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying the Static Route Configuration” on page 233.

Table 59: Defining Static Route Defaults

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Defaults level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Static , and then click Configure next to Defaults.	From the top of the configuration hierarchy, enter edit routing-options static defaults
Specify that the route is to be retained in the forwarding table after the routing process shuts down. By default, static routes are not retained.	1. Next to Retain, select the Yes check box. 2. Click OK .	Set the retain attribute: 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 No check box. 2. Click OK .	Set the no-readvertise attribute: set no-readvertise
Specify that the static route is to be included in the routing table whether the route is active or not. By default, passive routes are not included in the routing table.	1. From the Passive flag list, select Passive . 2. Click OK .	Set the passive attribute: set passive

Verifying the Static Route Configuration

Verify that the static routes are in the routing table and that those routes are active.

Displaying the Routing Table

Purpose Verify static route configuration as follows by displaying the routing table and checking its contents.

Action From the CLI, enter the show route terse command.

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 235
- Before You Begin on page 236
- Configuring a RIP Network with Quick Configuration on page 236
- Configuring a RIP Network with a Configuration Editor on page 239
- Verifying the RIP Configuration on page 247

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 101.

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-Enabled Interfaces

RIP 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 60.
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 247.

Table 60: RIP Routing Quick Configuration Summary

Field	Function	Your Action
RIP		
Enable RIP	Enables or disables RIP.	<ul style="list-style-type: none"> ■ To enable RIP, select the check box. ■ To disable RIP, clear the check box.
Advertise Default Route	Advertises the default route using RIPv2.	<ul style="list-style-type: none"> ■ To advertise the default route using RIPv2, select the check box. ■ To disable the default route advertisement, clear the check box.
RIP-Enabled Interfaces	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"> ■ To enable RIP on an interface, click the interface name to highlight it, and click the left arrow to add the interface to the RIP interfaces list. ■ To enable RIP on multiple interfaces at once, press Ctrl while you click multiple interface names to highlight them. Then click the left arrow to add the interfaces to the RIP interfaces list. ■ To enable RIP on all logical interfaces except the special fxp0 management interface, select All Interfaces in the Logical Interfaces list and click the left arrow. ■ To enable RIP on all the interfaces displayed in the Logical Interfaces list, click All to highlight every interface. Then click the left arrow to add the interfaces to the RIP interfaces list. ■ To disable RIP on one or more interfaces, highlight the interface or interfaces in the RIP interfaces box and click the right arrow to move them back to the Logical Interfaces list.

Configuring a RIP Network with a Configuration Editor

To configure the Services Router as a node in a RIP network, you must perform the following task marked *(Required)*.


- Configuring a Basic RIP Network (Required) on page 239
- Controlling Traffic in a RIP Network (Optional) on page 242
- Enabling Authentication for RIP Exchanges (Optional) on page 245

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

Configuring a Basic RIP Network (Required)

To use RIP on the Services Router, you must configure RIP on all the RIP interfaces within a network like the one shown in Figure 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 61.
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 242.
 - To authenticate RIP exchanges, see “Enabling Authentication for RIP Exchanges (Optional)” on page 245.
 - To check the configuration, see “Verifying the RIP Configuration” on page 247.

Table 61: Configuring a RIP Network

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

Table 61: Configuring a RIP Network (continued)

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

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 242
- Controlling Traffic with the Outgoing Metric on page 243

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 62.
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 245.
 - To check the configuration, see “Verifying the RIP Configuration” on page 247.

Table 62: Modifying the Incoming Metric

Task	J-Web Configuration Editor	CLI Configuration Editor
In the configuration hierarchy, navigate to the level of an interface in the alpha1 RIP group.	<ol style="list-style-type: none">1. In the configuration editor hierarchy, select Protocols > Rip, and click alpha1 in the Group name field.2. Click the interface name—for example, fe-0/0/0.0—in the Neighbor name field.	From the top of the configuration hierarchy, enter edit protocols rip group alpha1 neighbor fe-0/0/0
Increase the incoming metric to 3 .	In the Metric in box, type 3 , and click OK .	Set the incoming metric to 3 : set metric-in 3

Controlling Traffic with the Outgoing Metric

If an exported route was learned from a member of the same RIP group, the metric associated with that route is the normal RIP metric. For example, a RIP route with a metric of 5 learned from a neighbor configured with an incoming metric of 2 is advertised with a combined metric of 7 when advertised to neighbors in the same group. However, if this route was learned from a RIP neighbor in a different group or from a different protocol, the route is advertised with the metric value configured in the outgoing metric for that group. Figure 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 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 245.
 - To check the configuration, see “Verifying the RIP Configuration” on page 247.

Table 63: Modifying the Outgoing Metric

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the alpha1 level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Rip , and then click alpha1 in the Group name field.	From the top of the configuration hierarchy, enter edit protocols rip group alpha1
Increase the outgoing metric to 3.	In the Metric out box, type 3 , and click OK .	Set the outgoing metric to 3: set metric-out 3

Enabling Authentication for RIP Exchanges (Optional)

All RIPv2 protocol exchanges can be authenticated to guarantee that only trusted routers participate in the AS's routing. By default, this authentication is disabled. Authentication requires all routers within the RIP network or subnetwork to have the same authentication type and key (password) configured.

You can enable RIP authentication exchanges by either of the following methods:

- Enabling Authentication with Plain-Text Passwords on page 245
- Enabling Authentication with MD5 Authentication on page 246

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 64.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying the RIP Configuration” on page 247.

Table 64: Configuring Simple RIP Authentication

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to Rip level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Rip .	From the top of the configuration hierarchy, enter edit protocols rip

Table 64: Configuring Simple RIP Authentication (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Set the authentication type to simple .	From the Authentication type list, select simple .	Set the authentication type to simple : set authentication-type simple
Set the authentication key to a simple-text password. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.	In the Authentication key box, type a simple-text password, and click OK .	Set the authentication key to a simple-text password: set authentication-key <i>password</i>

Enabling Authentication with MD5 Authentication

To configure authentication that requires an MD5 password to be included in the transmitted packet, enable MD5 authentication by performing these steps on all RIP Services Routers in the network:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 65.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying the RIP Configuration” on page 247.

Table 65: Configuring MD5 RIP Authentication

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to Rip level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Rip .	From the top of the configuration hierarchy, enter edit protocols rip
Set the authentication type to MD5 .	From the Authentication type list, select md5 .	Set the authentication type to md5 : set authentication-type md5
Set the MD5 authentication key (password). The key can be from 1 through 16 contiguous characters long and can include any ASCII strings.	In the Authentication key box, type an MD5 authentication key, and click OK .	Set the MD5 authentication key: set authentication-key <i>password</i>

Verifying the RIP Configuration

To verify the RIP configuration, perform these tasks:

- Verifying the RIP-Enabled Interfaces on page 247
- Verifying the Exchange of RIP Messages on page 247
- Verifying Reachability of All Hosts in the RIP Network on page 248

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 routera-fxp0.englab.mycompany.net (192.168.71.246) 0.251 ms 0.235 ms 0.200 ms
```

What It Means

Each numbered row in the output indicates a router (“hop”) in the path to the host. The three time increments indicate the round-trip time (RTT) between the Services Router and the hop, for each traceroute packet.

To ensure that the RIP network is healthy, verify the following information:

- The final hop in the list is the host you want to reach.
- The number of expected hops to the host matches the number of hops in the traceroute output. The appearance of more hops than expected in the output indicates that a network segment is probably unreachable. It might also indicate that the incoming or outgoing metric on one or more hosts has been set unexpectedly.

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 251
- Before You Begin on page 253
- Configuring an OSPF Network with Quick Configuration on page 253
- Configuring an OSPF Network with a Configuration Editor on page 255
- Tuning an OSPF Network for Efficient Operation on page 263
- Verifying an OSPF Configuration on page 268

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 159. For information about configuring OSPF demand circuits, see “Configuring Dial-on-Demand Routing with OSPF Support (Optional)” on page 180.

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 101.

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 66.
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 268.

Table 66: OSPF Routing Quick Configuration Summary

Field	Function	Your Action
Router Identification		
Router Identifier (required)	Uniquely identifies the router.	Type the Services Router’s 32-bit IP address, in dotted decimal notation.
OSPF		
Enable OSPF	Enables or disables OSPF.	<ul style="list-style-type: none"> ■ To enable OSPF, select the check box. ■ To disable OSPF, clear the check box.
OSPF Area ID	Uniquely identifies the area within its AS.	Type a 32-bit numeric identifier for the area, or an integer. 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 66: 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"> ■ regular—A regular OSPF area, including the backbone area ■ stub—A stub area ■ nssa—A not-so-stubby area (NSSA)
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"> ■ To enable OSPF on an interface, click the interface name to highlight it, and click the left arrow to add the interface to the OSPF interfaces list. ■ To enable OSPF on multiple interfaces at once, press Ctrl while you click multiple interface names to highlight them. Then click the left arrow to add the interfaces to the OSPF interfaces list. ■ To enable OSPF on all logical interfaces except the special fxp0 management interface, select All Interfaces in the Logical Interfaces list and click the left arrow. ■ To enable OSPF on all the interfaces displayed in the Logical Interfaces list, click All to highlight every interface. Then click the left arrow to add the interfaces to the OSPF interfaces list. ■ To disable OSPF on one or more interfaces, highlight the interface or interfaces in the OSPF interfaces box and click the right arrow to move them back to the Logical Interfaces list.

Configuring an OSPF Network with a Configuration Editor

To configure the Services Router as a node in an OSPF network, you must perform the following tasks marked *(Required)*.

- Configuring the Router Identifier (Required) on page 256
- Configuring a Single-Area OSPF Network (Required) on page 256
- Configuring a Multiarea OSPF Network (Optional) on page 258
- Configuring Stub and Not-So-Stubby Areas (Optional) on page 261

To configure an OSPF demand circuit, see “Configuring Dial-on-Demand Routing with OSPF Support (Optional)” on page 180. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 159.)

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

Configuring the Router Identifier (Required)

The router identifier is the IP address that uniquely identifies the J-series Services Router.

OSPF uses the router identifier to elect a designated router, unless you manually specify a priority value. When the OSPF network first becomes active, by default, the router with the highest router identifier is elected the designated router.

To configure the router identifier for the Services Router:


1. Navigate to the top of the configuration hierarchy in either the J-Web or the CLI configuration editor.
2. Perform the configuration tasks described in Table 67.
3. Go on to “Configuring a Single-Area OSPF Network (Required)” on page 256.

Table 67: Configuring the Router Identifier

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
Enter the router ID value.	In the Router Id box, type the IP address of the Services Router, in dotted decimal notation.	Set the router-id value to the IP address of the Services Router, in dotted decimal notation. For example: set router-id 177.162.4.24
Apply your configuration changes.	Click OK to apply your entries to the configuration.	Changes in the CLI are applied automatically when you execute the set command.

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 68.
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 258.
 - To control external route advertisement in the AS, see “Configuring Stub and Not-So-Stubby Areas (Optional)” on page 261.
 - To configure an OSPF demand circuit, see “Configuring Dial-on-Demand Routing with OSPF Support (Optional)” on page 180. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 159.)
 - To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 263.
 - To check the configuration, see “Verifying an OSPF Configuration” on page 268.

Table 68: Configuring a Single-Area OSPF Network

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

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 259
- Creating Additional OSPF Areas on page 259
- Configuring Area Border Routers on page 260

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 256.

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 69.
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 260.
 - To control external route advertisement in the AS, see “Configuring Stub and Not-So-Stubby Areas (Optional)” on page 261.
 - To configure an OSPF demand circuit, see “Configuring Dial-on-Demand Routing with OSPF Support (Optional)” on page 180. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 159.)
 - To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 263.
 - To check the configuration, see “Verifying an OSPF Configuration” on page 268.

Table 69: Configuring a Multiarea OSPF Network

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Ospf level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Ospf .	From the top of the configuration hierarchy, enter edit protocols ospf
Create the additional area with a unique area ID, in dotted decimal notation.	<ol style="list-style-type: none"> 1. In the Area box, click Add new entry. 2. In the Area ID box, type 0.0.0.2. 	<ol style="list-style-type: none"> 1. Set the area ID to 0.0.0.2 and add an interface. For example: set area 0.0.0.2 interface fe-0/0/0
Add interfaces as needed to the OSPF area.	<ol style="list-style-type: none"> 1. In the Interface box, click Add new entry. 2. In the Interface name box, type the name of an interface on the Services Router and click OK. 3. Repeat Step 1 and Step 2 for each interface on this Services Router that you are adding to the backbone area. Only one interface is required. 	<ol style="list-style-type: none"> 2. Repeat Step 1 for each interface on this Services Router that you are adding to the area. Only one interface is required. <p>Changes in the CLI are applied automatically when you execute the set command.</p>

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 70.
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 261.
 - To configure an OSPF demand circuit, see “Configuring Dial-on-Demand Routing with OSPF Support (Optional)” on page 180. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 159.)

- To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 263.
- To check the configuration, see “Verifying an OSPF Configuration” on page 268.

Table 70: Configuring Area Border Routers


Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Ospf level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Ospf .	From the top of the configuration hierarchy, enter <code>edit protocols ospf</code>
Verify that the backbone area has at least one interface enabled for OSPF.	<p>Click 0.0.0.0 to display the Area ID 0.0.0.0 page, and verify that the backbone area has at least one interface enabled for OSPF.</p> <p>For example, Services Router B in Figure 60 has the following interfaces enabled for OSPF in the backbone area:</p> <ul style="list-style-type: none"> ■ Interface fe-0/0/0.0 ■ Interface fe-0/0/1.0 <p>To enable an interface on the backbone area, see “Configuring a Single-Area OSPF Network (Required)” on page 256.</p>	<p>View the configuration using the show command:</p> <p>show</p> <p>For example, Services Router B in Figure 60 has the following interfaces enabled for OSPF in the backbone area:</p> <pre>area 0.0.0.0 { interface fe-0/0/0.0; interface fe-0/0/1.0; }</pre> <p>To enable an interface on the backbone area, see “Configuring a Single-Area OSPF Network (Required)” on page 256.</p>
Create the additional area with a unique area ID, in dotted decimal format.	<ol style="list-style-type: none"> 1. In the Area box, click Add new entry. 2. In the Area ID box, type 0.0.0.2. 	<ol style="list-style-type: none"> 1. Set the area ID to 0.0.0.2 and add an interface. For example: <code>set area 0.0.0.2 interface fe-0/0/0</code>
Add interfaces as needed to the OSPF area.	<ol style="list-style-type: none"> 1. In the Interface box, click Add new entry. 2. In the Interface name box, type the name of an interface on the Services Router and click OK. 3. Repeat Step 1 and Step 2 for each interface on this Services Router that you are adding to the backbone area. Only one interface is required. 	<ol style="list-style-type: none"> 2. Repeat Step 1 for each interface on this Services Router that you are adding to the area. Only one interface is required. <p>Changes in the CLI are applied automatically when you execute the set command.</p>

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 259.
2. Configure an area border router to bridge the areas.
For instructions, see “Configuring Area Border Routers” on page 260.
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 71.
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 180. (You must have already configured an ISDN interface as described in “Configuring ISDN” on page 159.)

- To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 263.
- To check the configuration, see “Verifying an OSPF Configuration” on page 268.

Table 71: Configuring Stub Area and Not-So-Stubby Area Routers

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

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 264
- Controlling the Cost of Individual Network Segments on page 264
- Enabling Authentication for OSPF Exchanges on page 265
- Controlling Designated Router Election on page 267

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 72.

Table 72: Controlling Route Selection in the Forwarding Table by Setting Preferences

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Ospf level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Ospf .	From the top of the configuration hierarchy, enter edit protocols ospf
Set the external and internal route preferences.	<ol style="list-style-type: none"> 1. In the External preference box, type an external preference value—for example, 7. 2. In the Preference box, type an internal preference value—for example, 130. 3. Click OK. 	<ol style="list-style-type: none"> 1. Set the internal preference. For example: set preference 7 2. Set the external preference. For example: set external-preference 130 <p>Changes in the CLI are applied automatically when you execute the set command.</p>

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 73.

Table 73: Controlling the Cost of Individual Network Segments by Modifying the Metric

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the fe-0/0/0.0 level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Ospf > Area id 0.0.0.0 > Interface name fe-0/0/0.0 .	From the top of the configuration hierarchy, enter edit protocols ospf area 0.0.0.0 interface fe-0/0/0.0
Set the interface metric and the external and route preference.	<ul style="list-style-type: none">1. In the Metric box, type an interface metric value—for example, 5.2. Click OK.	<ul style="list-style-type: none">1. Set the interface metric. For example: set metric 52. Set the external preference. For example: set external-preference 130 <p>Changes in the CLI are applied automatically when you execute the set command.</p>

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 74.

Table 74: Enabling OSPF Authentication

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the 0.0.0.0 level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Ospf > Area id 0.0.0.0 .	From the top of the configuration hierarchy, enter <code>edit protocols ospf area 0.0.0.0</code>
Set the authentication type.	<ol style="list-style-type: none"> 1. From the Authentication type list, select the type of authentication to enable on the stub area: simple md5 2. Click OK. 	Set the authentication type to either simple or md5 . For example: <code>set authentication-type md5</code> Changes in the CLI are applied automatically when you execute the set command.

Table 74: Enabling OSPF Authentication (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the <i>interface-name</i> level in the configuration hierarchy.	In the configuration editor hierarchy under Protocols > Ospf > Area > 0.0.0.0 > interface , click an interface name.	From the top of the configuration hierarchy, enter edit protocols ospf area 0.0.0.0 interface <i>interface-name</i>
Set the authentication password (key) and, if applicable, the key identifier.	<ol style="list-style-type: none"> In the Key name box, type a password: For simple authentication, type from 1 through 8 ASCII characters. For MD5 authentication, type from 1 through 16 ASCII characters. For MD5 authentication only, in the Key ID box, type any value between 0 (the default) and 255 to associate with the MD5 password. Click OK. Repeat Step 1 through Step 3 for each interface in the stub area for which you are enabling authentication. 	<ol style="list-style-type: none"> Set the authentication password: For simple authentication, type from 1 through 8 ASCII characters. For MD5 authentication, type from 1 through 16 ASCII characters. For MD5 authentication only, set the key identifier to associate with the MD5 password to any value between 0 (the default) and 255. For example: set authentication-key Chey3nne key-id 2 Changes in the CLI are applied automatically when you execute the command. Repeat Step 1 and Step 2 for each interface in the stub area for which you are enabling authentication.

Controlling Designated Router Election

At designated router election, the router priorities are evaluated first, and the router with the highest priority is elected designated router.

By default, routers have a priority of 128. A priority of 0 marks the router as ineligible to become the designated router. To configure a router so it is always the designated router, set its priority to 255.

To change the priority of a Services Router to control designated router election:

- 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 75.

Table 75: 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 fe-0/0/1 level in the configuration hierarchy.	In the configuration editor hierarchy, select Protocols > Ospf > area id 0.0.0.3 > Interface name fe-0/0/1 .	From the top of the configuration hierarchy, enter edit protocols ospf area 0.0.0.3 interface fe-0/0/1
Set the Services Router priority.	<ol style="list-style-type: none"> 1. In the Priority box, type a value between 0 and 255. The default value is 128. 2. Click OK. 	<p>Set the priority to a value between 0 and 255. The default value is 128. For example:</p> <p>set priority 200</p> <p>Changes in the CLI are applied automatically when you execute the set command.</p>

Verifying an OSPF Configuration

To verify an OSPF configuration, perform these tasks:

- Verifying OSPF-Enabled Interfaces on page 268
- Verifying OSPF Neighbors on page 269
- Verifying the Number of OSPF Routes on page 270
- Verifying Reachability of All Hosts in an OSPF Network on page 271

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.

Sample Output

```
user@host> show ospf route
```

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 270.

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

Chapter 10

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 273
- Before You Begin on page 275
- Configuring BGP Sessions with Quick Configuration on page 275
- Configuring BGP Sessions with a Configuration Editor on page 277
- Verifying a BGP Configuration on page 286

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 219

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 221

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 101.

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 76.
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 286.

Table 76: BGP Routing Quick Configuration Summary

Field	Function	Your Action
Router Identification		
Router Identifier (required)	Uniquely identifies the router	Type the Services Router's 32-bit IP address, in dotted decimal notation.
BGP		
Enable BGP	Enables or disables BGP.	<ul style="list-style-type: none"> ■ To enable BGP, select the check box. ■ To disable BGP, clear the check box.
Autonomous System Number	Sets the unique numeric identifier of the AS in which the services router is configured.	<p>Type the Services Router's 32-bit AS number, in dotted decimal notation.</p> <p>If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter 3, the value assigned to the AS is 0.0.0.3.</p>
Peer Autonomous System Number	Sets the unique numeric identifier of the AS in which the peer host resides.	<p>Type the peer host's 32-bit AS number, in dotted decimal notation.</p> <p>If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter 3, the value assigned to the AS is 0.0.0.3.</p>
Peer Address	Specifies the IP address of the peer host's interface to which the BGP session is being established.	Type the IP address of the peer host's adjacent interface, in dotted decimal notation.
Local Address	Specifies the IP address of the local host's interface from which the BGP session is being established.	Type the IP address of the local host's adjacent interface, in dotted decimal notation.

Configuring BGP Sessions with a Configuration Editor

To configure the Services Router as a node in a BGP network, you must perform the following tasks marked *(Required)*.


- Configuring a Point-to-Point Peering Session (Required) on page 277
- Configuring BGP Within a Network (Required) on page 280
- Configuring a Route Reflector (Optional) on page 281
- Configuring BGP Confederations (Optional) on page 284

For information about using the J-Web and CLI configuration editors, see “Using Services Router 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 77.
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 280.
 - To configure route reflector clusters, see “Configuring a Route Reflector (Optional)” on page 281.
 - To subdivide autonomous systems (ASs), see “Configuring BGP Confederations (Optional)” on page 284.
 - To check the configuration, see “Verifying a BGP Configuration” on page 286.

Table 77: Configuring BGP Peering Sessions

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
Set the network's AS number to 17 .	<ol style="list-style-type: none"> 1. In the AS Number box, enter 17. 2. Click OK. 	Set the AS number to 17 : set autonomous-system 17
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
Create the BGP group external-peers , and add the external neighbor addresses to the group.	<ol style="list-style-type: none"> 1. In the Group box, click Add new entry. 2. In the Group name box, type the name of the group of external BGP peers—external-peers in this case. 3. In the Neighbor box, click Add new entry. 4. In the Address box, type the IP address of an external BGP peer, in dotted decimal notation, and click OK. 5. Repeat Step 3 and Step 4 for each BGP neighbor within the external group that you are configuring. 	<ol style="list-style-type: none"> 1. Create the group external-peers, and add the address of an external neighbor: set group external-peers neighbor 10.10.10.10 2. Repeat Step 1 for each BGP neighbor within the external peer group that you are configuring.
At the group level, set the AS number for the group external-peers to 22 . Because three of the peers in this group (peers A, B, and C) reside in one AS, you can set their AS number as a group.	<ol style="list-style-type: none"> 1. In the Peer as box, type the number of the AS in which most peers in the external-peers group reside. 2. Click OK. 	From the [edit protocols bgp] hierarchy level: set group external-peers peer-as 22
At the individual neighbor level, set the AS number for peer D to 79 . Because peer D is a member of the group external-peers , it inherits the peer AS number configured at the group level. You must override this value at the individual neighbor level.	<ol style="list-style-type: none"> 1. Under Neighbor, in the Address column, click the IP address of peer D—10.21.7.2 in this case. 2. In the Peer as box, type the AS number of the peer. 3. Click OK. 	From the [edit protocols bgp group external-peers] hierarchy level: set neighbor 10.21.7.2 peer-as 79
Set the group type to external .	<ol style="list-style-type: none"> 1. From the Type list, select external. 2. Click OK. 	From the [edit protocols bgp group external-peers] hierarchy level: set type external

- To check the configuration, see “Verifying a BGP Configuration” on page 286.

Table 78: Configuring IBGP Peering Sessions

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
Create the BGP group internal-peers , and add the internal neighbor addresses to the group. You must configure a full IBGP mesh, which requires that each peer be configured with every other internal peer as a BGP neighbor.	<ol style="list-style-type: none"> 1. In the Group box, click Add new entry. 2. In the Group name box, type the name of the group of internal BGP peers—internal-peers in this case. 3. In the Neighbor box, click Add new entry. 4. In the Address box, type the IP address of an internal BGP peer, in dotted decimal notation. 5. Click OK. 6. Repeat Step 3 and Step 4 for each internal BGP peer within the network. 	<ol style="list-style-type: none"> 1. Create the group internal-peers, and add the address of an internal neighbor: set group internal-peers neighbor 192.168.6.4 2. Repeat Step 1 for each internal BGP neighbor within the network.
Set the group type to internal .	<ol style="list-style-type: none"> 1. From the Type list, select internal. 2. Click OK. 	From the [edit protocols bgp group internal-peers] hierarchy level: set type internal
Configure a routing policy to advertise BGP routes.	See “Injecting OSPF Routes into the BGP Routing Table (Optional)” on page 428.	

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 277.
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 79.
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 284.

- To check the configuration, see “Verifying a BGP Configuration” on page 286.

Table 79: Configuring a Route Reflector

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

Table 79: Configuring a Route Reflector (continued)


Task	J-Web Configuration Editor	CLI Configuration Editor
<p>On the other routers in the cluster, create the BGP group cluster-peers, and add the internal IP address of the route reflector.</p> <p>You do not need to include the neighbor addresses of the other internal peers, or configure the cluster identifier on these route reflector clients. They need only be configured as internal neighbors.</p> <p>NOTE: If the other routers in the network are Services Routers, follow the steps in this row. Otherwise, consult the router documentation for instructions.</p>	<p>On a client Services Router in the cluster:</p> <ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Protocols > Bgp. 2. In the Group box, click Add new entry. 3. In the Group name box, type the name of the group in which the BGP peer is configured—cluster-peers in this case. 4. In the Neighbor box, click Add new entry. 5. In the Address box, type the IP address of the route reflector, in dotted decimal notation—in this case, 192.168.40.4. 6. Click OK. 	<p>On a client Services Router in the cluster:</p> <ol style="list-style-type: none"> 1. From the top of the configuration hierarchy, enter edit protocols bgp 2. Create the group cluster-peers, and add only the route reflector address to the group: set group cluster-peers neighbor 192.168.40.4
Configure a routing policy to advertise BGP routes.	See “Injecting OSPF Routes into the BGP Routing Table (Optional)” on page 428.	

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 80.
- 3. If you are finished configuring the router, commit the configuration.
- 4. To check the configuration, see “Verifying a BGP Configuration” on page 286.

Table 80: Configuring BGP Confederations

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

Table 80: 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"> 1. In the Members field, click Add new entry. 2. In the Value box, enter the sub-ASs that are members of this confederation. Separate multiple sub-ASs with a space. 	Add members to the confederation: set 17 members 64512 64513
Using EBGp, configure the peering session between the confederations (from router A to router B in this example). When setting the peer AS number for these sessions, use the sub-AS number rather than the AS number.	See “Configuring a Point-to-Point Peering Session (Required)” on page 277.	
Using IBGP, configure internal sessions within a sub-AS. You can configure an IBGP full mesh, or you can configure a route reflector.	<ul style="list-style-type: none"> ■ To configure an IBGP full mesh, see “Configuring BGP Within a Network (Required)” on page 280. ■ To configure a route reflector, see “Configuring a Route Reflector (Optional)” on page 281. 	

Verifying a BGP Configuration

To verify a BGP configuration, perform these tasks:

- Verifying BGP Neighbors on page 286
- Verifying BGP Groups on page 287
- Verifying BGP Summary Information on page 288
- Verifying Reachability of All Peers in a BGP Network on page 289

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

```

What It Means	<p>The output shows a summary of BGP session information. Verify the following information:</p> <ul style="list-style-type: none"> ■ For Groups, the total number of configured groups is shown. ■ For Peers, the total number of BGP peers is shown. ■ For Down Peers, the total number of unestablished peers is 0. If this value is not zero, one or more peering sessions are not yet established. ■ Under Peer, the IP address for each configured peer is shown. ■ Under AS, the peer AS for each configured peer is correct. ■ Under Up/Dwn State, the BGP state reflects the number of paths received from the neighbor, the number of these paths that have been accepted, and the number of routes being damped (such as 0/0/0). If the field is Active, it indicates a problem in the establishment of the BGP session.
----------------------	---

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

Purpose	By using the ping tool on each peer address in the network, verify that all peers in the network are reachable from each Services Router.
Action	<p>For each Services Router in the BGP network:</p> <ol style="list-style-type: none"> 1. In the J-Web interface, select Diagnose > Ping Host. 2. In the Remote Host box, type the name of a host for which you want to verify reachability from the Services Router. 3. Click Start. Output appears on a separate page.
Sample Output	<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>
What It Means	<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 293
- Configuring Signaling Protocols for Traffic Engineering on page 309
- Configuring Virtual Private Networks on page 321
- Configuring CLNS VPNs on page 345
- Configuring IPSec for Secure Packet Exchange on page 357

Chapter 11

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 293
- MPLS Overview on page 295
- Signaling Protocols Overview on page 300
- VPN Overview on page 304

MPLS and VPN Terms

To understand MPLS and VPNs, become familiar with the terms defined in Table 81 .

Table 81: 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 81: MPLS and VPN Terms (continued)

Term	Definition
inbound router	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> .
label	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).
Label Distribution Protocol (LDP)	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.
label-switched path (LSP)	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.)
label-switching router (LSR)	Any Services Router that is part of an LSP.
Layer 2 circuit	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.
Layer 2 VPN	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.
Layer 3 VPN	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.
link coloring	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.
Multiprotocol Label Switching (MPLS)	Method for engineering network traffic patterns by assigning short labels to network packets that describe how to forward the packets through the network.
multiple push	Addition by a Services Router of up to three labels to a packet as it enters a Multiprotocol Label Switching (MPLS) domain.
outbound router	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> .
penultimate hop popping (PHP)	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.
penultimate router	Second-to-last Services Router in an LSP. The penultimate router is responsible for label popping when penultimate hop popping (PHP) is configured.
pop	Removal by a Services Router of the top label from a packet as it exits the Multiprotocol Label Switching (MPLS) domain.
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).

Table 81: MPLS and VPN Terms (continued)

Term	Definition
provider router	Services Router in the service provider's network that does not attach to a customer edge (CE) device.
push	Addition of a label or stack of labels by a Services Router to a packet as it enters a Multiprotocol Label Switching (MPLS) domain.
Resource Reservation Protocol (RSVP)	Resource reservation setup protocol that interacts with integrated services on the Internet.
route distinguisher	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.
routing instance	Collection of routing tables, their interfaces, and the routing protocol parameters that control the information they contain.
swap	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.
swap and push	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.
traffic engineering database (TED)	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.
transit router	Any label-switching router (LSR) between the inbound and outbound Services Router of a label-switched path (LSP).
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.
VPN routing and forwarding (VRF) instance	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 296
- Label-Switched Paths on page 296
- Label-Switching Routers on page 297
- Labels on page 298

- Label Operations on page 298
- Penultimate Hop Popping on page 299
- LSP Establishment on page 299

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 300
- Resource Reservation Protocol on page 300

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 301
- Bandwidth Reservation Requirement on page 301
- Explicit Route Objects on page 301
- Constrained Shortest Path First on page 303
- Link Coloring on page 303

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 304
- VPN Routing Requirements on page 305
- VPN Routing Information on page 306
- Types of VPNs on page 307

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 12

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 309
- Before You Begin on page 310
- Configuring LDP and RSVP with a Configuration Editor on page 311
- Verifying an MPLS Configuration on page 316

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 101.
- Configure an interior gateway protocol (IGP) across your network. See “Configuring an OSPF Network” on page 251 or “Configuring a RIP Network” on page 235. 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 311
- Configuring RSVP-Signaled LSPs on page 313

For information about using the J-Web and CLI configuration editors, see “Using Services Router 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 82.

3. If you are finished configuring the router, commit the configuration.
4. Go on to “Verifying an LDP-Signaled LSP” on page 316.

Table 82: Configuring an LDP-Signaled LSP


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

Table 82: Configuring an LDP-Signaled LSP (continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Create the LDP instance on each Services Router in the MPLS network.	<ol style="list-style-type: none"> 1. Navigate to the Protocols > Ldp level in the configuration hierarchy. 2. Next to Interface, click Add new entry. 3. In the Interface name box, type the name of a transit interface—for example, fe-0/0/0. 4. Click OK. 5. Repeat Steps 1 through 4 for each transit interface on the routers in the MPLS network. 	<ol style="list-style-type: none"> 1. From the top of the configuration hierarchy, enter edit protocols ldp 2. Enable LDP on a transit interface. For example: set interface fe-0/0/0 3. Repeat Steps 1 and 2 for each transit interface on the routers in the MPLS network.
Set the keepalive interval to 5 seconds. The keepalive interval specifies the number of seconds between the transmission of keepalive messages along the LDP link.	<ol style="list-style-type: none"> 1. In the Keepalive interval box, type 5. 2. Click OK. 3. Repeat Steps 1 and 2 for each router in the MPLS network. 	<p>On each router in the MPLS network, enter</p> <p>set keepalive-interval 5</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 83.
3. If you are finished configuring the router, commit the configuration.
4. Go on to “Verifying an RSVP-Signaled LSP” on page 319.

Table 83: Configuring an RSVP-Signaled LSP

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

Table 83: Configuring an RSVP-Signaled LSP (continued)

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

Table 83: Configuring an RSVP-Signaled LSP (continued)

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

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 316
- Verifying an RSVP-Signaled LSP on page 319

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 316
- “Verifying LDP Sessions” on page 317
- “Verifying the Presence of LDP-Signaled LSPs” on page 318
- “Verifying Traffic Forwarding over the LDP-Signaled LSP” on page 318

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.



|                      |                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Sample Output</b> | <pre> user@c1&gt; 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> |
| <b>What It Means</b> | <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 319
- “Verifying RSVP Sessions” on page 319
- “Verifying the Presence of RSVP-Signaled LSPs” on page 320

### Verifying RSVP Neighbors

|                      |                                                                                                                                                                                                                                                   |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b>       | 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.                                                                                           |
| <b>Action</b>        | From the CLI, enter the <code>show rsvp neighbor</code> command.                                                                                                                                                                                  |
| <b>Sample Output</b> | <pre> user@r1&gt; 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> |
| <b>What It Means</b> | The output shows the IP addresses of the neighboring routers. Verify that each neighboring RSVP router loopback address is listed.                                                                                                                |

### Verifying RSVP Sessions

|                |                                                                                                                                           |
|----------------|-------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | 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 13

# 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 321
- Before You Begin on page 324
- Configuring VPNs with a Configuration Editor on page 324
- Verifying a VPN Configuration on page 342

## 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:

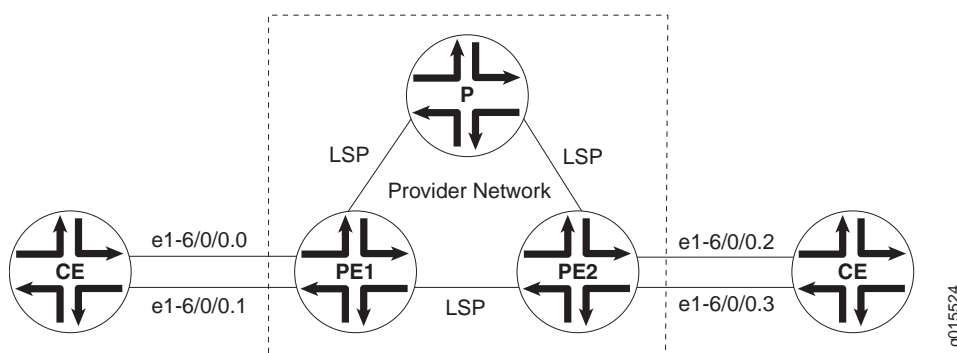
- Sample VPN Topology on page 322
- Basic Layer 2 VPN Configuration on page 322
- Basic Layer 2 Circuit Configuration on page 323

- Basic Layer 3 VPN Configuration on page 323

## 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 101.
- Determine the protocols to use in the VPN configuration. These protocols include
  - MPLS—See “Multiprotocol Label Switching Overview” on page 293 and the *JUNOS Routing Protocols Configuration Guide*.
  - BGP, EBGP, and internal BGP (IBGP)—See “Configuring BGP Sessions” on page 273 and the *JUNOS Routing Protocols Configuration Guide*.
  - LDP and Resource Reservation Protocol (RSVP)—See “Configuring Signaling Protocols for Traffic Engineering” on page 309 and the *JUNOS MPLS Applications Configuration Guide*.
  - OSPF—See “Configuring an OSPF Network” on page 251 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 84 to help you select the tasks for your VPN type. For information about using the J-Web and CLI configuration editors, see “Using Services Router Configuration Tools” on page 3.

- Configuring Interfaces Participating in a VPN on page 325
- Configuring Protocols Used by a VPN on page 327
- Configuring a VPN Routing Instance on page 335
- Configuring a VPN Routing Policy on page 337

**Table 84: VPN Configuration Task Summary**

| <b>Section</b>                                              | <b>Layer 3 VPN</b>                                                               | <b>Layer 2 VPN</b>                                      | <b>Layer 2 Circuit</b> |
|-------------------------------------------------------------|----------------------------------------------------------------------------------|---------------------------------------------------------|------------------------|
| “Configuring Interfaces Participating in a VPN” on page 325 | All Services Routers                                                             | All Services Routers                                    | All Services Routers   |
| “Configuring Protocols Used by a VPN” on page 327           | All Services Routers                                                             | All Services Routers                                    | All Services Routers   |
| “Configuring a VPN Routing Instance” on page 335            | PE Services Routers                                                              | PE Services Routers                                     | N/A                    |
| “Configuring a VPN Routing Policy” on page 337              | CE Services Routers<br>(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 101.

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 85 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 327.

**Table 85: Configuring an Interface for a VPN**

| Task                                                                                                                                                                                                                                                                                                       | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | CLI Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure IPv4.<br><br>(interfaces on all Services Routers)                                                                                                                                                                                                                                                | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Interfaces</b>.</li> <li>2. In the Interface name column, select the interface.</li> <li>3. For Layer 2 VPNs on the interface facing a CE router, select an encapsulation type, such as <b>ethernet-ccc</b> from the Encapsulation list. For Fast Ethernet interfaces, you also must select <b>Vlan tagging</b> from the Vlan tag mode list.</li> <li>4. In the Interface unit number column, select the logical interface.</li> <li>5. In the Family group, select <b>Inet</b> and click <b>Edit</b>.</li> <li>6. Next to Address, click <b>Add new entry</b></li> <li>7. In the Source box, type the IPv4 address—for example, <b>10.49.102.1/30</b>. For a loopback address on a Layer 2 configuration, select <b>Primary</b>.</li> <li>8. Click <b>OK</b> to return to the Unit page.</li> </ol> | <ul style="list-style-type: none"> <li>■ For all interfaces except loopback, and a Layer 2 VPN interface facing a CE router:<br/><br/>From the top of the configuration hierarchy, enter<br/><br/><code>edit interfaces interface-name unit logical_interface family inet address ipv4_address</code></li> <li>■ For a loopback address on a Layer 2 configuration:<br/><br/>From the top of the configuration hierarchy, enter<br/><br/><code>edit interfaces loO unit logical_interface family inet address ipv4_address primary</code></li> <li>■ For a Layer 2 VPN interface facing a CE router:<br/><br/>From the top of the configuration hierarchy, enter<br/><br/><code>set interfaces interface-name vlan-tagging encapsulation vlan-ccc unit logical_interface encapsulation vlan-ccc vlan-id id-number</code></li> </ul> |
| Configure the MPLS address family.<br><br>(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 <b>Mpls</b> in the Family group.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | At the [edit interfaces <i>interface</i> ] level, enter<br><br><code>set unit logical_interface family mpls</code>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| For Layer 2 VPNs and circuits, configure encapsulation.<br><br>If multiple logical units are configured, the encapsulation type is needed at the interface level only. It is always required at the unit level.<br><br>(for interfaces on a PE Services Router that communicate with a CE Services Router) | <ol style="list-style-type: none"> <li>1. On the Unit page, select an encapsulation type from the Encapsulation list.</li> <li>2. Click <b>OK</b>.</li> <li>3. On the Interface page, select an encapsulation type from the Encapsulation list.</li> <li>4. Click <b>OK</b> until you see the Configuration Interfaces page displaying all interfaces on the router.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <ol style="list-style-type: none"> <li>1. At the [edit interfaces <i>interface</i>] level, enter<br/><br/><code>set encapsulation encapsulation_type</code></li> <li>2. Enter<br/><br/><code>set unit logical_interface encapsulation encapsulation_type</code></li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |



## 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 86 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 327
- “Configuring a BGP Session” on page 329
- “Configuring Routing Options for VPNs” on page 330
- “Configuring an IGP and a Signaling Protocol” on page 331
- “Configuring LDP for Signaling” on page 331
- “Configuring RSVP for Signaling” on page 333
- “Configuring a Layer 2 Circuit” on page 334

**Table 86: VPN Protocol Configuration Task Summary**

| Section                                                                                                                                                                                                                                               | Layer 3 VPN                      | Layer 2 VPN                      | Layer 2 Circuit      |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|----------------------|
| “Configuring MPLS for VPNs” on page 327                                                                                                                                                                                                               | N/A unless you are using RSVP    | PE and provider Services Routers | PE Services Routers  |
| “Configuring a BGP Session” on page 329                                                                                                                                                                                                               | PE Services Routers              | PE Services Routers              | PE Services Routers  |
| “Configuring Routing Options for VPNs” on page 330                                                                                                                                                                                                    | All Services Routers             | All Services Routers             | All Services Routers |
| “Configuring an IGP and a Signaling Protocol” on page 331— <i>one</i> of the following tasks: <ul style="list-style-type: none"> <li>■ “Configuring LDP for Signaling” on page 331</li> <li>■ “Configuring RSVP for Signaling” on page 333</li> </ul> | PE and provider Services Routers | PE Services Routers              | PE Services Routers  |
| “Configuring a Layer 2 Circuit” on page 334                                                                                                                                                                                                           | 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 293 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 87 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 342.
5. Go on to “Configuring a BGP Session” on page 329.

**Table 87: 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.<br><br>(PE and provider Services Routers)                                                                                                                                                                                                                      | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Protocols &gt; Mpls &gt; Interface</b>.</li> <li>2. In the Interface name box, type <i>interface-name</i>.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                          | From the top of the configuration hierarchy, enter the following command for each interface you want to enable:<br><br><code>edit protocols mpls interface interface-name</code>                                                                                                                                                                 |
| 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.<br><br>The path name is defined on the source Services Router only and is unique between two routers.<br><br>(PE Services Router interface communicating with another PE Services Router) | <ol style="list-style-type: none"> <li>1. In the MPLS page, click <b>Add New Entry</b> in the Label switched path group.</li> <li>2. Type a path name in the Path name box and an IP address in the To box.</li> <li>3. Click <b>OK</b>.</li> <li>4. Next to Interface, click <b>Add New Entry</b>.</li> <li>5. Type <i>interface-name</i> in the Interface name box.</li> <li>6. Click <b>OK</b>.</li> <li>7. Repeat Steps 4 through 6 for each interface.</li> </ol> | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><br/> <code>edit protocols mpls label-switched-path path-name</code></li> <li>2. Enter<br/><br/> <code>set to ip-address</code></li> <li>3. Enter <code>up</code>.</li> <li>4. Enter<br/><br/> <code>interface interface-name</code></li> </ol> |

## 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 273 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 88 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 342.
5. Go on to “Configuring Routing Options for VPNs” on page 330.

**Table 88: Configuring an IBGP Session**

| <b>Task</b>                                                                                                    | <b>J-Web Configuration Editor</b>                                                                   | <b>CLI Configuration Editor</b>                                                                                      |
|----------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| Navigate to the top of the configuration hierarchy and configure the IBGP session.<br><br>(PE Services Router) | 1. In the configuration editor hierarchy, select <b>Protocols &gt; Bgp</b> .                        | 1. From the top of the configuration hierarchy, enter<br><br><code>edit protocols bgp group <i>group-name</i></code> |
|                                                                                                                | 2. Next to Group, click <b>Add New Entry</b> .                                                      |                                                                                                                      |
|                                                                                                                | 3. Type a name in the Group name box.                                                               | 2. Enter                                                                                                             |
|                                                                                                                | 4. From the Type list, select <b>Internal</b> .                                                     | <code>set type internal</code>                                                                                       |
|                                                                                                                | 5. In the Local address box, type the local loopback IP address.                                    | 3. Enter<br><br><code>set local-address <i>loopback-interface-ip-address</i></code>                                  |
|                                                                                                                | 6. In the Family group, select <b>L2vpn</b> for a Layer 2 VPN or <b>Inet vpn</b> for a Layer 3 VPN. | 4. Enter                                                                                                             |
|                                                                                                                | 7. Select <b>Unicast</b> .                                                                          | <code>set family <i>family-type</i> unicast</code>                                                                   |
|                                                                                                                | 8. Click <b>OK</b> .                                                                                | Replace <i>family-type</i> with <code>l2vpn</code> for a Layer 2 VPN or <code>inet-vpn</code> for a Layer 3 VPN.     |
|                                                                                                                | 9. In the Neighbor group, click <b>Add new entry</b> .                                              | 5. Enter <code>up</code> .                                                                                           |
|                                                                                                                | 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 <b>OK</b> until you return to the BGP page.                                               | <code>set neighbor <i>ip-address</i></code>                                                                          |

## 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 89.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 342.
5. Go on to “Configuring an IGP and a Signaling Protocol” on page 331.

**Table 89: Configuring Routing Options for a VPN**

| <b>Task</b>              | <b>J-Web Configuration Editor</b>                                                                                                                                                                            | <b>CLI Configuration Editor</b>                                                                                      |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------|
| Configure the AS number. | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, click <b>Routing Options</b>.</li> <li>2. In the AS number box, type the AS number.</li> <li>3. Click <b>OK</b>.</li> </ol> | <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 300.

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 251, “Configuring Static Routes” on page 223, and the *JUNOS Routing Protocols Configuration Guide*.

Configure the appropriate signaling protocol for your VPN:

- “Configuring LDP for Signaling” on page 331
- “Configuring RSVP for Signaling” on page 333

## 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 251.

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 90 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 325.

3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 342.
5. Go on to “Configuring a VPN Routing Instance” on page 335.

**Table 90: 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"> <li>1. In the configuration editor hierarchy, select <b>Protocols &gt; Ldp &gt; Interface</b>.</li> <li>2. In the Interface name column, type <i>interface-name</i>.</li> <li>3. Click <b>OK</b>.</li> <li>4. Repeat Steps 2 and 3 for each interface you want to enable.</li> </ol>                                                                                                                                                                                                                                                                                                                                                      | <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"> <li>1. In the configuration editor hierarchy, click <b>Protocols &gt; Ospf</b>.</li> <li>2. For Layer 2 VPN or circuit, select <b>Traffic engineering</b>.</li> <li>3. Next to Area group, click <b>Add new entry</b> and add the area.</li> <li>4. Next to Area group, select the area (0.0.0.0).</li> <li>5. Next to Interface group, select <b>Add new entry</b>.</li> <li>6. In the Interface name box, type <i>interface-name</i>.</li> <li>7. Click <b>OK</b>.</li> <li>8. Repeat Steps 5 through 7 to enable additional interfaces.</li> <li>9. Click <b>OK</b> twice to return to the Protocols page.</li> </ol> | <p>For OSPF:</p> <ol style="list-style-type: none"> <li>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> </li> <li>2. For Layer 2 VPN or circuit, move up to the [edit protocols ospf] level and enter <pre>set traffic-engineering</pre> </li> </ol> |

## 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 251.

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 91 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 342.
5. Go on to “Configuring a VPN Routing Instance” on page 335.

**Table 91: 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.<br>(PE Services Router)                                                                                                                    | For OSPF, follow these steps: <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Protocols &gt; Ospf</b>.</li> <li>2. Select <b>Traffic engineering</b>, and then click <b>Configure</b>.</li> <li>3. Select <b>Shortcuts</b>.</li> <li>4. Click <b>OK</b> until you return to the Protocols page.</li> </ol>                    | For OSPF, from the top of the configuration hierarchy, enter the following command for each interface you want to enable:<br><br>edit protocols ospf traffic-engineering shortcuts |
| Enable RSVP on interfaces that participate in the LSP.<br>(PE Services Router)<br>Enable interfaces on the source and destination points.<br>(provider Services Router)<br>Enable interfaces that connect the LSP between the PE Services Routers. | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Protocols &gt; Rsvp</b>.</li> <li>2. In the Interface group, click <b>Add New Entry</b>.</li> <li>3. Type an interface name.</li> <li>4. Click <b>OK</b>.</li> <li>5. Repeat Steps 2 through 4 for each interface you want to enable.</li> <li>6. Click <b>OK</b>.</li> </ol> | From the top of the configuration hierarchy, enter the following command for each interface you want to enable:<br><br>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 92 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 342.

**Table 92: 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.<br>(PE Services Router) | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Protocols &gt; L2circuit</b>.</li> <li>2. Next to Neighbor, click <b>Add new entry</b>.</li> <li>3. In the Neighbor box, enter the loopback address of the local router.</li> <li>4. Next to Interface, click <b>Add new entry</b>.</li> <li>5. In the <b>Interface</b> box, type the interface name of the remote PE router.</li> <li>6. In the Virtual circuit id box, type an ID number.</li> <li>7. Click <b>OK</b> until you return to the Protocols page.</li> </ol> | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><br/> <code>edit protocols l2circuit neighbor</code><br/> <code>interface-name interface interface-name</code><br/>           For <b>neighbor</b>, specify the local loopback address, and for <b>interface</b>, specify the interface name of the remote PE router.</li> <li>2. Enter<br/><br/> <code>set virtual-circuit-id id-number</code></li> </ol> |



## 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 93 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 342.
5. Go on to “Configuring a VPN Routing Policy” on page 337.

**Table 93: Configuring a VPN Routing Instance**

| <b>Task</b>                                                                                                                                                        | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                       | <b>CLI Configuration Editor</b>                                                                                                                                                                               |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the top of the configuration hierarchy and create the routing instance.<br><br>(PE Services Router)                                                    | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Routing instances &gt; Mpls</b>.</li> <li>2. In the Instance group, click <b>Add New Entry</b>.</li> <li>3. Type a name in the Instance name box.</li> </ol> | <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 <b>show route instance detail</b> command.<br><br>(PE Services Router) | In the Description box, type a description.                                                                                                                                                                                                             | <p>Enter</p> <p>set description "text"</p>                                                                                                                                                                    |
| Specify the instance type, either <b>l2vpn</b> for Layer 2 VPNs or <b>vrf</b> for Layer 3 VPNs.<br><br>(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.<br><br>(PE Services Router)                                                                                | <ol style="list-style-type: none"> <li>1. Next to Interface group, click <b>Add New Entry</b>.</li> <li>2. In the Interface name box, enter <i>interface-name</i>.</li> <li>3. Click <b>OK</b>.</li> </ol>                                              | <p>Enter</p> <p>set interface <i>interface-name</i></p>                                                                                                                                                       |
| Specify the route distinguisher.<br><br>(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"> <li>■ set route-distinguisher <i>as-number: number</i></li> <li>■ set route-distinguisher <i>ip-address: number</i></li> </ul> |

**Table 93: 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 338. | <ol style="list-style-type: none"> <li>Next to Vrf export group, select <b>Add new entry</b>.</li> <li>In the Value box, type the export routing policy name.</li> <li>Click <b>OK</b>.</li> <li>Next to Vrf import group, click <b>Add new entry</b>.</li> <li>In the Value box, type the import routing policy name.</li> <li>Click <b>OK</b>.</li> </ol> | <pre>set vrf-import import-policy-name vrf-export export-policy-name</pre>                                                                                                                                                        |
| (PE Services Router)                                                                                                          |                                                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                   |
| 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 341.    | <ol style="list-style-type: none"> <li>In the Vrf target box, click <b>Configure</b>.</li> <li>In the Community box, type the community (<b>target: community-id</b>, where <b>community-id</b> is <b>as-number: number</b> or <b>ip-address: number</b>).</li> <li>Click <b>OK</b>.</li> </ol>                                                             | <pre>set vrf-target target: community-id</pre> <p>Replace <b>community-id</b> with either of the following:</p> <ul style="list-style-type: none"> <li>■ <b>as-number: number</b></li> <li>■ <b>ip-address: number</b></li> </ul> |
| (PE Services Router)                                                                                                          |                                                                                                                                                                                                                                                                                                                                                             |                                                                                                                                                                                                                                   |

## 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 423 and the *JUNOS Routing Protocols Configuration Guide*.

- “Configuring a Routing Policy for Layer 2 VPNs” on page 338
- “Configuring a Routing Policy for Layer 3 VPNs” on page 341

## 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 94 and 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 342.

**Table 94: 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.<br><br>(PE Services Router) | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Policy options &gt; Policy statement</b>.</li> <li>2. In the Policy name box, type the policy name—for example, <code>import_vpn</code>.</li> </ol> | <p>From the top of the configuration hierarchy, enter</p> <pre>edit policy-options policy-statement import-policy-name</pre> |

**Table 94: Configuring an Import Routing Policy for Layer 2 VPNs (continued)**

| Task                                                           | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | CLI Configuration Editor                                                                                                                                                                                                            |
|----------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Define the term for accepting packets.<br>(PE Services Router) | <ol style="list-style-type: none"> <li>Next to Term group, click <b>Add new entry</b>.</li> <li>In the Term name box, type a term name—for example, <b>10</b>.</li> <li>Next to From, click <b>Configure</b>.</li> <li>Click <b>Add new entry</b>.</li> <li>Click <b>Protocol</b> and select <b>bgp</b> from the Value menu.</li> <li>Click <b>OK</b>.</li> <li>Next to Community, click <b>Add new entry</b>.</li> <li>Type the <i>community-name</i> in the Community Name box.</li> <li>Click <b>OK</b>.</li> <li>Next to Then, click <b>Configure</b>.</li> <li>From the Accept reject list, select <b>accept</b>.</li> <li>Click <b>OK</b> until you are at the Policy statement page.</li> </ol> | <ol style="list-style-type: none"> <li>Enter<br/><br/> <code>set term term-name-accept from protocol<br/>bgp community community-name</code></li> <li>Enter<br/><br/> <code>set term term-name-accept then accept</code></li> </ol> |
| Define the term for rejecting packets.<br>(PE Services Router) | <ol style="list-style-type: none"> <li>Next to the Term group, click <b>Add new entry</b>.</li> <li>In the Term name box, type a term name—for example, <b>20</b>.</li> <li>Next to Then, click <b>Configure</b>.</li> <li>From the Accept list, select <b>reject</b>.</li> <li>Click <b>OK</b> until you return to the Policy options page.</li> </ol>                                                                                                                                                                                                                                                                                                                                                | <ol style="list-style-type: none"> <li>Enter<br/><br/> <code>set term term-name-reject then reject</code></li> </ol>                                                                                                                |

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 95: Configuring an Export Routing Policy for Layer 2 VPNs**

| Task                                                               | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | CLI Configuration Editor                                                                                                                                                                         |
|--------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the export routing policy.<br><br>(PE Services Router)   | <ol style="list-style-type: none"> <li>Next to the Policy statement group, click <b>Add new entry</b>.</li> <li>In the Policy name box, type the policy name—for example, <code>export_vpn</code>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                 | <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.<br><br>(PE Services Router) | <ol style="list-style-type: none"> <li>Next to the Term group, click <b>Add new entry</b>.</li> <li>In the Term name box, type a term name—for example, <code>10</code>.</li> <li>Next to From, click <b>Configure</b>.</li> <li>Next to Community, click <b>Add new entry</b>.</li> <li>Type the <i>community-name</i> in the Community Name box.</li> <li>Click <b>OK</b>.</li> <li>Next to Then, click <b>Configure</b>.</li> <li>From the Accept reject list, select <b>accept</b>.</li> <li>Click <b>OK</b> twice until you are at the Policy statement page.</li> </ol> | <ol style="list-style-type: none"> <li>Enter <pre>set term term-name-accept from community add community-name</pre> </li> <li>Enter <pre>set term term-name-accept then accept</pre> </li> </ol> |
| Define the term for rejecting packets.<br><br>(PE Services Router) | <ol style="list-style-type: none"> <li>Next to the Term group, click <b>Add new entry</b>.</li> <li>In the Term name box, type a term name—for example, <code>20</code>.</li> <li>Next to Then, click <b>Configure</b>.</li> <li>From the Accept reject list, select <b>reject</b>.</li> <li>Click <b>OK</b> until you return to the Policy options page.</li> </ol>                                                                                                                                                                                                          | <ol style="list-style-type: none"> <li>Enter <pre>set term term-name-reject from community add community-name</pre> </li> <li>Enter <pre>set term term-name-reject then reject</pre> </li> </ol> |
| Define the community.<br><br>(PE Services Router)                  | <ol style="list-style-type: none"> <li>In the Community group, click <b>Add new entry</b>.</li> <li>In the Community name box, type a community name—for example, <code>VPN</code>.</li> <li>In the Members group, click <b>Add new entry</b>.</li> <li>In the Value box, type <code>target: community-id</code>, where <i>community-id</i> is <code>as-number : number</code> or <code>ip-address : number</code>.</li> <li>Click <b>OK</b> until you return to the Policy options page.</li> </ol>                                                                          | <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 96 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 342.

**Table 96: 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.<br><br>(CE Services Router) | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Policy options &gt; Policy statement</b>.</li> <li>2. In the Policy name box, type the policy name—for example, <code>loopback</code>.</li> </ol> | <p>From the top of the configuration hierarchy, enter</p> <pre>edit policy-options policy-statement policy-name</pre> |

**Table 96: Configuring a Routing Policy for Layer 3 VPNs (continued)**

| Task                                                               | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | CLI Configuration Editor                                                                                                                                                                                                                                                                                             |
|--------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Define the term for accepting packets.<br><br>(CE Services Router) | <ol style="list-style-type: none"> <li>1. In the Term group, click <b>Add new entry</b>.</li> <li>2. In the Term name box, type a term name—for example, <b>1</b>.</li> <li>3. Next to From, click <b>Configure</b>.</li> <li>4. Click <b>protocol</b>, then <b>Add new entry</b>.</li> <li>5. Select <b>direct</b> from the Value menu, and click <b>OK</b>.</li> <li>6.</li> <li>7. Next to Route Filter, click <b>Add new entry</b>.</li> <li>8. Type <i>local-loopback-address/netmask</i> in the Address box.</li> <li>9. Select <b>exact</b> from the Modifier list.</li> <li>10. Click <b>OK</b> twice.</li> <li>11. Next to Then, click <b>Configure</b>.</li> <li>12. From the Accept reject list, select <b>accept</b>.</li> <li>13. Click <b>OK</b> until you are at the Policy statement page.</li> </ol> | <ol style="list-style-type: none"> <li>1. Enter<br/><br/> <code>set term term-name-accept</code><br/> <code>from protocol direct route-filter</code><br/> <code>local-loopback-address/netmask exact</code> </li> <li>2. Enter<br/><br/> <code>set term term-name-accept</code> then <code>accept</code> </li> </ol> |
| Define the term for rejecting packets.<br><br>(CE Services Router) | <ol style="list-style-type: none"> <li>1. Next to the Term group, click <b>Add new entry</b>.</li> <li>2. In the Term name box, type a term name—for example, <b>2</b>.</li> <li>3. Next to Then, click <b>Configure</b>.</li> <li>4. From the Accept reject list, select <b>reject</b>.</li> <li>5. Click <b>OK</b> until you return to the Policy options page.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                          | <ol style="list-style-type: none"> <li>Enter<br/><br/> <code>set term term-name-reject</code> then <code>reject</code> </li> </ol>                                                                                                                                                                                   |

## 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 343



- Pinging a Layer 3 VPN on page 343
- Pinging a Layer 2 Circuit on page 343

### **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 14

# 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 345
- CLNS Overview on page 346
- Before You Begin on page 347
- Configuring CLNS with a Configuration Editor on page 347
- Verifying CLNS VPN Configuration on page 354

## CLNS Terms

Before configuring CLNS, become familiar with the terms defined in Table 97.

**Table 97: 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 97: CLNS Terms (continued)**

| <b>Term</b>                                                 | <b>Definition</b>                                                                                                                                                                                                                                                                                                                           |
|-------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>customer edge (CE) router</b>                            | 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.                                                                                                                                                                                           |
| <b>end system</b>                                           | A host in an Open Systems Interconnection (OSI) network.                                                                                                                                                                                                                                                                                    |
| <b>End System-to-Intermediate System (ES-IS)</b>            | 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.                                                                                                                                                 |
| <b>intermediate system</b>                                  | A router in an Open Systems Interconnection (OSI) network.                                                                                                                                                                                                                                                                                  |
| <b>International Organization for Standardization (ISO)</b> | Worldwide federation of standards bodies that promotes international standardization and published international agreements as International Standards.                                                                                                                                                                                     |
| <b>network layer reachability information (NLRI)</b>        | 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.                                                                                                                                                    |
| <b>network services access point (NSAP)</b>                 | 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. |
| <b>Open Systems Interconnection (OSI)</b>                   | Standard reference model for representing the way messages are transmitted between two points on a network.                                                                                                                                                                                                                                 |
| <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).                                                                                                                                                                                       |
| <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.                                                                                                                                                                   |

## 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 101.
- If applicable, configure BGP and VPNs. See “Configuring BGP Sessions” on page 273 and “Configuring Virtual Private Networks” on page 321.

## 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 348
- Configuring ES-IS on page 349
- Configuring IS-IS for CLNS on page 350
- Configuring CLNS Static Routes on page 352
- Configuring BGP for CLNS on page 353



**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 335.

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 98.
3. Go on to one of the following tasks:
  - “Configuring IS-IS for CLNS” on page 350
  - “Configuring CLNS Static Routes” on page 352
  - “Configuring BGP for CLNS” on page 353
  - “Verifying CLNS VPN Configuration” on page 354

**Table 98: 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 <b>aaaa</b> . | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Routing instances</b>.</li> <li>2. Next to Instance, click <b>Add new entry</b>.</li> <li>3. In the Instance name box, type <b>aaaa</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> | From the top of the configuration hierarchy, enter<br>edit routing-instances aaaa |
| Specify the instance type <b>vrf</b> for Layer 3 VPNs.                                           | In the Instance type list, select <b>vrf</b> .                                                                                                                                                                                                                              | Enter<br>set instance-type vrf                                                    |

**Table 98: Configuring a VPN Routing Instance for CLNS (continued)**

| <b>Task</b>                                                                                                                                     | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>CLI Configuration Editor</b>                                                                                                                                                      |
|-------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Specify the interfaces that belong to the routing instance <b>aaaa</b> —for example, <b>lo0.1</b> , <b>e1-2/0/0.0</b> , and <b>t1-3/0/0.0</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>lo0.1</b>.</li> <li>3. Click <b>OK</b>.</li> <li>4. Next to Interface, click <b>Add New Entry</b>.</li> <li>5. In the Interface name box, type <b>e1-2/0/0.0</b>.</li> <li>6. Click <b>OK</b>.</li> <li>7. Next to Interface, click <b>Add New Entry</b>.</li> <li>8. In the Interface name box, type <b>t1-3/0/0.0</b>.</li> <li>9. Click <b>OK</b>.</li> </ol> | <p>Enter</p> <ol style="list-style-type: none"> <li>1. <b>set interface lo0.1</b></li> <li>2. <b>set interface e1-2/0/0.0</b></li> <li>3. <b>set interface t1-3/0/0.0</b></li> </ol> |
| Specify the route distinguisher—for example, <b>10.255.245.1:1</b> .                                                                            | In the Rd type box, type <b>10.255.245.1:1</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <p>Enter</p> <p><b>set route-distinguisher 10.255.245.1:1</b></p>                                                                                                                    |
| Specify the policy for the Layer 3 VRF table—for example, <b>target:11111:1</b> .                                                               | <ol style="list-style-type: none"> <li>1. Next to Vrf target, click <b>Configure</b>.</li> <li>2. In the Community box, type <b>target:11111:1</b>.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                              | <p>Enter</p> <p><b>set vrf-target target:11111:1</b></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 99.
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 350

- “Configuring CLNS Static Routes” on page 352
- “Configuring BGP for CLNS” on page 353
- “Verifying CLNS VPN Configuration” on page 354

**Table 99: Configuring ES-IS**

| Task                                                                           | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                         | CLI Configuration Editor                                                                         |
|--------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|
| Navigate to the <b>Routing instances</b> level in the configuration hierarchy. | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Routing instances</b>.</li> <li>2. Under Instance name, click <b>aaaa</b>.</li> </ol>                                                                                                                                                                   | <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"> <li>1. Next to Protocols, click <b>Configure</b>.</li> <li>2. Next to Esis, click <b>Configure</b>.</li> <li>3. Next to Interface, click <b>Add new entry</b>.</li> <li>4. In the Interface name box, type <b>all</b>.</li> <li>5. Click <b>OK</b> until you return to the Protocols statement page.</li> </ol> | <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 423.

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 100.
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 352
  - “Configuring BGP for CLNS” on page 353
  - “Verifying CLNS VPN Configuration” on page 354



**Table 100: Configuring IS-IS to Exchange CLNS Routes**

| <b>Task</b>                                                                                  | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>CLI Configuration Editor</b>                                                                                                                |
|----------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Routing instances</b> level in the configuration hierarchy.               | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Routing instances</b>.</li> <li>2. Under Instance name, click <b>aaaa</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                  | <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"> <li>1. Next to Protocols, click <b>Configure</b>.</li> <li>2. Next to Isis, click <b>Configure</b>.</li> <li>3. Next to CLNS routing, select the <b>Yes</b> box.</li> </ol>                                                                                                                                                                                                                                                                                                                                    | <p>Enter</p> <pre>set protocols isis clns-routing</pre>                                                                                        |
| Enable IS-IS on all interfaces.                                                              | <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>all</b>.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                          | <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"> <li>1. Next to No ipv4 routing, select the <b>Yes</b> box.</li> <li>2. Next to No ipv6 routing, select the <b>Yes</b> box.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                             | <p>Enter</p> <pre>set protocols isis no-ipv4-routing no-ipv6-routing</pre>                                                                     |
| Define the BGP export policy name—for example, <b>dist-bgp</b> —and the family and protocol. | <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 <b>dist-bgp</b>.</li> <li>4. Next to From, click <b>Configure</b>.</li> <li>5. In the Family list, select <b>iso</b>.</li> <li>6. Next to Protocol, click <b>Add new entry</b>.</li> <li>7. In the Value list, select <b>bgp</b>.</li> <li>8. Click <b>OK</b> until you return to the Policy statement page.</li> </ol> | <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 100: Configuring IS-IS to Exchange CLNS Routes (continued)**

| <b>Task</b>                              | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                          | <b>CLI Configuration Editor</b>                                                                                                |
|------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------|
| Define the action for the export policy. | <ol style="list-style-type: none"> <li>1. Next to Then, click <b>Configure</b>.</li> <li>2. In the Accept reject list, select <b>accept</b>.</li> <li>3. Click <b>OK</b> until you return to the Configuration page.</li> </ol>                                                                                                                                                                            | <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"> <li>1. In the configuration editor hierarchy, select <b>Routing instances</b>.</li> <li>2. Next to aaaa, click <b>Protocols</b>.</li> <li>3. Next to Isis, click <b>Edit</b>.</li> <li>4. Next to Export, click <b>Add new entry</b>.</li> <li>5. In the Value box, type <b>dist-bgp</b>.</li> <li>6. Click <b>OK</b> until you return to the Instance page.</li> </ol> | <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.

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 101.

This procedure, as well as the configuration provided in “Verifying CLNS VPN Configuration” on page 354, 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

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 353
  - “Verifying CLNS VPN Configuration” on page 354

**Table 101: Configuring Static CLNS Routes**

| <b>Task</b>                                                                    | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | <b>CLI Configuration Editor</b>                                                                                                                        |
|--------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Routing instances</b> level in the configuration hierarchy. | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Routing instances</b>.</li> <li>2. Under Instance name, click <b>aaaa</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <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"> <li>1. Next to Routing options, click <b>Configure</b>.</li> <li>2. Next to Rib, click <b>Add new entry</b>.</li> <li>3. In the Rib name box, type <b>aaaa.iso.0</b>.</li> <li>4. Next to Static, click <b>Configure</b>.</li> <li>5. Next to Iso route, click <b>Add new entry</b>.</li> <li>6. In the Destination box, type <b>47.0005.80ff.f800.0000.bbbb.1022/104</b>.</li> <li>7. From the Next hop list, select <b>Next hop</b>.</li> <li>8. Next to Next hop, click <b>Add new entry</b>.</li> <li>9. In the Value box, type <b>47.0005.80ff.f800.0000.aaaa.1000.1921.6800.4196.00</b>.</li> <li>10. Click <b>OK</b>.</li> </ol> | <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 102.
3. If you are finished configuring the router, commit the configuration.
4. To verify the configuration, see “Verifying CLNS VPN Configuration” on page 354.

**Table 102: Configuring BGP to Carry CLNS VPN NLRI Messages**

| <b>Task</b>                                                                   | <b>J-Web Configuration Editor</b>                                                                                                                                                       | <b>CLI Configuration Editor</b>                                                               |
|-------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| 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                                            |
| Define a BGP group name—for example, pedge-pegde.                             | <ol style="list-style-type: none"> <li>Next to Group, click <b>Add new entry</b>.</li> <li>In the Group name box, type <b>pedge-pegde</b>.</li> </ol>                                   | <pre>set protocols bgp group pedge-pegde 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"> <li>Next to Neighbor, click <b>Add new entry</b>.</li> <li>In the Address box, type <b>10.255.245.215</b>.</li> </ol>                                |                                                                                               |
| Define the family.                                                            | <ol style="list-style-type: none"> <li>Under Family, next to Iso vpn, click <b>Configure</b>.</li> <li>Next to Unicast, select the <b>Yes</b> box.</li> <li>Click <b>OK</b>.</li> </ol> |                                                                                               |

## 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 15

# 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 357
- Before You Begin on page 358
- Configuring an IPsec Tunnel with Quick Configuration on page 358
- Configuring an IPsec Tunnel with a Configuration Editor on page 360
- Verifying the IPsec Tunnel Configuration on page 370

## 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 101.
- Configure one or more routing protocols. See “Configuring Static Routes” on page 223, “Configuring a RIP Network” on page 235, “Configuring an OSPF Network” on page 251, or “Configuring BGP Sessions” on page 273.

## 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**

|  |  | ? |
|--|--|---|

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 103.
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 370.

**Table 103: IPSec Tunnels Quick Configuration Summary**

| Field                             | Function                                                                                                                                                                                                                                  | Your Action                                                                                                                                                                                      |
|-----------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Tunnel Information</b>         |                                                                                                                                                                                                                                           |                                                                                                                                                                                                  |
| 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"> <li>1. In the text box at the bottom of the list, enter an IP address or address prefix, in dotted decimal notation.</li> <li>2. Click <b>Add</b>.</li> </ol> |

## 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 361
- Configuring IPSec Service Sets on page 362
- Configuring an IPSec Stateful Firewall Filter Rule on page 366
- Configuring a NAT Pool on page 368

## 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 104.
3. Go on to “Configuring IPSec Service Sets” on page 362.

**Table 104: Configuring IPSec Interfaces**

| 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 |

**Table 104: Configuring IPSec Interfaces (continued)**

| <b>Task</b>                                                                                                                                                                                                                                                      | <b>J-Web Configuration Editor</b>                                            | <b>CLI Configuration Editor</b>                                      |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------|----------------------------------------------------------------------|
| Configure the inside services interface for the IPSec tunnel.                                                                                                                                                                                                    | 1. In the Interface field, click <b>Add new entry</b> .                      | 1. Configure the services interface as an inside-service interface:  |
| On the J-series Services Router, the services interface is always <b>sp-0/0/0.unit</b> . The logical interface must have a unit number other than 0. By default, the J-Web Quick Configuration uses the unit number 1001 for inside-service logical interfaces.  | 2. In the Interface name field, type <b>sp-0/0/0</b> , and click <b>OK</b> . | <b>set sp-0/0/0 unit 1001 service-domain inside</b>                  |
|                                                                                                                                                                                                                                                                  | 3. In the Interface field, click <b>sp-0/0/0</b> .                           | 2. Configure the services interface as an <b>inet</b> interface:     |
|                                                                                                                                                                                                                                                                  | 4. In the Unit field, click <b>Add new entry</b> .                           | <b>set sp-0/0/0 unit 1001 family inet</b>                            |
|                                                                                                                                                                                                                                                                  | 5. In the Interface unit number field, type <b>1001</b> .                    |                                                                      |
|                                                                                                                                                                                                                                                                  | 6. In the Service domain box, select <b>inside</b> from the list.            |                                                                      |
|                                                                                                                                                                                                                                                                  | 7. In the Family field, click <b>inet</b> .                                  |                                                                      |
|                                                                                                                                                                                                                                                                  | 8. Select the <b>Primary</b> box, and click <b>OK</b> .                      |                                                                      |
| Configure the outside services interface for the IPSec tunnel.                                                                                                                                                                                                   | 1. In the Interface field, click <b>Add new entry</b> .                      | 1. Configure the services interface as an outside-service interface: |
| On the J-series Services Router, the services interface is always <b>sp-0/0/0.unit</b> . The logical interface must have a unit number other than 0. By default, the J-Web Quick Configuration uses the unit number 2001 for outside-service logical interfaces. | 2. In the Interface name field, type <b>sp-0/0/0</b> , and click <b>OK</b> . | <b>set sp-0/0/0 unit 2001 service-domain outside</b>                 |
|                                                                                                                                                                                                                                                                  | 3. In the Interface field, click <b>sp-0/0/0</b> .                           | 2. Configure the services interface as an <b>inet</b> interface:     |
|                                                                                                                                                                                                                                                                  | 4. In the Unit field, click <b>Add new entry</b> .                           | <b>set sp-0/0/0 unit 2001 family inet</b>                            |
|                                                                                                                                                                                                                                                                  | 5. In the Interface unit number field, type <b>2001</b> .                    |                                                                      |
|                                                                                                                                                                                                                                                                  | 6. In the Service domain box, select <b>outside</b> from the list.           |                                                                      |
|                                                                                                                                                                                                                                                                  | 7. In the Family field, click <b>inet</b> .                                  |                                                                      |
|                                                                                                                                                                                                                                                                  | 8. Select the <b>Primary</b> box, and click <b>OK</b> .                      |                                                                      |

## 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 105.
3. Go on to “Configuring an IPsec Stateful Firewall Filter Rule” on page 366.

**Table 105: Configuring IPsec Service Sets**

| Task                                                     | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | CLI Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the next-hop service set for the IPsec tunnel. | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, click <b>Services</b>.</li> <li>2. In the Service sets field, click <b>Add new entry</b>.</li> <li>3. In the Service set name field, type the name of the service set. The name can be any unique string.</li> <li>4. In the Service type choice field, select <b>Next hop service</b> from the list.</li> <li>5. In the Nested configuration field, click <b>Next hop service</b>.</li> <li>6. In the Inside service interface field, type the services interface, including unit number, for the inside-service interface—for example, <code>sp-0/0/0.1001</code>.</li> <li>7. Click <b>OK</b>.</li> <li>8. In the Nested configuration field, click <b>Next hop service</b>.</li> <li>9. In the Outside service interface field, type the services interface, including the unit number—for example, <code>sp-0/0/0.2002</code>.</li> <li>10. 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 services</code></li> <li>2. Set the inside-service interface:<br/><br/><code>set service-set service-set-name<br/>next-hop-service<br/>inside-service-interface<br/>sp-0/0/0.1001</code></li> <li>3. Set the outside-service interface:<br/><br/><code>set service-set service-set-name<br/>next-hop-service<br/>outside-service-interface<br/>sp-0/0/0.2001</code></li> </ol> |

**Table 105: Configuring IPSec Service Sets (continued)**

| <b>Task</b>                                                                                                                                                                                            | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                              |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the local gateway for the IPSec service set.                                                                                                                                                 | <ol style="list-style-type: none"> <li>1. In the Ipsec vpn options field, click <b>Configure</b>.</li> <li>2. In the Local gateway box, type the IP address of the local tunnel endpoint, in dotted decimal notation—for example, 1.1.1.1.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                           | <p>Set the local gateway address for the service set:</p> <pre>set service-set service-set-name ipsec-vpn-options local-gateway 1.1.1.1</pre>                                                                                                                                                |
| <p>Configure IPSec rules to set the remote gateway on all traffic to 2.2.2.2.</p> <p>Because the rule applies to all traffic, you must only configure the action (or then statement) for the term.</p> | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, click <b>Services &gt; Ipsec-vpn</b>.</li> <li>2. In the Rule field, click <b>Add new entry</b>.</li> <li>3. In the Rule name field, type the name of the rule. The rule name can be any unique string.</li> <li>4. In the term field, click <b>Add new entry</b>.</li> <li>5. In the Term name field, type the name of the term. It can be any unique string.</li> <li>6. To configure an action, click <b>Then</b>.</li> <li>7. In the Remote gateway field, type the remote gateway address, in dotted decimal notation—for example, 2.2.2.2.</li> <li>8. Click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/>edit services ipsec-vpn</li> <li>2. Configure a rule with a term that sets the remote gateway to 2.2.2.2:<br/><br/>set rule rule-name term term-name then remote-gateway 2.2.2.2</li> </ol> |

**Table 105: Configuring IPSec Service Sets (continued)**

| <b>Task</b>                                                                                                                                                                                                                                                                                                              | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>Configure an security association with a static IKE key.</p> <p>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.</p> <p>The IKE key is configured as <b>ike policy</b> and then applied using the <b>dynamic</b> statement.</p> | <ol style="list-style-type: none"> <li>From the top of the configuration hierarchy, select <b>Services &gt; Ipsec-vpn &gt; Ike</b>.</li> <li>In the Policy field, click <b>Add new entry</b>.</li> <li>In the Name box, type the name of the IKE policy. It can be any unique string.</li> <li>Click <b>Pre-shared key</b>.</li> <li>In the Key choice field, select <b>Ascii text</b> from the list.</li> <li>In the Ascii text box, enter the IKE key in plain text.</li> <li>Click <b>OK</b>.</li> <li>Navigate to the IPSec rule configured previously. From the top of the configuration hierarchy, click <b>Services &gt; Ipsec-vp &gt; rule-name &gt; term term-name &gt; then</b>.</li> <li>Click <b>Dynamic</b>.</li> <li>In the Ike-policy box, type the name of the IKE policy you configured.</li> <li>Click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>From the top of the configuration hierarchy, enter<br/><code>edit services ipsec-vpn ike</code></li> <li>Configure the IKE pre-shared key in ASCII text format:<br/><code>set policy policy-name pre-shared-key ascii-text ike-key</code></li> <li>Navigate to the IPSec rule configured previously. From the top of the configuration hierarchy, enter<br/><code>edit services ipsec-vpn rule-name term term-name</code> then.</li> <li>Configure a dynamic security association that applies the IKE policy:<br/><code>set dynamic ike-policy policy-name</code></li> </ol> |

**Table 105: Configuring IPSec Service Sets (continued)**

| <b>Task</b>                                                                        | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                          | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                      |
|------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the IPSec rule so that it acts on input traffic.                         | <ol style="list-style-type: none"> <li>From the top of the configuration hierarchy, click <b>Services &gt; Ipsec-vpn &gt; Rule &gt; rule-name</b>.</li> <li>In the Match direction field, select <b>Input</b> from the list.</li> <li>Click <b>OK</b>.</li> </ol>                                                                                                                                                                          | <ol style="list-style-type: none"> <li>From the top of the configuration hierarchy, enter<br/><code>edit services ipsec-vpn rule rule-name</code></li> <li>Set the match direction for the rule:<br/><code>set match-direction input</code></li> </ol>               |
| Apply the IPSec rule to all traffic through the previously configured service set. | <ol style="list-style-type: none"> <li>From the top of the configuration hierarchy, click <b>Services &gt; Service-set &gt; service-set-name</b>.</li> <li>In the Ipsec vpn rules choice field, select <b>Ipsec vpn rules</b> from the list.</li> <li>In the Ipsec vpn rules field, click <b>Add new entry</b>.</li> <li>In the Rule name box, type the name of the previously configured IPSec rule.</li> <li>Click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>From the top of the configuration hierarchy, enter<br/><code>edit services service-set service-set-name</code></li> <li>Apply the IPSec rule previously configured:<br/><code>set ipsec-vpn-rules rule-name</code></li> </ol> |

### 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 437.

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 106.
- Go on to “Configuring a NAT Pool” on page 368.



**Table 106: Configuring an IPSec Stateful Firewall Filter Rule**

| Task                                                               | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | CLI Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create the stateful firewall rule and apply it to inbound traffic. | <ol style="list-style-type: none"> <li>From the top of the configuration hierarchy, click <b>Services &gt; Stateful firewall</b>.</li> <li>In the rule field, click <b>Add new entry</b>.</li> <li>In the Rule name box, type the name of the rule. It can be any unique string.</li> <li>In the Match direction field, select <b>Input</b> from the list.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <ol style="list-style-type: none"> <li>From the top of the configuration hierarchy, enter<br/><code>edit services stateful-firewall</code></li> <li>Create the firewall rule and apply it to input traffic:<br/><br/><code>set rule <i>rule-name</i> match-direction input</code></li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Create the firewall term to match only desired traffic.            | <ol style="list-style-type: none"> <li>In the Term field, click <b>Add new entry</b>.</li> <li>In the Term name box, type the name of the term. It can be any unique string.</li> <li>Click <b>From</b>.</li> <li>In the Destination address field, click <b>Add new entry</b>.</li> <li>In the address field, select <b>Enter specific value</b> from the list.</li> <li>In the Address box, type the IP address of the local tunnel endpoint, in dotted decimal notation, and click <b>OK</b>.</li> <li>In the Source address field, click <b>Add new entry</b>.</li> <li>In the address field, select <b>Enter specific value</b> from the list.</li> <li>In the Address box, type the IP address of the remote tunnel endpoint, in dotted decimal notation, and click <b>OK</b>.</li> <li>In the Applications field, click <b>Add new entry</b>.</li> <li>In the Application name field, type <code>junos-ipsec-esp</code>, and click <b>OK</b>.</li> <li>In the Applications field, click <b>Add new entry</b>.</li> <li>In the Application name field, type <code>junos-ike</code>, and click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>Create the firewall term and match all packets with a destination address that matches the local tunnel endpoint:<br/><br/><code>set term <i>term-name</i> from destination-address <i>local-tunnel-end-point-address</i></code></li> <li>Match all packets with a source address that matches the remote tunnel endpoint:<br/><br/><code>set term <i>term-name</i> from source-address <i>remote-tunnel-end-point-address</i></code></li> <li>Match all packets using IPSec as an application protocol:<br/><br/><code>set term <i>term-name</i> from applications junos-ipsec-esp</code></li> <li>Match all packets using IKE as an application protocol:<br/><br/><code>set term <i>term-name</i> from applications junos-ike</code></li> </ol> |

**Table 106: Configuring an IPSec Stateful Firewall Filter Rule (continued)**

| <b>Task</b>                                                 | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                             |
|-------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the firewall term to accept only desired traffic. | <ol style="list-style-type: none"> <li>1. Click <b>OK</b> to return to the Term name page, and click <b>Then</b>.</li> <li>2. In the Designation field, select <b>Accept</b> from the list, select the <b>Yes</b> box.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                        | <p>Set the match action to accept:</p> <pre>set term <i>term-name</i> then accept</pre>                                                                                                                                                                                                     |
| Create the firewall term to reject all other traffic.       | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, click <b>Services &gt; Stateful firewall &gt; Rule &gt; <i>rule-name</i></b></li> <li>2. In the Term field, click <b>Add new entry</b>.</li> <li>3. In the Term name field, type the name of the term. The name can be any unique string.</li> <li>4. Click <b>Then</b>.</li> <li>5. In the Designation field, select <b>Discard</b> from the list.</li> </ol> | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/> <pre>edit services stateful-firewall rule <i>rule-name</i></pre> </li> <li>2. Configure a term to discard all traffic:<br/> <pre>set term <i>term-name</i> then discard</pre> </li> </ol> |

## 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 107.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying the IPSec Tunnel Configuration” on page 370.

**Table 107: Configuring a NAT Pool for IPSec**

| <b>Task</b>                                                                                | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                               |
|--------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the NAT pool from which the addresses for Network Address Translation are taken. | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, click <b>Services &gt; Nat</b>.</li> <li>2. In the Pool field, click <b>Add new entry</b>.</li> <li>3. In the Pool name field, type the name of the NAT pool. It can be any unique string less than 64 characters long.</li> <li>4. In the Address choice field, select <b>Address</b> from the list.</li> <li>5. In the Address box, type the IP address of the local tunnel endpoint, in dotted decimal notation.</li> </ol> | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><code>edit services nat</code></li> <li>2. Add the local tunnel endpoint to the NAT address pool:<br/><br/><code>set pool <i>pool-name</i> address 1.1.1.1</code></li> </ol> |

**Table 107: Configuring a NAT Pool for IPSec (continued)**

| <b>Task</b>                                                                                                                         | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|-------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the router so that all outgoing traffic is matched against the IP address of the local tunnel endpoint.                   | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, click <b>Services &gt; Nat</b>.</li> <li>2. In the Rule field, click <b>Add new entry</b>.</li> <li>3. In the Rule name field, type the name of the rule. The name can be any unique string.</li> <li>4. In the Match direction field, select <b>Output</b> from the list.</li> <li>5. In the Term field, click <b>Add new entry</b>.</li> <li>6. In the Term name field, type the name of the term. The name can be any unique string.</li> <li>7. Click <b>From</b>.</li> <li>8. In the Source address field, click <b>Add new entry</b>.</li> <li>9. In the address field, select <b>Enter specific value</b> from the list.</li> <li>10. In the Address box, type the IP address of the local tunnel endpoint, in dotted decimal notation, and 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 services nat</code></li> <li>2. Configure a NAT rule and apply it to all output traffic:<br/><code>set rule rule-name match-direction output</code></li> <li>3. Configure the rule to match traffic with a source address that is the same as the local tunnel endpoint:<br/><code>set rule rule-name term term-name from source-address 1.1.1.1</code></li> </ol> |
| Configure the router so that the source address for traffic through the local endpoint is translated to the local endpoint address. | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, click <b>Services &gt; Nat &gt; Rule &gt; rule-name Term &gt; term-name</b></li> <li>2. Click <b>Then</b>.</li> <li>3. Click <b>Translated</b>.</li> <li>4. In the Source pool field, type the name of the NAT pool in which the local tunnel endpoint is configured.</li> <li>5. In the Source field, select <b>Static</b> from the list.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                          | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><code>edit services nat rule rule-name term term-name</code></li> <li>2. Configure the source pool:<br/><code>set then translated source-pool pool-name</code></li> <li>3. Configure the type of translation:<br/><code>set then translated translation-type source static</code></li> </ol>                                                                                  |

## 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 375
- Configuring a Multicast Network on page 385





## Chapter 16

# 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 385.

- Multicast Terms on page 375
- Multicast Architecture on page 378
- Dense and Sparse Routing Modes on page 380
- Strategies for Preventing Routing Loops on page 380
- Multicast Protocol Building Blocks on page 381

## Multicast Terms

---

To understand multicast routing, you must be familiar with the terms defined in Table 108. See Figure 75 for a general view of some of the elements commonly used in an IP multicast network architecture.

**Table 108: Multicast Terms**

| <b>Term</b>                                        | <b>Definition</b>                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 108: Multicast Terms (continued)**

| <b>Term</b>                                   | <b>Definition</b>                                                                                                                                                                                                                                   |
|-----------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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***

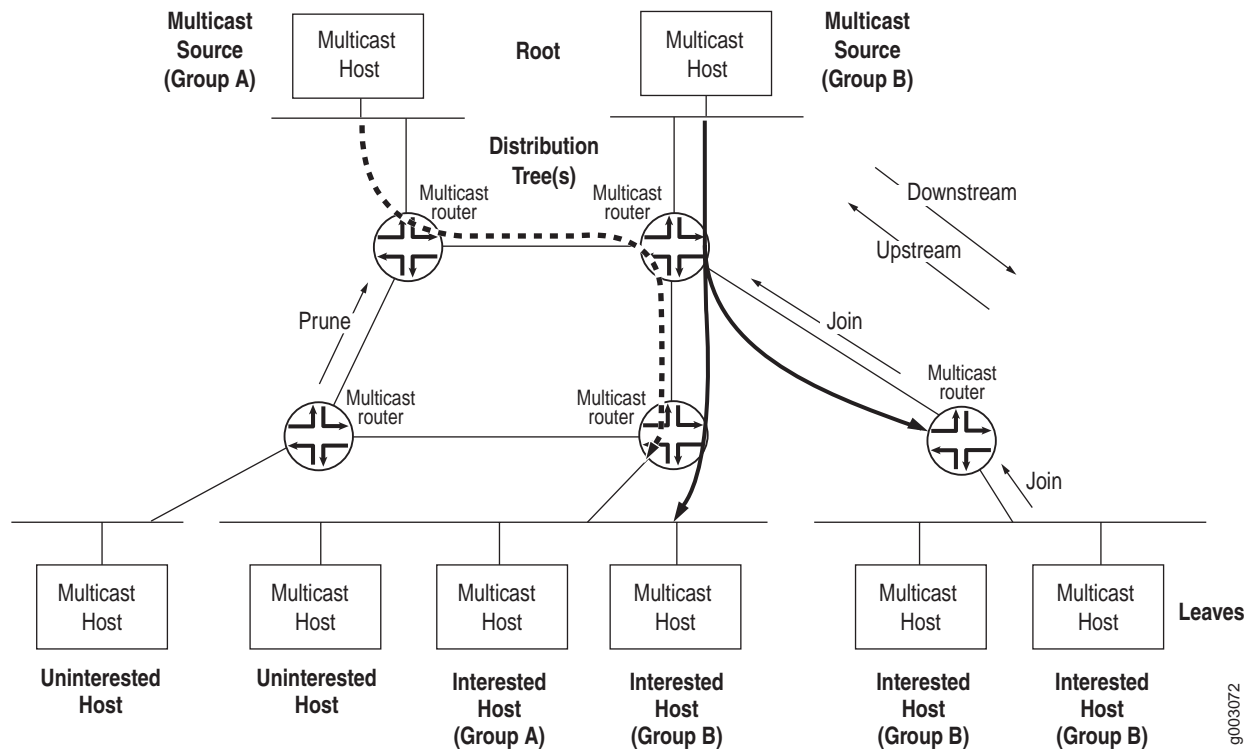
A single upstream interface on the router leads toward the source to receive multicast packets. The downstream interfaces on the router lead toward the receivers to transmit packets. A router can have as many downstream interfaces as it has logical interfaces, minus 1. To prevent looping, the router's 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 109.



**CAUTION:** A common multicast guideline is *not to run dense mode on a WAN under any circumstances*.

**Table 109: 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 110 lists and summarizes these protocols.

**Table 110: 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 110: 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<br>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 110: 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).<br><br>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 17

# 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 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 386
- Configuring a Multicast Network with a Configuration Editor on page 386
- Verifying a Multicast Configuration on page 392

## 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 375.
- 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 Services Router Configuration Tools” on page 3.

- Configuring SAP and SDP (Optional) on page 386
- Configuring IGMP (Required) on page 387
- Configuring the PIM Static RP (Optional) on page 388
- Configuring a PIM RPF Routing Table (Optional) on page 390

### 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 111.
3. Go on to “Configuring IGMP (Required)” on page 387.

**Table 111: Configuring SAP and SDP**

| Task                                                                                                                                                                                            | J-Web Configuration Editor                                                                                                                                                                                                                                                                      | CLI Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Listen</b> level in the configuration hierarchy.                                                                                                                             | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Protocols &gt; Sap</b>.</li> <li>2. Click <b>Add new entry</b> next to Listen.</li> </ol>                                                                                                            | From the top of the configuration hierarchy, enter<br><br><b>edit protocols sap</b>                                                                                                                                                                                                                                                                                                                                                                                           |
| (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"> <li>1. In the Address box, type the multicast address the Services Router can listen to session announcements on, in dotted decimal notation.</li> <li>2. In the Port box, type the port number in decimal notation.</li> <li>3. Click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>1. Set the <b>address</b> value to the IP address that the Services Router can listen to session announcements on, in dotted decimal notation. For example:<br/><br/><b>set listen 224.2.127.254</b></li> <li>2. Set the <b>port</b> value to the number of the port that the Services Router can listen to session announcements on, in decimal notation. For example:<br/><br/><b>set listen 224.2.127.254 port 9875.</b></li> </ol> |

## 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 112.
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 388.
  - To check the configuration, see “Verifying a Multicast Configuration” on page 392.

**Table 112: Explicitly Configuring the IGMP version**

| Task                                                                                                                                                            | J-Web Configuration Editor                                                                                                                                                                                                                            | CLI Configuration Editor                                                                                                                                                                                                                                                                                                             |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Interface</b> level in the configuration hierarchy.                                                                                          | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Protocols &gt; Igmp</b>.</li> <li>2. Click <b>Add new entry</b> next to Interface.</li> </ol>                                                              | <p>From the top of the configuration hierarchy, enter</p> <pre>edit protocols igmp</pre>                                                                                                                                                                                                                                             |
| 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"> <li>1. In the Interface name box, type the name of the interface, or <b>all</b>.</li> <li>2. In the Version box, type the version number: <b>1</b>, <b>2</b>, or <b>3</b>.</li> <li>3. Click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>1. Set the <b>interface</b> value to the interface name, or <b>all</b>. For example:           <pre>set igmp interface all</pre> </li> <li>2. Set the <b>version</b> value to <b>1</b>, <b>2</b>, or <b>3</b>. For example:           <pre>set igmp interface all version 2</pre> </li> </ol> |

### 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 113.
3. Go on to “Configuring a PIM RPF Routing Table (Optional)” on page 390.

**Table 113: Configuring PIM Sparse Mode and the RP**

| Task                                                                   | J-Web Configuration Editor                                                                                                                                                              | CLI Configuration Editor                                                                                          |
|------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Interface</b> level in the configuration hierarchy. | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Protocols &gt; Pim</b>.</li> <li>2. Click <b>Add new entry</b> next to Interface.</li> </ol> | From the top of the configuration hierarchy, enter<br><br><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:<br><br><code>set pim interface all</code> |
| Apply your configuration changes.                                      | Click <b>OK</b> 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 <b>Interface</b> level in the configuration hierarchy.   | Click <b>Add new entry</b> next to Interface.                                                                                                                                           | Remain at the<br><br><code>edit protocols pim interface</code><br><br>configuration hierarchy level.              |
| Disable PIM on the network management interface.                       | <ol style="list-style-type: none"> <li>1. In the Interface name box, type <code>fe-0/0/0</code>.</li> <li>2. Select the check box next to Disable.</li> </ol>                           | Disable the <code>fe-0/0/0</code> interface:<br><br><code>set pim interface fe-0/0/0 unit 0 disable</code>        |
| Apply your configuration changes.                                      | Click <b>OK</b> to apply your entries to the configuration.                                                                                                                             | Changes in the CLI are applied automatically when you execute the <code>set</code> command.                       |

**Table 113: Configuring PIM Sparse Mode and the RP (continued)**

| Task                                                            | J-Web Configuration Editor                                                                                                                                                                                                                                              | CLI Configuration Editor                                                                                                                         |
|-----------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Rp</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Protocols &gt; Pim &gt; Rp</b> .                                                                                                                                                                                       | From the top of the configuration hierarchy, enter<br><br>edit protocols pim rp                                                                  |
| Configure the IP address of the RP.                             | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to Static.</li> <li>2. Click <b>Add new entry</b> next to Address.</li> <li>3. In the Addr box, type the IP address of the RP in dotted decimal notation.</li> <li>4. Click <b>OK</b>.</li> </ol> | Set the <b>address</b> value to the IP address of the RP in dotted decimal notation. For example:<br><br><b>set static address 192.168.14.27</b> |

### 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 114.
3. If you are finished configuring the router, commit the configuration.
4. To check the configuration, see “Verifying a Multicast Configuration” on page 392.



**Table 114: Configuring a PIM RPF Routing Table**

| 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                                                                                                                     |
| Configure a new group for the RPF routing table.                                                                        | Next to Rib groups, click <b>Add new entry</b> .                                                                                                                                                                                                                                                              | Enter<br><br>edit rib-groups                                                                                                                                                                       |
| Configure a name for the RPF routing table group, and use <b>inet.2</b> for its export routing table.                   | <ol style="list-style-type: none"> <li>1. In the Ribgroup name box, type a name for the RPF routing table group—for example, <b>multicast-rpf-rib</b>.</li> <li>2. In the Export rib box, type <b>inet.2</b>.</li> </ol>                                                                                      | Type the name for the RPF routing table and set the export routing table to <b>inet.2</b> . For example:<br><br>set multicast-rpf-rib export-rib inet.2                                            |
| Configure an import routing table routing information base (RIB) group for the RPF routing table.                       | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Import rib.</li> <li>2. In the Value box, type <b>inet.2</b>.</li> <li>3. Click <b>OK</b> three times.</li> </ol>                                                                                                                | Set the import routing table to <b>inet.2</b> . For example:<br><br>set multicast-rpf-rib import-rib inet.2                                                                                        |
| Navigate to the <b>Rib group</b> level in the configuration hierarchy.                                                  | In the configuration editor hierarchy, select <b>Protocols &gt; Pim &gt; Rib group</b> .                                                                                                                                                                                                                      | From the top of the configuration hierarchy, enter<br><br>edit protocols pim                                                                                                                       |
| Apply the RPF routing table to PIM.                                                                                     | <ol style="list-style-type: none"> <li>1. In the Inet box, type the name of the RPF routing table group—for example, <b>multicast-rpf-rib</b>.</li> <li>2. Click <b>OK</b> three times.</li> </ol>                                                                                                            | Enter<br><br>set rib-group multicast-rpf-rib                                                                                                                                                       |
| Create a RIB group for the interface routes.                                                                            | <ol style="list-style-type: none"> <li>1. Navigate to the <b>Routing options</b> level in the configuration hierarchy.</li> <li>2. Next to Rib groups, click <b>Add new entry</b>.</li> </ol>                                                                                                                 | From the top of the configuration hierarchy, enter<br><br>edit routing-options rib-groups.                                                                                                         |
| Configure a name for the RPF routing table group, and use <b>inet.2</b> and <b>inet.0</b> for its import routing table. | <ol style="list-style-type: none"> <li>1. In the Ribgroup name box, type a name for the RPF routing table group—for example, <b>if-rib</b>.</li> <li>2. Click <b>Add new entry</b> next to Import rib.</li> <li>3. In the Value box, type <b>inet.2 inet.0</b>.</li> <li>4. Click <b>OK</b> twice.</li> </ol> | Type the name for the RPF routing table and set the export routing table to <b>inet.2</b> and <b>inet.0</b> . For example:<br><br>set if-rib import-rib inet.2<br><br>set if-rib import-rib inet.0 |
| Add the RIB group to the interface routes.                                                                              | <ol style="list-style-type: none"> <li>1. On the <b>Routing options</b> page, select <b>Interface routes &gt; Rib group</b>.</li> <li>2. In the Inet box, type the name of the interface RIB group—for example, <b>if-rib</b>.</li> <li>3. Click <b>OK</b>.</li> </ol>                                        | From the top of the configuration hierarchy, enter<br><br>edit routing-options interface-routes<br><br>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 392
- Verifying the IGMP Version on page 392
- Verifying the PIM Mode and Interface Configuration on page 393
- Verifying the PIM RP Configuration on page 393
- Verifying the RPF Routing Table Configuration on page 394

### Verifying SAP and SDP Addresses and Ports

| <b>Purpose</b>       | Verify that SAP and SDP are configured to listen on the correct group addresses and ports.                                                                                                                                                                                                                                                                                                                                                                                        |               |      |               |      |
|----------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------|------|---------------|------|
| <b>Action</b>        | From the CLI, enter the <code>show sap listen</code> command.                                                                                                                                                                                                                                                                                                                                                                                                                     |               |      |               |      |
| <b>Sample Output</b> | <pre>user@host&gt; show sap listen</pre> <table> <tr> <th>Group Address</th><th>Port</th></tr> <tr> <td>224.2.127.254</td><td>9875</td></tr> </table>                                                                                                                                                                                                                                                                                                                             | Group Address | Port | 224.2.127.254 | 9875 |
| Group Address        | Port                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |               |      |               |      |
| 224.2.127.254        | 9875                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |               |      |               |      |
| <b>What It Means</b> | <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"> <li>■ Each group address configured, especially the default 224.2.127.254, is listed.</li> <li>■ Each port configured, especially the default 9875, is listed.</li> </ul> <p>For more information about <code>show sap listen</code>, see the <i>JUNOS Routing Protocols and Policies Command Reference</i>.</p> |               |      |               |      |

### Verifying the IGMP Version

|                      |                                                                                                                                                                                                                                                                                                                                        |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b>       | Verify that IGMP version 2 is configured on all applicable interfaces.                                                                                                                                                                                                                                                                 |
| <b>Action</b>        | From the CLI, enter the <code>show igmp interface</code> command.                                                                                                                                                                                                                                                                      |
| <b>Sample Output</b> | <pre>user@host&gt; show igmp interface</pre> <pre>Interface: fe-0/0/0.0   Querier: 192.168.4.36   State:           Up Timeout:      197 Version:  2 Groups:      0</pre> <pre>Configured Parameters: IGMP Query Interval: 125.0 IGMP Query Response Interval: 10.0 IGMP Last Member Query Interval: 1.0 IGMP Robustness Count: 2</pre> |

```

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
pime.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 rpscommand`.

**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 Routing Policy, Firewall Filters, and Class of Service**

- Policy, Firewall Filter, and Class-of-Service Overview on page 397
- Configuring Routing Policies on page 423
- Configuring Firewall Filters and NAT on page 437
- Configuring Class of Service with DiffServ on page 477



## Chapter 18

# 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 Network Interfaces and Class of Service Configuration Guide*.

- Policy, Firewall Filter, and CoS Terms on page 397
- Routing Policy Overview on page 399
- Firewall Filter Overview on page 404
- Class-of-Service Overview on page 413

## 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 115.

**Table 115: Policy, Firewall Filter, and CoS Terms**

| <b>Term</b>                                    | <b>Definition</b>                                                                                                                                                                                                                                                                                                                              |
|------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>assured forwarding (AF)</b>                 | 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.                                                                                                                                                    |
| <b>behavior aggregate (BA) classifier</b>      | 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).                                                                          |
| <b>best-effort (BE)</b>                        | 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.                                                                                                                         |
| <b>class of service (CoS)</b>                  | 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.                                                                                                                                                                            |
| <b>Differentiated Services (DiffServ)</b>      | 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). |
| <b>DiffServ code point (DSCP)</b>              | 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.                                                                                                                                                                                                    |
| <b>drop profile</b>                            | 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.                                                                                                                                                             |
| <b>expedited forwarding (EF)</b>               | CoS packet forwarding class that provides end-to-end service with low loss, low latency, low jitter, and assured bandwidth.                                                                                                                                                                                                                    |
| <b>multifield (MF) classifier</b>              | 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.                                             |
| <b>network address port translation (NAPT)</b> | 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.                                                                                                                                      |
| <b>Network Address Translation (NAT)</b>       | 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.                                                                                                                              |
| <b>network control (NC)</b>                    | CoS packet forwarding class that is typically high priority because it supports protocol control.                                                                                                                                                                                                                                              |
| <b>PLP bit</b>                                 | 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.                |
| <b>policer</b>                                 | 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.                                    |
| <b>policing</b>                                | Applying rate and burst size limits to traffic on an interface.                                                                                                                                                                                                                                                                                |
| <b>random early detection (RED)</b>            | 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.                                                                                                                              |
| <b>rule</b>                                    | Guide that the Services Router follows when applying services. A rule consists of a match direction and one or more terms.                                                                                                                                                                                                                     |
| <b>service set</b>                             | Collection of services. Examples of services include stateful firewall filters and Network Address Translation (NAT).                                                                                                                                                                                                                          |



**Table 115: Policy, Firewall Filter, and CoS Terms (continued)**

| <b>Term</b>                      | <b>Definition</b>                                                                                                                                                                                                                                                                          |
|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>stateful firewall filter</b>  | 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. |
| <b>stateless firewall filter</b> | 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.                                                                         |
| <b>term</b>                      | Firewall filters contain one or more terms that specify filter match conditions and actions.                                                                                                                                                                                               |
| <b>trusted network</b>           | 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.                                                                                                     |
| <b>untrusted network</b>         | 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 399
- Applying Routing Policies on page 404

### 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 400
- “Routing Policy Match Conditions” on page 400
- “Routing Policy Actions” on page 402
- “Default and Final Actions” on page 404

## 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 116 summarizes the routing policy match conditions.

**Table 116: 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 116: 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).<br><br>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><br>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).<br><br>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"> <li>■ <b>egp</b>—Path information originated from another AS.</li> <li>■ <b>igp</b>—Path information originated from within the local AS.</li> <li>■ <b>incomplete</b>—Path information was learned by some other means.</li> </ul> |
| policy [ <i>policy-names</i> ]                                | Name of one or more policies to evaluate as a subroutine.                                                                                                                                                                                                                                                                                                                           |
| preference <i>preference</i><br>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 116: Summary of Routing Policy Match Conditions (continued)**

| Match Condition                                                            | Description                                                                                                                                                                                                                                                                                                                                                 |
|----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| prefix-list <i>name</i>                                                    | Named list of IP addresses configured at the <b>Policy-options</b> level in the configuration hierarchy.<br><br>This match condition can be used on import policies only.                                                                                                                                                                                   |
| protocol <i>protocol</i>                                                   | 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: <b>aggregate</b> , <b>bgp</b> , <b>direct</b> , <b>dvmrp</b> , <b>isis</b> , <b>local</b> , <b>ospf</b> , <b>pim-dense</b> , <b>pim-sparse</b> , <b>rip</b> , <b>ripng</b> , or <b>static</b> .                            |
| route-filter <i>destination-prefix match-type &lt;actions&gt;</i>          | 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.<br><br>Route filters can be used on import policies only.            |
| route-type <i>value</i>                                                    | Type of route. The value can be either <b>external</b> or <b>internal</b> .                                                                                                                                                                                                                                                                                 |
| source-address-filter <i>destination-prefix match-type &lt;actions&gt;</i> | 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.<br><br>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 117 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 117: Summary of Key Routing Policy Actions**

| Action                                              | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|-----------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Flow Control Actions</b>                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 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 <b>accept</b> or <b>reject</b> action specified in the <b>then</b> statement is ignored. Any actions specified in the <b>then</b> statement that manipulate route characteristics are applied to the route.                                                                                                                                                                                                                                                                                         |
| next policy                                         | Skips to and evaluates the next routing policy. Any <b>accept</b> or <b>reject</b> action specified in the <b>then</b> statement is ignored. Any actions specified in the <b>then</b> statement that manipulate route characteristics are applied to the route.                                                                                                                                                                                                                                                                                                          |
| <b>Route Manipulation Actions</b>                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| 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><br>color2 <i>preference</i> | Sets the preference value to the specified value. The <b>color</b> and <b>color2</b> 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 117: 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 <b>metric</b> (for the first metric value) and continuing with <b>metric2</b> , <b>metric3</b> , and <b>metric4</b> .<br><br>For BGP routes, <b>metric</b> corresponds to the MED, and <b>metric2</b> 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.<br><br>If you specify <i>address</i> as <b>self</b> , 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*.

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 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.

---

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).

This section contains the following topics:

- Stateful and Stateless Firewall Filters on page 405
- Process for Configuring a Stateful Firewall Filter and NAT on page 406
- Summary of Stateful Firewall Filter and NAT Match Conditions and Actions on page 406
- Planning a Stateless Firewall Filter on page 408
- Stateless Firewall Filter Match Conditions, Actions, and Action Modifiers on page 409

## **Stateful and Stateless Firewall Filters**

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 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 357 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 406.

## Summary of Stateful Firewall Filter and NAT Match Conditions and Actions

Table 118 lists the match conditions you can specify in stateful firewall filter and NAT terms. Table 119 and Table 120 list actions you can specify in stateful firewall filter and NAT terms.



**Table 118: 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 119: 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 <b>accept</b> action.<br><br>You can specify the IP option as text or a numeric value: <b>any</b> (0), <b>ip-security</b> (130), <b>ip-stream</b> (8), <b>loose-source-route</b> (3), <b>route-record</b> (7), <b>router-alert</b> (148), <b>strict-source-route</b> (9), and <b>timestamp</b> (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 <b>discard</b> .                                                                                                                                                                                                                                                                                                                                                                                                                              |

**Table 120: NAT Actions**

| Actions                                             | Description                                                 |
|-----------------------------------------------------|-------------------------------------------------------------|
| syslog                                              | Record information in the system logging facility.          |
| translated destination-pool<br><i>nat-pool-name</i> | Translate the destination address using the specified pool. |
| translated source-pool<br><i>nat-pool-name</i>      | Translate the source address using the specified pool.      |

**Table 120: NAT Actions (continued)**

| Actions                                                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
|------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| translation-type (destination <i>type</i>   source <i>type</i> ) | <p>Translate the destination and source port using the specified type:</p> <ul style="list-style-type: none"> <li>■ <b>destination static</b>—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 <b>destination-pool</b> name. The referenced pool must contain exactly one address and no <b>port</b> configuration at the [edit nat pool] hierarchy level.</li> <li>■ <b>source dynamic</b>—Translate the source address with port mapping by means of NAT. You must specify a <b>source-pool</b> name. The referenced pool must include a <b>port</b> configuration at the [edit nat pool] hierarchy level.</li> <li>■ <b>source static</b>—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 <b>source-pool</b> name. The referenced pool must contain exactly one address and no <b>port</b> configuration at the [edit nat pool] hierarchy level.</li> </ul> |
| 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 409. For more information about stateless firewall filters, see the *JUNOS Policy Framework Configuration Guide*.

## Stateless Firewall Filter Match Conditions, Actions, and Action Modifiers

Table 121 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 122 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 121: Stateless Firewall Filter Match Conditions**

| Match Condition                       | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|---------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Numeric Range Match Conditions</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <i>keyword-except</i>                 | <p>Negates a match. For example, <b>destination-port-except</b> <i>number</i> .</p> <p>The following keywords accept the <b>-except</b> extension: <b>destination-port</b>, <b>dscp</b>, <b>esp-spi</b>, <b>forwarding-class</b>, <b>fragment-offset</b>, <b>icmp-code</b>, <b>icmp-type</b>, <b>interface-group</b>, <b>ip-options</b>, <b>packet-length</b>, <b>port</b>, <b>precedence</b>, <b>protocol</b> and <b>source-port</b>.</p>                                                                                                      |
| <i>destination-port number</i>        | <p>TCP or User Datagram Protocol (UDP) destination port field. You cannot specify both the <b>port</b> and <b>destination-port</b> match conditions in the same term. Normally, you specify this match in conjunction with the <b>protocol tcp</b> or <b>protocol udp</b> 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 <b>telnet</b> or <b>23</b>.</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 <b>assured-forwarding</b>, <b>best-effort</b>, <b>expedited-forwarding</b>, or <b>network-control</b>.</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 <b>protocol icmp</b> match statement to determine which protocol is being used on the port.</p> <p>This value or keyword provides more specific information than <b>icmp-type</b>. Because the value's meaning depends on the associated <b>icmp-type</b>, you must specify <b>icmp-type</b> along with <b>icmp-code</b>.</p> <p>In place of the numeric value, you can specify a text synonym. For example, you can specify <b>ip-header-bad</b> or <b>0</b>.</p> |
| <i>icmp-type number</i>               | <p>ICMP packet type field. Normally, you specify this match in conjunction with the <b>protocol icmp</b> 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 <b>time-exceeded</b> or <b>11</b>.</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 <b>port</b> match and either the <b>destination-port</b> or <b>source-port</b> match conditions in the same term. Normally, you specify this match in conjunction with the <b>protocol tcp</b> or <b>protocol udp</b> 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 <b>bgp</b> or <b>179</b>.</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 <b>immediate</b> or <b>0x40</b>.</p>                                                                                                                                                                                                                                                                                                      |

**Table 121: 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 <b>ospf</b> or <b>89</b> .                                                                                                                                                                                                                                                                                           |
| source-port <i>number</i>                      | TCP or UDP source port field. You cannot specify the <b>port</b> and <b>source-port</b> match conditions in the same term. Normally, you specify this match in conjunction with the <b>protocol tcp</b> or <b>protocol udp</b> match statement to determine which protocol is being used on the port.<br><br>In place of the numeric value, you can specify a text synonym. For example, you can specify <b>http</b> or <b>80</b> . |
| <b>Address Match Conditions</b>                |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| address <i>prefix</i>                          | IP source or destination address field. You cannot specify both the <b>address</b> and the <b>destination-address</b> or <b>source-address</b> match conditions in the same term.                                                                                                                                                                                                                                                   |
| destination-address <i>prefix</i>              | IP destination address field. You cannot specify the <b>destination-address</b> and <b>address</b> match conditions in the same term.                                                                                                                                                                                                                                                                                               |
| destination-prefix-list <i>prefix-list</i>     | IP destination prefix list field. You cannot specify the <b>destination-prefix-list</b> and <b>prefix-list</b> match conditions in the same term.                                                                                                                                                                                                                                                                                   |
| prefix-list <i>prefix-list</i>                 | IP source or destination prefix list field. You cannot specify both the <b>prefix-list</b> and the <b>destination-prefix-list</b> or <b>source-prefix-list</b> match conditions in the same term.                                                                                                                                                                                                                                   |
| source-address <i>prefix</i>                   | IP source address field. You cannot specify the <b>source-address</b> and <b>address</b> match conditions in the same rule.                                                                                                                                                                                                                                                                                                         |
| source-prefix-list <i>prefix-list</i>          | IP source prefix list field. You cannot specify the <b>source-prefix-list</b> and <b>prefix-list</b> match conditions in the same term.                                                                                                                                                                                                                                                                                             |
| <b>Bit-Field Match Conditions with Values</b>  |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| 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 <b>more-fragments</b> or <b>0x2000</b> .                                                                                                                                                                                                                                                                        |
| ip-options <i>number</i>                       | IP options. In place of the numeric value, you can specify a text synonym. For example, you can specify <b>record-route</b> or <b>7</b> .                                                                                                                                                                                                                                                                                           |
| tcp-flags <i>number</i>                        | TCP flags. Normally, you specify this match in conjunction with the <b>protocol tcp</b> 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 <b>syn</b> or <b>0x02</b> .                                                                                                                                              |
| <b>Bit-Field Text Synonym Match Conditions</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| 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 <b>fragment-offset 0-8191</b> .                                                                                                                                                                                         |

**Table 121: 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)".<br><br>This condition does not implicitly check that the protocol is TCP. To do so, specify the <b>protocol tcp</b> match condition. |
| tcp-initial     | First TCP packet of a connection. This match condition is a synonym for "(syn & !ack)".<br><br>This condition does not implicitly check that the protocol is TCP. To do so, specify the <b>protocol tcp</b> match condition.                       |

**Table 122: Stateless Firewall Filter Bit-Field Logical Operators**

| Logical Operator | Description |
|------------------|-------------|
| (...)            | Grouping    |
| !                | Negation    |
| & or +           | Logical AND |
| or ,             | Logical OR  |

Table 123 lists the actions and action modifiers you can specify in stateless firewall filter terms.

**Table 123: 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 <b>then</b> 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: <b>administratively-prohibited</b> (default), <b>bad-host-tos</b> , <b>bad-network-tos</b> , <b>host-prohibited</b> , <b>host-unknown</b> , <b>host-unreachable</b> , <b>network-prohibited</b> , <b>network-unknown</b> , <b>network-unreachable</b> , <b>port-unreachable</b> , <b>precedence-cutoff</b> , <b>precedence-violation</b> , <b>protocol-unreachable</b> , <b>source-host-isolated</b> , <b>source-route-failed</b> , or <b>tcp-reset</b> . If you specify <b>tcp-reset</b> , a TCP reset is returned if the packet is a TCP packet. Otherwise, nothing is returned. |
| routing-instance<br><i>routing-instance</i> | Routes the packet using the specified routing instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Action Modifiers</b>                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |

**Table 123: 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 <b>show firewall log</b> command at the CLI.                                                                                                                   |
| loss-priority <i>priority</i>      | Sets the scheduling priority of the packet. The priority can be <b>low</b> or <b>high</b> .                                                                                                                                                                                |
| 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 <b>discard</b> .                                                                                                                                        |

## 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 477.

This overview contains the following topics. For more information about CoS and DiffServ, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.

- Benefits of DiffServ CoS on page 413
- DSCPs and Forwarding Service Classes on page 414
- JUNOS CoS Functions on page 415
- How Forwarding Classes and Schedulers Work on page 417

### 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 124 identifies the forwarding service classes and aliases that correspond to the 21 DSCPs.

**Table 124: Default Forwarding Service Class-to-DSCP Mapping**

| DiffServ<br>Service Class<br>Alias | IP DSCP | Forwarding Service Class and Use                                                                                                                                                                                                                                                                                                                                           |
|------------------------------------|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ef                                 | 101110  | <p><b>Expedited forwarding</b>—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> |



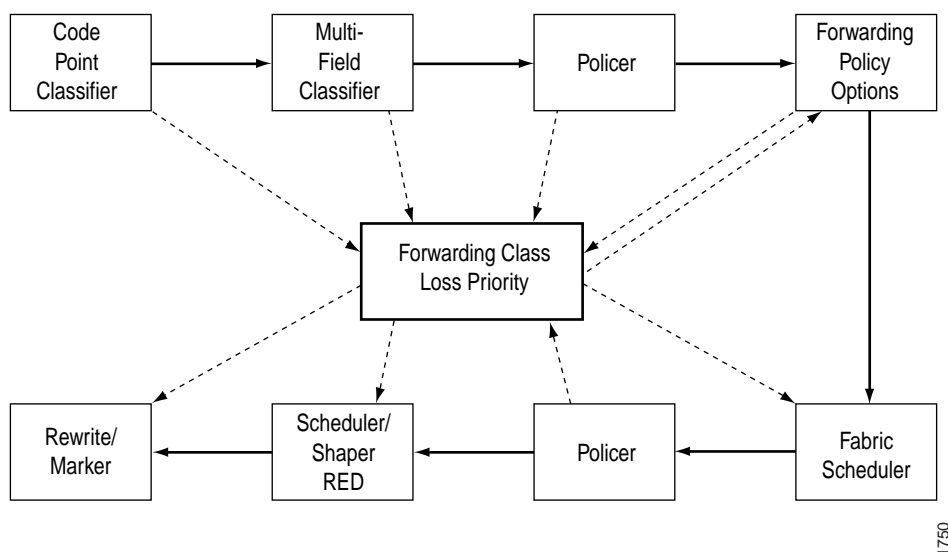
**Table 124: Default Forwarding Service Class-to-DSCP Mapping (continued)**

| <b>DiffServ<br/>Service Class<br/>Alias</b> | <b>IP DSCP</b> | <b>Forwarding Service Class and Use</b>                                                                                                                                                                                                                                                                                                                                                                               |
|---------------------------------------------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| af11                                        | 001010         | <b>Assured forwarding</b> —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         | <b>Best-effort</b> —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         | <b>Conversational services</b> —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.<br><br>Packetized voice is a good example of a conversational service.                                                                                                                     |
| cs2                                         | 010000         |                                                                                                                                                                                                                                                                                                                                                                                                                       |
| cs3                                         | 011000         |                                                                                                                                                                                                                                                                                                                                                                                                                       |
| cs4                                         | 100000         |                                                                                                                                                                                                                                                                                                                                                                                                                       |
| cs5                                         | 101000         |                                                                                                                                                                                                                                                                                                                                                                                                                       |
| nc1/cs6                                     | 110000         | <b>Network control</b> —The Services Router delivers packets in this service class with a low priority. (These packets are not delay sensitive.)<br><br>Typically, these packets represent routing protocol hello or keepalive messages. Because loss of these packets jeopardizes proper network operation, delay is preferable to discard.<br><br>(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 76 shows the components of the JUNOS CoS features, illustrating the sequence in which they interact. Table 125 defines the components and explains their use.

**Figure 76: Packet Flow Through JUNOS CoS-Configurable Components****Table 125: 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"> <li>■ Behavior aggregate (BA) or code point traffic classifiers—Allow you to set the forwarding class and PLP based on DSCP.</li> <li>■ 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.</li> </ul> |
| 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"> <li>■ Allow you to associate forwarding classes with next hops.</li> <li>■ Allow you to create classification overrides, which assign forwarding classes to sets of prefixes.</li> </ul>                                                                                                                                                                                                                                                                                                                                         |

**Table 125: JUNOS CoS Components (continued)**

| <b>CoS Component</b>                     | <b>Use</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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"> <li>■ 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.</li> <li>■ 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.</li> <li>■ 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.</li> </ul> |
| 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 417
- “Default Scheduler Settings” on page 418
- “Default Behavior Aggregate (BA) Classifiers” on page 419
- “DSCP Rewrites” on page 420
- “Sample BA Classification” on page 420

### **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 126 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 126: 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 127 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 477.

**Table 127: 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 127: 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 485.

## Sample BA Classification

Table 128 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 477.

**Table 128: 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 128: Sample BA Classification Forwarding Classes and Queues (continued)**

| <b>DSCP Alias</b> | <b>DSCP Bits</b> | <b>Forwarding Class</b> | <b>PLP</b> | <b>Queue</b> |
|-------------------|------------------|-------------------------|------------|--------------|
| 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 19

# 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 424
- Configuring a Routing Policy with a Configuration Editor on page 424

## 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 399.
- 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 101.
- Configure an Interior Gateway Protocol (IGP) and Border Gateway Protocol (BGP), if necessary. See “Configuring BGP Sessions” on page 273.
- Configure the router interface to reject or accept routes, if necessary. See “Configuring Firewall Filters and NAT” on page 437.
- Configure static routes, if necessary. See “Configuring Static Routes” on page 223.

## 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 Services Router Configuration Tools” on page 3.

- Configuring the Policy Name (Required) on page 425
- Configuring a Policy Term (Required) on page 425
- Rejecting Known Invalid Routes (Optional) on page 426
- Injecting OSPF Routes into the BGP Routing Table (Optional) on page 428
- Grouping Source and Destination Prefixes in a Forwarding Class (Optional) on page 430
- Configuring a Policy to Prepend the AS Path (Optional) on page 431
- Configuring Damping Parameters (Optional) on page 433

## 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 129.
3. Go on to “Configuring a Policy Term (Required)” on page 425.

**Table 129: Configuring the Policy Name**

| Task                                                                          | J-Web Configuration Editor                                                                        | CLI Configuration Editor                                                              |
|-------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| Navigate to the <b>Policy statement</b> level in the configuration hierarchy. | In the J-Web configuration editor hierarchy, select <b>Policy options &gt; Policy statement</b> . | From the top of the CLI configuration hierarchy, enter<br><br>edit policy-options     |
| Enter the policy name.                                                        | In the Policy name box, type the name of the policy.                                              | Type the <b>policy-name</b> value. For example:<br><br>set policy-statement policy1   |
| Apply your configuration changes.                                             | Click <b>OK</b> to apply your entries to the configuration.                                       | Changes in the CLI are applied automatically when you execute the <b>set</b> command. |

## 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 130.
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 426.
- To advertise additional routes, see “Injecting OSPF Routes into the BGP Routing Table (Optional)” on page 428.
- To create a forwarding class, see “Grouping Source and Destination Prefixes in a Forwarding Class (Optional)” on page 430.
- To make a route less preferable to BGP, see “Configuring a Policy to Prepend the AS Path (Optional)” on page 431.
- To suppress route information, see “Configuring Damping Parameters (Optional)” on page 433.

**Table 130: Configuring a Policy Term**

| Task                                                                          | J-Web Configuration Editor                                                                                                                                                       | CLI Configuration Editor                                                                                   |
|-------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Policy statement</b> level in the configuration hierarchy. | In the J-Web configuration editor hierarchy, select <b>Policy options &gt; Policy statement</b> .                                                                                | From the top of the CLI configuration hierarchy, enter<br><br>edit policy-options policy-statement policy1 |
| Create and name a policy term.                                                | <ol style="list-style-type: none"> <li>1. In the Term box, click <b>Add new entry</b>.</li> <li>2. In the Term name box, type the name of a term and click <b>OK</b>.</li> </ol> | Create and name a policy term. For example:<br><br>set term term1                                          |

### 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 131 lists route list match types.

**Table 131: Route List Match Types**

| Match Type | Match If ...                                                                                                                                              |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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. |

**Table 131: Route List Match Types (continued)**

| Match Type                                                        | Match If ...                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|-------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| 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"> <li>■ The route shares the same most-significant bits (described by <i>prefix-length</i>) of the first destination prefix.</li> <li>■ 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.</li> <li>■ The number of bits in the route's prefix length is less than or equal to the number of bits in the second prefix.</li> </ul> <p>You do not use the <b>through</b> 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, 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 accept routes less than 8 bits in length:

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 132.
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 428.
  - To create a forwarding class, see “Grouping Source and Destination Prefixes in a Forwarding Class (Optional)” on page 430.
  - To make a route less preferable to BGP, see “Configuring a Policy to Prepend the AS Path (Optional)” on page 431.

- To suppress route information, see “Configuring Damping Parameters (Optional)” on page 433.

**Table 132: Creating a Policy to Reject Known Invalid Routes**

| Task                                                              | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                              | CLI Configuration Editor                                                                                                                                                                             |
|-------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Term</b> level in the configuration hierarchy. | In the J-Web configuration editor hierarchy, select <b>Policy options &gt; Policy statement &gt; Term</b> .                                                                                                                                                                                                                                                                             | From the top of the CLI configuration hierarchy, enter<br><br>edit policy-options policy-statement rejectpolicy1 term rejectterm1                                                                    |
| Specify the routes to accept.                                     | <ol style="list-style-type: none"> <li>1. In the From option, click <b>Configure</b>.</li> <li>2. In the Route filter box, click <b>Add new entry</b>.</li> <li>3. In the Address box, enter the prefix of the routes.</li> <li>4. Click <b>OK</b>.</li> </ol>                                                                                                                          | Accept routes less than 8 bits in length:<br><br>set from route-filter 0/0 up to /7 accept                                                                                                           |
| Accept these routes.                                              | <ol style="list-style-type: none"> <li>1. In the Then option, click <b>Configure</b>.</li> <li>2. In the Accept option, select the <b>Yes</b> check box.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                        |                                                                                                                                                                                                      |
| Specify the routes to reject.                                     | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Policy options &gt; Policy statement &gt; Term</b>.</li> <li>2. In the From option, click <b>Configure</b>.</li> <li>3. In the Route filter box, click <b>Add new entry</b>.</li> <li>4. In the Value box, enter the prefix of the routes to reject.</li> <li>5. Click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>1. Specify routes less than 8 bits in length:<br/><br/>set from route-filter /8 orlonger</li> <li>2. Reject these routes:<br/><br/>set then reject</li> </ol> |
| Reject these routes.                                              | <ol style="list-style-type: none"> <li>1. In the Then option, click <b>Configure</b>.</li> <li>2. In the Reject option, select the <b>Yes</b> check box.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                        |                                                                                                                                                                                                      |

### ***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 redistribute OSPF routes from area 1 only into BGP and 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 133.
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 430.
  - To make a route less preferable to BGP, see “Configuring a Policy to Prepend the AS Path (Optional)” on page 431.
  - To suppress route information, see “Configuring Damping Parameters (Optional)” on page 433.

**Table 133: Creating a Policy to Inject OSPF Routes into BGP**

| Task                                                                             | J-Web Configuration Editor                                                                                                                                                                                                                        | CLI Configuration Editor                                                                                                          |
|----------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Term</b> level in the configuration hierarchy.                | In the J-Web configuration editor hierarchy, select <b>Policy options &gt; Policy statement &gt; Term</b> .                                                                                                                                       | From the top of the CLI configuration hierarchy, enter<br><br>edit policy-options policy-statement injectpolicy1 term injectterm1 |
| Specify the OSPF routes.                                                         | <ol style="list-style-type: none"> <li>1. In the From option, click <b>Configure</b>.</li> <li>2. In the Protocol box, click <b>Add new entry</b>.</li> <li>3. In the Value drop box, select <b>OSPF</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> | Specify the OSPF match condition:<br><br>set from ospf                                                                            |
| Specify the routes from a particular OSPF area.                                  | <ol style="list-style-type: none"> <li>1. In the Area option, type <b>1</b>.</li> <li>2. Click <b>OK</b>.</li> </ol>                                                                                                                              | Specify Area 1 as a match condition:<br><br>set from area 1                                                                       |
| Specify that the route is to be accepted if the previous conditions are matched. | <ol style="list-style-type: none"> <li>1. Next to Then, click <b>Configure</b>.</li> <li>2. From the Accept reject box, Select <b>Accept</b>.</li> </ol>                                                                                          | Specify the action to accept:<br><br>set then accept                                                                              |

**Table 133: Creating a Policy to Inject OSPF Routes into BGP (continued)**

| <b>Task</b>                                                                    | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                            | <b>CLI Configuration Editor</b>                                                       |
|--------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| Set the default option to reject other OSPF routes.                            | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Policy options &gt; Policy statement &gt; Term</b>.</li> <li>2. In the Then option, click <b>Configure</b>.</li> <li>3. From the Accept reject box, Select <b>Reject</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> | Changes in the CLI are applied automatically when you execute the <b>set</b> command. |
| Navigate to the <b>Protocol &gt; Bgp</b> level in the configuration hierarchy. | In the J-Web configuration editor hierarchy, select <b>Protocols &gt; Bgp</b> .                                                                                                                                                                                                                              | From the top of the CLI configuration hierarchy, enter:<br><br>edit protocols bgp     |
| Apply the routing policy <b>policy1</b> to BGP.                                | <ol style="list-style-type: none"> <li>1. In the Export box, click <b>Add new entry</b>.</li> <li>2. In the Value option, enter <b>policy1</b>.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                      | Specify the OSPF match condition:<br><br>set export policy1                           |

### **Grouping Source and Destination Prefixes in a Forwarding Class (Optional)**

Create a forwarding class that includes packets based on both the destination address and the source address in the packet.

To configure and apply a routing policy to group source and destination prefixes in a forwarding class:

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 134.
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 make a route less preferable to BGP, see “Configuring a Policy to Prepend the AS Path (Optional)” on page 431.
  - To suppress route information, see “Configuring Damping Parameters (Optional)” on page 433.



**Table 134: Creating a Policy to Group Source and Destination Prefixes in a Forwarding Class**

| Task                                                                          | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                   | CLI Configuration Editor                                                                                                                                                                                                                                                        |
|-------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Term</b> level in the configuration hierarchy.             | In the J-Web configuration editor hierarchy, select <b>Policy options &gt; Policy statement &gt; Term</b> .                                                                                                                                                                                                                                                                  | From the top of the CLI configuration hierarchy, enter<br><br>edit policy-options policy-statement policy1 term term1                                                                                                                                                           |
| Specify the routes to include in the route filter.                            | <ol style="list-style-type: none"> <li>1. In the From option, click <b>Configure</b>.</li> <li>2. In the Route filter box, click <b>Add new entry</b>.</li> <li>3. In the Value box, enter the source and destination prefixes.</li> <li>4. Click <b>OK</b>.</li> </ol>                                                                                                      | <ol style="list-style-type: none"> <li>1. Specify source routes 10.210.0.0/16 or longer:<br/><br/>set from route-filter 10.210.0.0/16 orlonger</li> <li>2. Specify destination routes 10.215.0.0/16 or longer:<br/><br/>set from route-filter 10.215.0.0/16 orlonger</li> </ol> |
| Group the source and destination prefixes.                                    | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Policy options &gt; Policy statement &gt; Term</b>.</li> <li>2. In the Then option, click <b>Configure</b>.</li> <li>3. In the Forwarding class box, enter the forwarding class name.</li> <li>4. Click <b>OK</b>.</li> </ol>                                                     | Specify the forwarding class name:<br><br>set then forwarding class forwarding-class-name1                                                                                                                                                                                      |
| Navigate to the <b>Forwarding table</b> level in the configuration hierarchy. | In the J-Web configuration editor hierarchy, select <b>Routing options &gt; Forwarding table</b> .                                                                                                                                                                                                                                                                           | From the top of the CLI configuration hierarchy, enter<br><br>edit routing-options forwarding-table                                                                                                                                                                             |
| Apply the policy to the forwarding table.                                     | <ol style="list-style-type: none"> <li>1. In the Export box, click <b>Add new entry</b>.</li> <li>2. In the Value box, enter the name of the policy.</li> <li>3. Click <b>OK</b>.</li> </ol> <p>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.</p> | Specify source routes 10.210.0.0/16 or longer:<br><br>set export policy1<br><br>You can refer to the same routing policy one or more times in the same or a different <b>export</b> 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 prepend 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 135.
3. If you are finished configuring the router, commit the configuration.
4. To suppress route information, see “Configuring Damping Parameters (Optional)” on page 433.

**Table 135: Creating a Policy to Prepend AS Numbers**

| Task                                                              | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                              | CLI Configuration Editor                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Term</b> level in the configuration hierarchy. | In the J-Web configuration editor hierarchy, select <b>Policy options &gt; Policy statement &gt; Term</b> .                                                                                                                                                                                                                                                             | From the top of the CLI configuration hierarchy, enter<br><br>edit policy-options policy-statement<br>prependpolicy1 term prependterm1                                                                                                                                                                                                                                      |
| Specify the routes to prepend AS numbers to.                      | <ol style="list-style-type: none"> <li>1. In the From option, click <b>Configure</b>.</li> <li>2. In the Value box, enter the prefixes you wish to prepend.</li> <li>3. In the Route filter box, click <b>Add new entry</b>.</li> <li>4. Click <b>OK</b>.</li> </ol>                                                                                                    | <ol style="list-style-type: none"> <li>1. Prepend routes 172.168.0.0/12 or longer:<br/><br/>set from route-filter<br/>172.16.0.0/12 orlonger</li> <li>2. Prepend routes 192.168.0.0/16 or longer:<br/><br/>set from route-filter<br/>192.168.0.0/16 orlonger</li> <li>3. Prepend routes 10.0.0.0/8 or longer:<br/><br/>set from route-filter 10.0.0.0/8 orlonger</li> </ol> |
| Specify the AS numbers to prepend.                                | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Policy options &gt; Policy statement &gt; Term</b>.</li> <li>2. In the Then option, click <b>Configure</b>.</li> <li>3. In the AS path prepend box, enter the string of AS numbers to prepend. Separate each AS number with a space.</li> <li>4. Click <b>OK</b>.</li> </ol> | Specify the AS numbers to prepend, and enclose them inside double quotation marks:<br><br>set then as-path-prepend “1 1 1 1”                                                                                                                                                                                                                                                |

**Table 135: Creating a Policy to Prepend AS Numbers (continued)**

| Task                                                                                 | J-Web Configuration Editor                                                                                                                                                                                                                                                               | CLI Configuration Editor                                                                                                                                                        |
|--------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Protocols &gt; BGP &gt;</b> level in the configuration hierarchy. | In the J-Web configuration editor hierarchy, select <b>Protocols &gt; BGP &gt;</b> .                                                                                                                                                                                                     | From the top of the CLI configuration hierarchy, enter<br><br>edit protocols bgp                                                                                                |
| Apply the policy as an import policy for all BGP routes.                             | <ol style="list-style-type: none"> <li>1. In the Import box, click <b>Add new entry</b>.</li> <li>2. In the Value box, enter the name of the policy.</li> <li>3. Click <b>OK</b>.</li> </ol> <p>The routing policy is evaluated when routes are being imported to the routing table.</p> | <p>Apply the policy:</p> <p>set import prependpolicy1</p> <p>You can refer to the same routing policy one or more times in the same or a different <b>import</b> statement.</p> |

### 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.)

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, 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 136.
3. If you are finished configuring the router, commit the configuration.

**Table 136: Creating a Policy to Accept and Apply Damping on Routes**

| <b>Task</b>                                                                 | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                       | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                  |
|-----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Term</b> level in the configuration hierarchy.           | In the J-Web configuration editor hierarchy, select <b>Policy options &gt; Policy statement &gt; Term</b> .                                                                                                                                                                                                                                                             | From the top of the CLI configuration hierarchy, enter<br><br>edit policy-options policy-statement dampenpolicy1 term dampenterm1                                                                                                                                                                                                                                |
| Specify the routes to dampen.                                               | <ol style="list-style-type: none"> <li>1. In the From option, click <b>Configure</b>.</li> <li>2. In the Value box, enter the prefixes you wish to dampen.</li> <li>3. In the Route filter box, click <b>Add new entry</b>.</li> <li>4. In the Value box, enter the prefixes you wish to dampen.</li> <li>5. Click <b>OK</b>.</li> </ol>                                | <ol style="list-style-type: none"> <li>1. Dampen routes 172.168.0.0/16 or longer:<br/><br/>set from route-filter 172.16.0.0/12 orlonger</li> <li>2. Dampen routes 192.168.0.0/16 or longer:<br/><br/>set from route-filter 192.168.0.0/16 orlonger</li> <li>3. Dampen routes 10.0.0.0/8 or longer:<br/><br/>set from route-filter 10.0.0.0/8 orlonger</li> </ol> |
| Specify the damping parameters group to apply to the route filter.          | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Policy options &gt; Policy statement &gt; Term</b>.</li> <li>2. In the Then option, click <b>Configure</b>.</li> <li>3. In the AS path prepend box, enter the string of AS numbers to prepend. Separate each AS number with a space.</li> <li>4. Click <b>OK</b>.</li> </ol> | Specify the AS numbers to prepend, and enclose inside them inside double quotation marks:<br><br>set then as-path-prepend "1 1 1 1"                                                                                                                                                                                                                              |
| Navigate to the <b>Policy options</b> level in the configuration hierarchy. | In the J-Web configuration editor hierarchy, select <b>Policy options</b> .                                                                                                                                                                                                                                                                                             | From the top of the CLI configuration hierarchy, enter<br><br>edit policy-options                                                                                                                                                                                                                                                                                |

**Table 136: Creating a Policy to Accept and Apply Damping on Routes (continued)**

| <b>Task</b>                                                                                                                                                                                            | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>CLI Configuration Editor</b>                                                                                                                                                                                                   |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create a damping parameter group.                                                                                                                                                                      | <ol style="list-style-type: none"> <li>1. In the Damping box, click <b>Add new entry</b>.</li> <li>2. In the Damping object name box, enter the name of the damping parameter group.</li> <li>3. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                          | <p>Create and configure the damping parameter groups:</p> <pre>edit damping group1 half-life 30 suppress 3000 reuse 750 max-suppress 60</pre> <pre>edit damping group2 half-life 40 suppress 400 reuse 1000 max-suppress 45</pre> |
| Configure a damping parameter group.                                                                                                                                                                   | <ol style="list-style-type: none"> <li>1. In the Half life box, enter the half life duration, in minutes.</li> <li>2. In the Max suppress box, enter the maximum holddown time, in minutes.</li> <li>3. In the Reuse box, enter the reuse threshold, for this damping group.</li> <li>4. In the Suppress box, enter the cutoff threshold, for this damping group.</li> <li>5. To disable damping for this damping group, select the <b>Disable</b> check box.</li> <li>6. Click <b>OK</b>.</li> </ol> | <pre>edit damping group3 disable</pre>                                                                                                                                                                                            |
| Navigate to the <b>BGP</b> level in the configuration hierarchy.                                                                                                                                       | In the J-Web configuration editor hierarchy, select <b>Protocols &gt; Bgp</b> .                                                                                                                                                                                                                                                                                                                                                                                                                       | <p>From the top of the CLI configuration hierarchy, enter</p> <pre>edit protocols bgp</pre>                                                                                                                                       |
| Enable damping.                                                                                                                                                                                        | <ol style="list-style-type: none"> <li>1. Select the <b>Damping</b> check box.</li> <li>2. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                | <p>Enable damping:</p> <pre>set damping</pre>                                                                                                                                                                                     |
| Navigate to the <b>Neighbor</b> 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 <b>172.16.15.14</b> . | In the J-Web configuration editor hierarchy, select <b>Protocols &gt; Bgp &gt; Group Group1 &gt; Neighbor 172.16.15.14</b> .                                                                                                                                                                                                                                                                                                                                                                          | <p>From the top of the CLI configuration hierarchy, enter</p> <pre>edit protocols bgp group group1 neighbor 172.16.15.14</pre>                                                                                                    |
| Apply the policy as an import policy for the BGP neighbor.                                                                                                                                             | <ol style="list-style-type: none"> <li>1. In the Import box, click <b>Add new entry</b>.</li> <li>2. In the Value box, enter the name of the policy.</li> <li>3. Click <b>OK</b>.</li> </ol> <p>The routing policy is evaluated when routes are imported to the routing table.</p>                                                                                                                                                                                                                    | <p>Apply the policy:</p> <pre>set import dampenpolicy1</pre> <p>You can refer to the same routing policy one or more times in the same or a different <b>import</b> statement.</p>                                                |



## Chapter 20

# 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. To configure a stateless firewall filter, use a configuration editor.

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 438
- Configuring a Stateful Firewall Filter with Quick Configuration on page 438
- Configuring a Stateful Firewall Filter with a Configuration Editor on page 442
- Configuring a Stateless Firewall Filter with a Configuration Editor on page 448
- Verifying Firewall Filter Configuration on page 465

## 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 404.
- 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 101.

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 77 and Figure 78 show the Firewall/NAT Quick Configuration main and application pages.



**Figure 77: Firewall/NAT Quick Configuration Main 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

### 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**

fxp0.0

**Trusted Interfaces**

fe-0/0/0.0

**Figure 78: 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

|  |  |  |
|--|--|--|
|  |  |  |
|--|--|--|

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 137.
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 465.
- To verify a stateful firewall filter, see “Verifying Firewall Filter Configuration” on page 465.

**Table 137: Firewall/NAT Quick Configuration Pages Summary**

| Field                                    | Function                                                                                                                                                | Your Action                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Stateful Firewall</b>                 |                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Enable Stateful Firewall                 | Enables stateful firewall filter configuration.                                                                                                         | To enable stateful firewall filter configuration, select the check box.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>Trusted Interfaces</b>                |                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 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"> <li>■ 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.</li> <li>■ 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.</li> </ul> |
| <b>Network Address Translation (NAT)</b> |                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| 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.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>Outside Applications Allowed</b>      |                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|                                          | Add or delete applications that are allowed to operate from the untrusted network to the trusted network.                                               | <p>Click <b>Add</b> to move to the Firewall/NAT Quick Configuration application page. When you have finished entering information into this page, click <b>OK</b> to save it.</p> <p>To cancel your entries, click <b>Cancel</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |

**Table 137: Firewall/NAT Quick Configuration Pages Summary (continued)**

| <b>Field</b>                       | <b>Function</b>                                                                                        | <b>Your Action</b>                                                                                                                                                                                                                                 |
|------------------------------------|--------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Application</b>                 |                                                                                                        |                                                                                                                                                                                                                                                    |
| 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.                                                                                                                                       |
| <b>Source Address</b>              |                                                                                                        |                                                                                                                                                                                                                                                    |
| 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 <b>Add</b> button, then click <b>Add</b>.</p> <p>To delete an IP address and prefix, select them in the Source Addresses and Prefixes box, then click <b>Delete</b>.</p>      |
| <b>Destination Address</b>         |                                                                                                        |                                                                                                                                                                                                                                                    |
| 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 <b>Add</b> button, then click <b>Add</b>.</p> <p>To delete an IP address and prefix, select them in the Destination Addresses and Prefixes box, then click <b>Delete</b>.</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 138.

**Table 138: Sample Stateful Firewall Filter and NAT Rules**

| Rule            | Type   | Term or Terms                                                                                                                                                                                                                                                                              |
|-----------------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| to-wan-rule     | Output | <ul style="list-style-type: none"> <li>■ <b>app-term</b>—Accepts packets from any of the applications defined by the JUNOS default group <code>junos-algs-outbound</code> application set.</li> <li>■ <b>accept-all-term</b>—Accepts packets that do not match <b>app-term</b>.</li> </ul> |
| from-wan-rule   | Input  | <ul style="list-style-type: none"> <li>■ <b>wan-src-addr-term</b>—Accepts input packets with a source prefix of <code>192.168.33.0/24</code>.</li> <li>■ <b>discard-all-term</b>—Discards all packets.</li> </ul>                                                                          |
| nat-to-wan-rule | Output | <b>private-public-term</b> —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 406.

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 139.
3. To apply the stateful firewall filter and NAT to the interface, perform the configuration tasks described in Table 140.
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 465.
  - To verify the stateful firewall filter, see “Verifying a Stateful Firewall Filter” on page 470.

**Table 139: Configuring a Stateful Firewall Filter and NAT**

| Task                                                                                                    | J-Web Configuration Editor                                                                                                                                                                                                                                                    | CLI Configuration Editor                                                                                                                                                                                    |
|---------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Stateful firewall</b> level in the configuration hierarchy.                          | In the configuration editor hierarchy, select <b>Services &gt; Stateful firewall</b> .                                                                                                                                                                                        | From the top of the configuration hierarchy, enter <code>edit services stateful-firewall</code> .                                                                                                           |
| Define <b>to-wan-rule</b> and set its match direction.                                                  | <ol style="list-style-type: none"> <li>1. Next to Rule, click <b>Add new entry</b>.</li> <li>2. In the Rule name box, type <b>to-wan-rule</b>.</li> <li>3. From the Match direction list, select <b>output</b>.</li> </ol>                                                    | Set the rule name, match direction, term name, and match condition:<br><br><code>set rule to-wan-rule match-direction output</code><br><code>term app-term from application-sets junos-algs-outbound</code> |
| Define <b>app-term</b> for the <b>to-wan-rule</b> rule.                                                 | <ol style="list-style-type: none"> <li>1. Next to Term, click <b>Add new entry</b>.</li> <li>2. In the Term name box, type <b>app-term</b>.</li> </ol>                                                                                                                        |                                                                                                                                                                                                             |
| Define the match condition for <b>app-term</b> —the default <b>junos-algs-outbound</b> application set. | <ol style="list-style-type: none"> <li>1. Next to From, click <b>Configure</b>.</li> <li>2. Next to Application sets, click <b>Add new entry</b>.</li> <li>3. In the Application set name box, type <b>junos-algs-outbound</b>.</li> <li>4. Click <b>OK</b> twice.</li> </ol> |                                                                                                                                                                                                             |
| Define an action for <b>app-term</b> .                                                                  | <ol style="list-style-type: none"> <li>1. On the Term <b>app-term</b> page, next to Then, click <b>Configure</b>.</li> <li>2. In the Designation list, select <b>Accept</b>.</li> <li>3. Click <b>OK</b> twice.</li> </ol>                                                    | Set the action:<br><br><code>set rule to-wan-rule term app-term then accept</code>                                                                                                                          |

**Table 139: Configuring a Stateful Firewall Filter and NAT (continued)**

| <b>Task</b>                                                                                                                   | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                     | <b>CLI Configuration Editor</b>                                                                                                                                                               |
|-------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Define <b>accept-all-term</b> for <b>to-wan-rule</b> .                                                                        | <ol style="list-style-type: none"> <li>1. On the Rule <b>to-wan-rule</b> page, next to Term, click <b>Add new entry</b>.</li> <li>2. In the Term name box, type <b>accept-all-term</b>.</li> </ol>                                                                                                                                    | <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 <b>accept-all-term</b> . The action is taken only if a packet does not match <b>app-term</b> .           | <ol style="list-style-type: none"> <li>1. Next to Then, click <b>Configure</b>.</li> <li>2. From the Designation list, select <b>Accept</b>.</li> <li>3. Next to Accept, select the check box.</li> <li>4. Click <b>OK</b> three times.</li> </ol>                                                                                    |                                                                                                                                                                                               |
| Define <b>from-wan-rule</b> and set its match direction.                                                                      | <ol style="list-style-type: none"> <li>1. On the Rule page, next to Rule, click <b>Add new entry</b>.</li> <li>2. In the Rule name box, type <b>from-wan-rule</b>.</li> <li>3. From the Match direction list, select <b>input</b>.</li> </ol>                                                                                         | <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 <b>wan-src-addr-term</b> for the <b>from-wan-rule</b> rule.                                                            | <ol style="list-style-type: none"> <li>1. Next to Term, click <b>Add new entry</b>.</li> <li>2. In the Term name box, type <b>wan-src-addr-term</b>.</li> </ol>                                                                                                                                                                       |                                                                                                                                                                                               |
| Define the match condition for <b>wan-src-addr-term</b> .                                                                     | <ol style="list-style-type: none"> <li>1. Next to From, click <b>Configure</b>.</li> <li>2. Next to Source address, click <b>Add new entry</b>.</li> <li>3. From the Address list, select <b>Enter Specific Value—&gt;</b>.</li> <li>4. In the Prefix box, type <b>192.168.33.0/24</b>.</li> <li>5. Click <b>OK</b> twice.</li> </ol> |                                                                                                                                                                                               |
| Define an action for <b>wan-src-addr-term</b> .                                                                               | <ol style="list-style-type: none"> <li>1. On the Term <b>wan-src-addr-term</b> page, next to Then, click <b>Configure</b>.</li> <li>2. In the Designation list, select <b>Accept</b>.</li> <li>3. Click <b>OK</b> twice.</li> </ol>                                                                                                   | <p>Set the action:</p> <p>set rule from-wan-rule term wan-src-addr-term then accept</p>                                                                                                       |
| Define <b>discard-all-term</b> for <b>from-wan-rule</b> .                                                                     | <ol style="list-style-type: none"> <li>1. On the Rule <b>from-wan-rule</b> page, next to Term, click <b>Add new entry</b>.</li> <li>2. In the Term name box, type <b>discard-all-term</b>.</li> </ol>                                                                                                                                 | <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 <b>discard-all-term</b> . The action is taken only if a packet does not match <b>wan-src-addr-term</b> . | <ol style="list-style-type: none"> <li>1. Next to Then, click <b>Configure</b>.</li> <li>2. From the Designation list, select <b>Discard</b>.</li> <li>3. Click <b>OK</b> three times.</li> </ol>                                                                                                                                     |                                                                                                                                                                                               |

**Table 139: Configuring a Stateful Firewall Filter and NAT (continued)**

| <b>Task</b>                                                           | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    | <b>CLI Configuration Editor</b>                                                                                                                                                                             |
|-----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Nat</b> level in the configuration hierarchy.      | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Services</b>.</li> <li>2. Next to NAT, click <b>Configure</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                 | From the top of the configuration hierarchy, enter <code>edit services nat</code> .                                                                                                                         |
| Define the <b>public-pool</b> address pool name and range.            | <ol style="list-style-type: none"> <li>1. Next to Pool, click <b>Add new entry</b>.</li> <li>2. In the Pool name box, type <code>public-pool</code>.</li> <li>3. From the Address choice list, select <b>Address range</b>.</li> <li>4. In the High box, type <code>10.148.2.32</code>. In the Low box, <code>10.148.2.1</code>.</li> </ol>                                                                                                                                                                                                                          | <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"> <li>1. Next to Port, click <b>Configure</b>.</li> <li>2. From the Port choice list, select <b>Automatic</b>.</li> <li>3. Click <b>OK</b> twice.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                        | <p>Configure the source port translation to be automatic:</p> <pre>set pool public-pool port automatic</pre>                                                                                                |
| Define <b>nat-to-wan-rule</b> and <b>private-public-term</b> .        | <ol style="list-style-type: none"> <li>1. On the Nat page, next to Rule, click <b>Add new entry</b>.</li> <li>2. In the Rule name box, type <code>nat-to-wan-rule</code>.</li> <li>3. From the Match direction list, select <b>output</b>.</li> <li>4. Next to Term, select <b>Add new entry</b>.</li> <li>5. In the Term name box, type <code>private-public-term</code>.</li> <li>6. Next to Then, select <b>Configure</b>.</li> <li>7. Next to Translated, select <b>Configure</b>.</li> <li>8. In the Source pool box, type <code>public-pool</code>.</li> </ol> | <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 <b>private-public-term</b> .    | <ol style="list-style-type: none"> <li>1. Next to Translation type, select the check box.</li> <li>2. Select <b>Configure</b>.</li> <li>3. From the Source list, select <b>dynamic</b>.</li> <li>4. Click <b>OK</b> five times.</li> </ol>                                                                                                                                                                                                                                                                                                                           | <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 140: Applying a Stateful Firewall Filter and NAT to an Interface**

| <b>Task</b>                                                                                                             | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                   | <b>CLI Configuration Editor</b>                                                                                                                   |
|-------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Services</b> level in the configuration hierarchy.                                                   | 1. In the configuration editor hierarchy, select <b>Services</b> .                                                                                                                                                                                                                                                                                                  | From the top of the configuration hierarchy, enter <b>edit services</b> .                                                                         |
| Define <b>wan-service-set</b> and assign the stateful firewall filter rule <b>to-wan-rule</b> to the service set.       | 1. Next to Service set, click <b>Add new entry</b> .<br>2. In the Service set name box, type <b>wan-service-set</b> .<br>3. From the Stateful firewall rules choice list, select <b>Stateful firewall rules</b> .<br>4. Next to Stateful firewall rules, click <b>Add new entry</b> .<br>5. In the Rule name box, type <b>to-wan-rule</b> .<br>6. Click <b>OK</b> . | Define the service set and assign the rule:<br><br><b>set service-set wan-service-set stateful-firewall-rules to-wan-rule</b>                     |
| Assign the stateful firewall filter rule <b>from-wan-rule</b> to the service set.                                       | 1. Next to Stateful firewall rules, click <b>Add new entry</b> .<br>2. In the Rule name box, type <b>from-wan-rule</b> .<br>3. Click <b>OK</b> .                                                                                                                                                                                                                    | Define the service set and assign the rule:<br><br><b>set service-set wan-service-set stateful-firewall-rules from-wan-rule</b>                   |
| Assign the NAT rule <b>nat-to-wan-rule</b> to the service set.                                                          | 1. From the Nat rules choice list, select <b>Nat rules</b> .<br>2. Next to Nat rules, click <b>Add new entry</b> .<br>3. In the Rule name box, type <b>nat-to-wan-rule</b> .<br>4. Click <b>OK</b> .                                                                                                                                                                | Assign the rule to the service set:<br><br><b>set service-set wan-service-set nat-rules nat-to-wan-rule</b>                                       |
| Define the service set type and virtual interface <b>sp-0/0/0</b> as the service interface for <b>wan-service-set</b> . | 1. From the Service type choice list, select <b>Interface service</b> .<br>2. Next to Interface service, click <b>Configure</b> .<br>3. In the Service interface box, type <b>sp-0/0/0</b> .<br>4. Click <b>OK</b> .                                                                                                                                                | Define the service set type and the service interface:<br><br><b>set service-set wan-service-set interface-service service-interface sp-0/0/0</b> |

**Table 140: Applying a Stateful Firewall Filter and NAT to an Interface (continued)**

| <b>Task</b>                                                                                                                                                                                                                                              | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | <b>CLI Configuration Editor</b>                                                                                                                                                                                                        |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure the <b>sp-0/0/0</b> service interface.                                                                                                                                                                                                         | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>interfaces</b>.</li> <li>2. Next to Interface, click <b>Add new entry</b>.</li> <li>3. In the Interface name box, type <b>sp-0/0/0</b>.</li> <li>4. Next to Unit, click <b>Add new entry</b>.</li> <li>5. In the Interface unit number box, type <b>0</b>.</li> <li>6. Next to Inet, select the check box.</li> <li>7. Click <b>Configure</b>.</li> <li>8. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                       | <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 <b>Inet</b> level of the T1 interface—the untrusted interface in this example—and apply <b>wan-service-set</b> to the input and output sides of the <b>t1-0/0/0</b> interface. | <ol style="list-style-type: none"> <li>1. In the configuration editor hierarchy, select <b>Interfaces &gt; t1-0/0/0 &gt; Unit &gt; 0 &gt; Family &gt; Inet</b>.</li> <li>2. Next to Service, click <b>Configure</b>.</li> <li>3. Next to Input, click <b>Configure</b>.</li> <li>4. Next to Service set, click <b>Add new entry</b>.</li> <li>5. In the Service set name box, type <b>wan-service-set</b>.</li> <li>6. Click <b>OK</b>.</li> <li>7. Next to Output, click <b>Configure</b>.</li> <li>8. Next to Service set, click <b>Add new entry</b>.</li> <li>9. In the Service set name box, type <b>wan-service-set</b>.</li> <li>10. Click <b>OK</b>.</li> </ol> | <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 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 409.

- Stateless Firewall Filter Strategies on page 449
- Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources on page 450

- Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods on page 453
- Configuring a Routing Engine Firewall Filter to Handle Fragments on page 459
- Applying a Stateless Firewall Filter to an Interface on page 464

## 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 450 and “Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods” on page 453.

### 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 459.

## 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 141 lists the terms that are configured in this sample filter.

**Table 141: 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 <b>ssh-term</b> or <b>bgp-term</b> , creates a firewall filter log and system logging records, then discards all packets. To view the log, enter the <b>show firewall log</b> operational mode command. (For more information, see “Displaying Firewall Filter Logs” on page 471.) |

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 142.
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 465.
  - To apply the firewall filter to the Routing Engine, see “Applying a Stateless Firewall Filter to an Interface” on page 464.
  - To verify the firewall filter, see “Verifying a Services, Protocols, and Trusted Sources Firewall Filter” on page 473.

**Table 142: Configuring a Protocols and Services Firewall Filter for the Routing Engine**

| <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.                                                          | In the configuration editor hierarchy, select <b>Firewall</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | From the top of the configuration hierarchy, enter <b>edit firewall</b> .                                                                                                                   |
| Define <b>protect-RE</b> and <b>ssh-term</b> , and define the protocol, destination port, and source address match conditions. | <ol style="list-style-type: none"> <li>Next to Filter, click <b>Add new entry</b>.</li> <li>In the Filter name box, type <b>protect-RE</b>.</li> <li>Next to Term, click <b>Add New Entry</b>.</li> <li>In the Rule name box, type <b>ssh-term</b>.</li> <li>Next to From, click <b>Configure</b>.</li> <li>In the Protocol choice list, select <b>Protocol</b>.</li> <li>Next to Protocol, click <b>Add new entry</b>.</li> <li>In the Value keyword list, select <b>tcp</b>.</li> <li>Click <b>OK</b>.</li> <li>In the Destination port choice list, select <b>Destination port</b>.</li> <li>Next to Destination port, click <b>Add new entry</b>.</li> <li>In the Value keyword list, select <b>ssh</b>.</li> <li>Click <b>OK</b>.</li> <li>Next to Source address, click <b>Add new entry</b>.</li> <li>In the Address box, type <b>192.168.122.0/24</b>.</li> <li>Click <b>OK</b> twice.</li> </ol> | <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 <b>ssh-term</b> .                                                                                       | <ol style="list-style-type: none"> <li>On the Term <b>ssh-term</b> page, next to Then, click <b>Configure</b>.</li> <li>In the Designation list, select <b>Accept</b>.</li> <li>Click <b>OK</b> twice.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <p>Set the actions:</p> <pre>set family inet filter protect-RE term ssh-term then accept</pre>                                                                                              |

**Table 142: Configuring a Protocols and Services Firewall Filter for the Routing Engine (continued)**

| <b>Task</b>                                                                                              | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>CLI Configuration Editor</b>                                                                                                                                                        |
|----------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Define <b>bgp-term</b> , and define the protocol, destination port, and source address match conditions. | <ol style="list-style-type: none"> <li>1. On the Filter <b>protect-RE</b> page, next to Term, click <b>Add New Entry</b>.</li> <li>2. In the Rule name box, type <b>bgp-term</b>.</li> <li>3. Next to From, click <b>Configure</b>.</li> <li>4. In the Protocol choice list, select <b>Protocol</b>.</li> <li>5. Next to Protocol, click <b>Add new entry</b>.</li> <li>6. In the Value keyword list, select <b>tcp</b>.</li> <li>7. Click <b>OK</b>.</li> <li>8. In the Destination port choice list, select <b>Destination port</b>.</li> <li>9. Next to Destination port, click <b>Add new entry</b>.</li> <li>10. In the Value keyword list, select <b>bgp</b>.</li> <li>11. Click <b>OK</b>.</li> <li>12. Next to Source address, click <b>Add new entry</b>.</li> <li>13. In the Address box, type <b>10.2.1.0/24</b>.</li> <li>14. Click <b>OK</b> twice.</li> </ol> | <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 <b>bgp-term</b> .                                                                  | <ol style="list-style-type: none"> <li>1. On the Term <b>bgp-term</b> page, next to Then, click <b>Configure</b>.</li> <li>2. In the Designation list, select <b>Accept</b>.</li> <li>3. Click <b>OK</b> twice.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <p>Set the action:</p> <pre>set family inet filter protect-RE term bgp-term then accept</pre>                                                                                          |
| Define <b>discard-rest-term</b> and its action.                                                          | <ol style="list-style-type: none"> <li>1. On the Filter <b>protect-RE</b> page, next to Term, click <b>Add New Entry</b>.</li> <li>2. In the Rule name box, type <b>discard-rest-term</b>.</li> <li>3. Next to Then, click <b>Configure</b>.</li> <li>4. Next to Log, select the check box.</li> <li>5. Next to Syslog, select the check box.</li> <li>6. In the Designation list, select <b>Discard</b>.</li> <li>7. Click <b>OK</b> four times.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                | <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 450), 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 143 lists the terms that are configured in this sample filter.

**Table 143: Sample Stateless Firewall Filter protect-RE Terms to Protect Against Floods**

| <b>Term</b>         | <b>Purpose</b>                                                                                                                                                                                                                                                                                                                                             | <b>Policer</b>                                                                                                                                                          |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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"> <li>■ Connection request packets (SYN and ACK flag bits equal 1 and 0)</li> <li>■ Connection release packets (FIN flag bit equals 1)</li> <li>■ Connection reset packets (RST flag bit equals 1)</li> </ul> | 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"> <li>■ Echo request packets</li> <li>■ Echo response packets</li> <li>■ Unreachable packets</li> <li>■ Time-exceeded packets</li> </ul>                                                                                     | 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 144.
3. To configure the prefix lists and the firewall filter, perform the configuration tasks described in Table 145.
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 465.
  - To apply the firewall filter to the Routing Engine, see “Applying a Stateless Firewall Filter to an Interface” on page 464.
  - To verify the firewall filter, see “Verifying a TCP and ICMP Flood Firewall Filter” on page 474.



**Table 144: Configuring Policers for TCP and ICMP**

| Task                                                                                                                                                                                                                                                                 | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | CLI Configuration Editor                                                                                                                                              |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Firewall</b> level in the configuration hierarchy.                                                                                                                                                                                                | In the configuration editor hierarchy, select <b>Firewall</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | From the top of the configuration hierarchy, enter <b>edit firewall</b> .                                                                                             |
| Define <b>tcp-connection-policer</b> and set its rate limits.<br><br>You can use the following abbreviations when specifying the bandwidth limit: <ul style="list-style-type: none"> <li>■ k (1000)</li> <li>■ m (1,000,000)</li> <li>■ g (1,000,000,000)</li> </ul> | <ol style="list-style-type: none"> <li>Next to Policer, click <b>Add new entry</b>.</li> <li>In the Policer name box, type <b>tcp-connection-policer</b>.</li> <li>Next to Filter specific, select the check box.</li> <li>Next to If Exceeding, select the check box and click <b>Configure</b>.</li> <li>In the Burst size limit box, type <b>15k</b>. The burst size limit can be from 1,500 through 100,000,000 bytes.</li> <li>In the Bandwidth list, select <b>Bandwidth limit</b>.</li> <li>In the Bandwidth limit box, type <b>500k</b>. The bandwidth limit can be from 32,000 through 32,000,000,000 bps.</li> <li>Click <b>OK</b>.</li> </ol> | Set the policer name and its rate limits:<br><br><pre>set policer tcp-connection-policer filter-specific if-exceeding burst-size-limit 15k bandwidth-limit 500k</pre> |
| Define the policer action for <b>tcp-connection-policer</b> .                                                                                                                                                                                                        | <ol style="list-style-type: none"> <li>On the Policer <b>tcp-connection-policer</b> page, next to Then, click <b>Configure</b>.</li> <li>Next to Discard, select the check box.</li> <li>Click <b>OK</b> twice.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                               | Set the policer action:<br><br><pre>set policer tcp-connection-policer then discard</pre>                                                                             |

**Table 144: 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.<br><br>You can use the following abbreviations when specifying the bandwidth limit:<br><br><ul style="list-style-type: none"> <li>■ k (1000)</li> <li>■ m (1,000,000)</li> <li>■ g (1,000,000,000)</li> </ul> | <ol style="list-style-type: none"> <li>On the Firewall page, next to Policer, click <b>Add new entry</b>.</li> <li>In the Policer name box, type <code>icmp-policer</code>.</li> <li>Next to Filter specific, select the check box.</li> <li>Next to If Exceeding, select the check box and click <b>Configure</b>.</li> <li>In the Burst size limit box, type <b>15k</b>. The burst size limit can be from 1,500 through 100,000,000 bytes.</li> <li>In the Bandwidth list, select <b>Bandwidth limit</b>.</li> <li>In the Bandwidth limit box, type <b>1m</b>. The bandwidth limit can be from 32,000 through 32,000,000,000 bps.</li> <li>Click <b>OK</b>.</li> </ol> | Set the policer name and its rate limits:<br><br><pre>set policer icmp-policer filter-specific if-exceeding burst-size-limit 15k bandwidth-limit 1m</pre> |
| Define the policer action for <code>icmp-policer</code> .                                                                                                                                                                                                               | <ol style="list-style-type: none"> <li>On the Policer <code>icmp-policer</code> page, next to Then, click <b>Configure</b>.</li> <li>Next to Discard, select the check box.</li> <li>Click <b>OK</b> three times.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                             | Set the policer action:<br><br><pre>set policer icmp-policer then discard</pre>                                                                           |

**Table 145: Configuring a TCP and ICMP Flood Firewall Filter for the Routing Engine**

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

**Table 145: Configuring a TCP and ICMP Flood Firewall Filter for the Routing Engine (continued)**

| Task                                                                                                         | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                   | CLI Configuration Editor                                                                                                                                                                                           |
|--------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Define the prefix list trusted-addresses.                                                                    | <ol style="list-style-type: none"> <li>Next to Prefix list, click <b>Add new entry</b>.</li> <li>In the Name box, type trusted-addresses.</li> <li>Next to Prefix list item, click <b>Add new entry</b>.</li> <li>In the Prefix box, type 192.168.122.0/24.</li> <li>Click <b>OK</b>.</li> <li>Next to Prefix list item, click <b>Add new entry</b>.</li> <li>In the Prefix box, type 10.2.1.0/24.</li> <li>Click <b>OK</b> three times.</li> </ol>                          | <p>Set the prefix list:</p> <pre>set prefix-list trusted-addresses 192.168.122.0/24  set prefix-list trusted-addresses 10.2.1.0/24</pre>                                                                           |
| Navigate to the <b>Firewall</b> level in the configuration hierarchy.                                        | In the configuration editor hierarchy, select <b>Firewall</b> .                                                                                                                                                                                                                                                                                                                                                                                                              | From the top of the configuration hierarchy, enter <b>edit firewall</b> .                                                                                                                                          |
| Define <b>protect-RE</b> and <b>tcp-connection-term</b> , and define the source prefix list match condition. | <ol style="list-style-type: none"> <li>Next to Filter, click <b>Add new entry</b>.</li> <li>In the Filter name box, type <b>protect-RE</b>.</li> <li>Next to Term, click <b>Add New Entry</b>.</li> <li>In the Rule name box, type <b>tcp-connection-term</b>.</li> <li>Next to From, click <b>Configure</b>.</li> <li>Next to Source prefix list, click <b>Add new entry</b>.</li> <li>In the Name box, type <b>trusted-addresses</b>.</li> <li>Click <b>OK</b>.</li> </ol> | <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 <b>tcp-connection-term</b> .                          | <ol style="list-style-type: none"> <li>In the TCP flags box, type <b>(syn &amp; !ack)   fin   rst</b>.</li> <li>In the Protocol choice list, select <b>Protocol</b>.</li> <li>Next to Protocol, click <b>Add new entry</b>.</li> <li>In the Value keyword list, select <b>tcp</b>.</li> <li>Click <b>OK</b>.</li> </ol>                                                                                                                                                      | <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 &amp; !ack)   fin   rst"</pre> |

**Table 145: Configuring a TCP and ICMP Flood Firewall Filter for the Routing Engine (continued)**

| <b>Task</b>                                 | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                               | <b>CLI Configuration Editor</b>                                                                                                          |
|---------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------|
| Define the actions for tcp-connection-term. | <ol style="list-style-type: none"> <li>On the Term tcp-connection-term page, next to Then, click <b>Configure</b>.</li> <li>In the Policier box, type tcp-connection-policer.</li> <li>In the Designation list, select <b>Accept</b>.</li> <li>Click <b>OK</b> twice.</li> </ol>                                                                                                                                                | <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"> <li>On the Filter protect-RE page, next to Term, click <b>Add New Entry</b>.</li> <li>In the Rule name box, type icmp-term.</li> <li>Next to From, click <b>Configure</b>.</li> <li>In the Protocol choice list, select <b>Protocol</b>.</li> <li>Next to Protocol, click <b>Add new entry</b>.</li> <li>In the Value keyword list, select <b>icmp</b>.</li> <li>Click <b>OK</b>.</li> </ol> | <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 145: Configuring a TCP and ICMP Flood Firewall Filter for the Routing Engine (continued)**

| <b>Task</b>                            | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <b>CLI Configuration Editor</b>                                                                                                                                           |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Define the ICMP type match conditions. | <ol style="list-style-type: none"> <li>1. In the Icmp type choice list, select <b>Icmp type</b>.</li> <li>2. Next to Icmp type, click <b>Add new entry</b>.</li> <li>3. In the Value keyword list, select <b>echo-request</b>.</li> <li>4. Click <b>OK</b>.</li> <li>5. Next to Icmp type, click <b>Add new entry</b>.</li> <li>6. In the Value keyword list, select <b>echo-reply</b>.</li> <li>7. Click <b>OK</b>.</li> <li>8. Next to Icmp type, click <b>Add new entry</b>.</li> <li>9. In the Value keyword list, select <b>unreachable</b>.</li> <li>10. Click <b>OK</b>.</li> <li>11. Next to Icmp type, click <b>Add new entry</b>.</li> <li>12. In the Value keyword list, select <b>time-exceeded</b>.</li> <li>13. Click <b>OK</b>.</li> </ol> | <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 icmp-term.      | <ol style="list-style-type: none"> <li>1. On the <b>icmp-term</b> page, next to Then, click <b>Configure</b>.</li> <li>2. In the Count box, type <b>icmp-counter</b>.</li> <li>3. In the Policer box, type <b>icmp-policer</b>.</li> <li>4. In the Designation list, select <b>Accept</b>.</li> <li>5. Click <b>OK</b> four times.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                             | <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 146 lists the terms that are configured in this sample filter.

**Table 146: 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 147.
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 465.
  - To apply the firewall filter to the Routing Engine, see “Applying a Stateless Firewall Filter to an Interface” on page 464.
  - To verify the firewall filter, see “Verifying a Firewall Filter That Handles Fragments” on page 475.

**Table 147: Configuring a Fragments Firewall Filter for the Routing Engine**

| <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.                                                                                           | In the configuration editor hierarchy, select <b>Firewall</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | From the top of the configuration hierarchy, enter <b>edit firewall</b> .                                                                                              |
| Define <b>fragment-RE</b> and <b>small-offset-term</b> , and define the fragment offset match condition.<br><br>The fragment offset can be from 1 through 8191. | <ol style="list-style-type: none"> <li>Next to Filter, click <b>Add new entry</b>.</li> <li>In the Filter name box, type <b>fragment-RE</b>.</li> <li>Next to Term, click <b>Add New Entry</b>.</li> <li>In the Rule name box, type <b>small-offset-term</b>.</li> <li>Next to From, click <b>Configure</b>.</li> <li>In the Fragment offset choice list, select <b>Fragment offset</b>.</li> <li>Next to Fragment offset, select <b>Add New Entry</b>.</li> <li>In the Range box, type <b>1-5</b>.</li> <li>Click <b>OK</b> twice.</li> </ol> | <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 <b>small-offset-term</b> .                                                                                                                | <ol style="list-style-type: none"> <li>On the Term <b>small-offset-term</b> page, next to Then, click <b>Configure</b>.</li> <li>Next to Syslog, select the check box.</li> <li>In the Designation list, select <b>Discard</b>.</li> <li>Click <b>OK</b> twice.</li> </ol>                                                                                                                                                                                                                                                                     | <p>Set the action:</p> <pre>set family inet filter fragment-RE term small-offset-term then syslog discard</pre>                                                        |

**Table 147: Configuring a Fragments Firewall Filter for the Routing Engine (continued)**

| <b>Task</b>                                                                                                                   | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>CLI Configuration Editor</b>                                                                                                                                                                                                     |
|-------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Define <b>not-fragmented-term</b> , and define the fragment, protocol, destination port, and source address match conditions. | <ol style="list-style-type: none"> <li>1. On the Filter <b>fragment-RE</b> page, next to Term, click <b>Add New Entry</b>.</li> <li>2. In the Term name box, type <b>not-fragmented-term</b>.</li> <li>3. Next to From, click <b>Configure</b>.</li> <li>4. In the Fragment flags box, type <b>0x0</b>.</li> <li>5. In the Fragment offset choice list, select <b>Fragment offset</b>.</li> <li>6. Next to Fragment offset, select <b>Add New Entry</b>.</li> <li>7. In the Range box, type <b>0</b>.</li> <li>8. Click <b>OK</b>.</li> <li>9. In the Protocol choice list, select <b>Protocol</b>.</li> <li>10. Next to Protocol, click <b>Add new entry</b>.</li> <li>11. In the Value keyword list, select <b>tcp</b>.</li> <li>12. Click <b>OK</b>.</li> <li>13. In the Destination port choice list, select <b>Destination port</b>.</li> <li>14. Next to Destination port, click <b>Add new entry</b>.</li> <li>15. In the Value keyword list, select <b>bgp</b>.</li> <li>16. Click <b>OK</b>.</li> <li>17. Next to Source address, click <b>Add new entry</b>.</li> <li>18. In the Address box, type <b>10.2.1.0/24</b>.</li> <li>19. Click <b>OK</b> twice.</li> </ol> | <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 <b>not-fragmented-term</b> .                                                                            | <ol style="list-style-type: none"> <li>1. On the Term <b>not-fragmented-term</b> page, next to Then, click <b>Configure</b>.</li> <li>2. In the Designation list, select <b>Accept</b>.</li> <li>3. Click <b>OK</b> twice.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <p>Set the action:</p> <pre>set family inet filter fragment-RE term not-fragmented-term then accept</pre>                                                                                                                           |



**Table 147: 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"> <li>On the Filter <code>fragment-RE</code> page, next to Term, click <b>Add New Entry</b>.</li> <li>In the Rule name box, type <code>first-fragment-term</code>.</li> <li>Next to From, click <b>Configure</b>.</li> <li>Next to First fragment, select the check box.</li> <li>In the Protocol choice list, select <b>Protocol</b>.</li> <li>Next to Protocol, click <b>Add new entry</b>.</li> <li>In the Value keyword list, select <b>tcp</b>.</li> <li>Click <b>OK</b>.</li> <li>In the Destination port choice list, select <b>Destination port</b>.</li> <li>Next to Destination port, click <b>Add new entry</b>.</li> <li>In the Value keyword list, select <b>bgp</b>.</li> <li>Click <b>OK</b>.</li> <li>Next to Source address, click <b>Add new entry</b>.</li> <li>In the Address box, type <code>10.2.1.0/24</code>.</li> <li>Click <b>OK</b> twice.</li> </ol> | <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"> <li>On the Term <code>first-fragment-term</code> page, next to Then, click <b>Configure</b>.</li> <li>In the Designation list, select <b>Accept</b>.</li> <li>Click <b>OK</b> twice.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | <p>Set the action:</p> <pre>set family inet filter fragment-RE term first-fragment-term then accept</pre>                                                                                                     |

**Table 147: Configuring a Fragments Firewall Filter for the Routing Engine (continued)**

| <b>Task</b>                                                                | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>CLI Configuration Editor</b>                                                                                                                    |
|----------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|
| Define <code>fragment-term</code> and define the fragment match condition. | <ol style="list-style-type: none"> <li>1. On the Filter <code>fragment-RE</code> page, next to Term, click <b>Add New Entry</b>.</li> <li>2. In the Rule name box, type <code>fragment-term</code>.</li> <li>3. Next to From, click <b>Configure</b>.</li> <li>4. In the Fragment offset choice list, select <b>Fragment offset</b>.</li> <li>5. Next to Fragment offset, select <b>Add New Entry</b>.</li> <li>6. In the Range box, type <code>6-8191</code>.</li> <li>7. Click <b>OK</b> twice.</li> </ol> | Set the term name and define match conditions:<br><br><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"> <li>1. On the Term <code>fragment-term</code> page, next to Then, click <b>Configure</b>.</li> <li>2. In the Designation list, select <b>Accept</b>.</li> <li>3. Click <b>OK</b> four times.</li> </ol>                                                                                                                                                                                                                                                                   | Set the action:<br><br><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 148.
2. If you are finished configuring the router, commit the configuration.

**Table 148: Applying a Firewall Filter to the Routing Engine Interface**

| <b>Task</b>                                                             | <b>J-Web Configuration Editor</b>                                                                                                                                                 | <b>CLI Configuration Editor</b>                                                 |
|-------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Navigate to the <b>Inet</b> level in the configuration hierarchy.       | In the configuration editor hierarchy, select <b>Interfaces &gt; lo0 &gt; Unit &gt; 0 &gt; Family &gt; Inet</b> .                                                                 | From the top of the configuration hierarchy, apply the filter to the interface: |
| Apply <b>protect-RE</b> as an input filter to the <b>lo0</b> interface. | <ol style="list-style-type: none"> <li>Next to Filter, click <b>Configure</b>.</li> <li>In the Input box, type <b>protect-RE</b>.</li> <li>Click <b>OK</b> five times.</li> </ol> | 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 465
- Verifying a Stateful Firewall Filter on page 470
- Displaying Firewall Filter Logs on page 471
- Displaying Firewall Filter Statistics on page 472
- Verifying a Services, Protocols, and Trusted Sources Firewall Filter on page 473
- Verifying a TCP and ICMP Flood Firewall Filter on page 474
- Verifying a Firewall Filter That Handles Fragments on page 475

### Displaying Firewall Filter Configurations

|                |                                                                                                                                                |
|----------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | Verify the configuration of the firewall filter. You can analyze the flow of the firewall filter terms by displaying the entire configuration. |
| <b>Action</b>  | From the J-Web interface, select <b>Configuration &gt; View and Edit &gt; View Configuration Text</b> .                                        |

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 442
- Stateless `protect-RE` filter configured in “Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources” on page 450
- Stateless `protect-RE` filter configured in “Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods” on page 453
- Stateless `fragment-RE` filter configured in “Configuring a Routing Engine Firewall Filter to Handle Fragments” on page 459

#### 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 442.

- 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 450 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 453 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 450.

**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 453.

**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 459.

**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 21

# 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 149.

**Table 149: 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 Network Interfaces and Class of Service Configuration Guide*.

- Before You Begin on page 478
- Configuring CoS with DiffServ with a Configuration Editor on page 478
- Verifying a DiffServ Configuration on page 508

## 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 397.
- 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 Services Router Configuration Tools” on page 3.

- Configuring a Policer for a Firewall Filter (*Required*) on page 479
- Configuring and Applying a Firewall Filter for a Multifield Classifier (*Required*) on page 480
- Assigning Forwarding Classes to Output Queues (*Required*) on page 484
- Configuring and Applying Rewrite Rules (*Required*) on page 485
- Configuring and Applying Behavior Aggregate Classifiers (*Required*) on page 490
- Configuring RED Drop Profiles for Assured Forwarding Congestion Control (*Required*) on page 494
- Configuring Schedulers (*Optional*) on page 496
- Configuring and Applying Scheduler Maps (*Optional*) on page 500



- Configuring and Applying Virtual Channels (Optional) on page 503
- Configuring and Applying Adaptive Shaping (Optional) on page 507

**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 437 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 150.
3. Go on to “Configuring and Applying a Firewall Filter for a Multifield Classifier (Required)” on page 480.

**Table 150: Configuring a Policer for a Firewall Filter**

| Task                                                                  | J-Web Configuration Editor                                                                                                                                            | CLI Configuration Editor                                                |
|-----------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| Navigate to the <b>Firewall</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Firewall</b> .                                                                                                       | From the top of the configuration hierarchy, enter<br><br>edit firewall |
| Create and name the policer for expedited forwarding.                 | <div>1. Click <b>Add new entry</b> next to Policer.</div> <div>2. In the Policer name box, type a name for the EF policer—for example, <code>ef-policer</code>.</div> | Enter<br><br>edit policer ef-policer                                    |

**Table 150: Configuring a Policer for a Firewall Filter (continued)**

| <b>Task</b>                                                                          | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>CLI Configuration Editor</b>                                                                       |
|--------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------|
| Enter the burst limit and bandwidth for the policer.                                 | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to If exceeding.</li> <li>2. In the Burst size limit box, type a limit for the burst size allowed—for example, 2k.</li> <li>3. From the <b>Bandwidth</b> list, select a limit or percentage—for example, <b>bandwidth-percent</b>.</li> <li>4. In the Bandwidth percent box, type a percentage for the bandwidth allowed for this type of traffic—for example, 10.</li> <li>5. Click <b>OK</b>.</li> </ol> | <p>Enter</p> <p>set if-exceeding burst-limit-size 2k</p> <p>set if-exceeding bandwidth-percent 10</p> |
| Enter the loss priority for packets exceeding the limits established by the policer. | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to Then.</li> <li>2. From the Loss priority list, select <b>high</b>.</li> <li>3. Click <b>OK</b> three times.</li> </ol>                                                                                                                                                                                                                                                                                  | <p>Enter</p> <p>set then loss-priority high</p>                                                       |

### **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 151.

**Table 151: Sample mf-classifier Firewall Filter Terms**

| Rule (Term)          | Purpose                                                                                                                                                                                                                    | Contents                                                                                                                     |
|----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------|
| assured forwarding   | Detects packets destined for <b>192.168.44.55</b> , assigns them to an assured forwarding class, and gives them a low likelihood of being dropped.                                                                         | Match condition: destination address <b>192.168.44.55</b><br>Forwarding class: <b>af-class</b><br>Loss priority: low         |
| expedited-forwarding | Detects packets destined for <b>192.168.66.77</b> , 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 479. | Match condition: destination address <b>192.168.66.77</b><br>Forwarding class: <b>ef-class</b><br>Policer: <b>ef-policer</b> |
| network control      | Detects packets with a network control precedence and forwards them to the network control class.                                                                                                                          | Match condition: precedence <b>net-control</b><br>Forwarding class: <b>nc-class</b>                                          |
| best-effort-data     | Detects all other packets and assigns them to the best effort class.                                                                                                                                                       | Forwarding class: <b>be-class</b>                                                                                            |

For more information about firewalls filters see “Configuring Firewall Filters and NAT” on page 437 and the *JUNOS Policy Framework Configuration Guide*.

To configure a firewall filter for a multfield 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 152.
3. Go on to “Assigning Forwarding Classes to Output Queues (Required)” on page 484.

**Table 152: Configuring and Applying a Firewall Filter for a Multifield Classifier**

| Task                                                                  | J-Web Configuration Editor                                                                                                                                                                                                                                                           | CLI Configuration Editor                                                |
|-----------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------|
| Navigate to the <b>Firewall</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Firewall</b> .                                                                                                                                                                                                                      | From the top of the configuration hierarchy, enter<br><br>edit firewall |
| Create and name the multifield classifier filter.                     | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Filter.</li> <li>2. In the Filter name box, type a name for the multifield classifier filter—for example, <b>mf-classifier</b>.</li> <li>3. Select the check box next to Interface specific.</li> </ol> | Enter<br><br>edit filter mf-classifier<br>set interface-specific        |

**Table 152: Configuring and Applying a Firewall Filter for a Multifield Classifier (continued)**

| <b>Task</b>                                                          | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                | <b>CLI Configuration Editor</b>                                                                                                                                                                                              |
|----------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create and name the term for the assured forwarding traffic class.   | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Term.</li> <li>2. In the Rule name box, type a name for the assured forwarding term—for example, <b>assured-forwarding</b>.</li> </ol>                                                                                                                                              | <p>Enter</p> <p><b>edit term assured-forwarding</b></p>                                                                                                                                                                      |
| Create the match condition for the assured forwarding traffic class. | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to From.</li> <li>2. Click <b>Add new entry</b> next to Destination address.</li> <li>3. In the Address box, type the destination address for assured forwarding traffic in dotted decimal notation—for example, <b>192.168.44.55</b>.</li> <li>4. Click <b>OK</b> three times.</li> </ol> | <p>Enter</p> <p><b>set from destination-address 192.168.44.55</b></p>                                                                                                                                                        |
| Create the priority for the assured forwarding traffic class.        | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to Then.</li> <li>2. In the Forwarding class box, type the forwarding class for assured forwarding DiffServ traffic—for example, <b>af-class</b>.</li> <li>3. From the Loss priority list, select <b>low</b>.</li> <li>4. Click <b>OK</b> twice.</li> </ol>                                | <p>From the top of the configuration hierarchy, enter</p> <p><b>edit firewall filter mf-classifier term assured-forwarding</b></p> <p><b>set then forwarding-class af-class</b></p> <p><b>set then loss-priority low</b></p> |
| Create and name the term for the expedited forwarding traffic class. | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Term.</li> <li>2. In the Rule name box, type a name for the expedited term—for example, <b>expedited-forwarding</b>.</li> </ol>                                                                                                                                                     | <p>Enter</p> <p><b>edit term expedited-forwarding</b></p>                                                                                                                                                                    |
| Create the match condition for the assured forwarding traffic class. | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to From.</li> <li>2. Click <b>Add new entry</b> next to Destination address.</li> <li>3. In the Address box, type the destination address for assured forwarding traffic in dotted decimal notation—for example, <b>192.168.66.77</b>.</li> <li>4. Click <b>OK</b> twice.</li> </ol>       | <p>Enter</p> <p><b>set from destination-address 192.168.66.77</b></p>                                                                                                                                                        |

**Table 152: Configuring and Applying a Firewall Filter for a Multifield Classifier (continued)**

| <b>Task</b>                                                                                                                              | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>CLI Configuration Editor</b>                                                                                                                                                                                                 |
|------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create the priority and apply the policer for the expedited forwarding traffic class.                                                    | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to Then.</li> <li>2. In the Forwarding class box, type the forwarding class for expedited forwarding DiffServ traffic—for example, <b>ef-class</b>.</li> <li>3. In the Policer box, type the name of the EF policer previously configured for expedited forwarding DiffServ traffic—<b>ef-policer</b>.<br/><br/>(See “Configuring a Policer for a Firewall Filter (Required)” on page 479.)</li> <li>4. Click <b>OK</b> twice.</li> </ol> | <p>From the top of the configuration hierarchy, enter</p> <p><b>edit firewall filter mf-classifier term expedited-forwarding</b></p> <p><b>set then forwarding-class ef-class</b></p> <p><b>set then policer ed-policer</b></p> |
| Create and name the term for the network control traffic class.                                                                          | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Term.</li> <li>2. In the Rule name box, type a name for the network control term—for example, <b>network-control</b>.</li> </ol>                                                                                                                                                                                                                                                                                                   | <p>Enter</p> <p><b>edit term network-control</b></p>                                                                                                                                                                            |
| Create the match condition for the network control traffic class.                                                                        | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to From.</li> <li>2. From the Precedence choice list, select <b>Precedence</b>.</li> <li>3. Click <b>Add new entry</b> next to Precedence.</li> <li>4. From the Value keyword list, select <b>net-control</b>.</li> <li>5. Click <b>OK</b> twice.</li> </ol>                                                                                                                                                                              | <p>Enter</p> <p><b>set from traffic-class net-control</b></p>                                                                                                                                                                   |
| Create the forwarding class for the network control traffic class.                                                                       | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to Then.</li> <li>2. In the Forwarding class box, type the forwarding class for network control traffic—for example, <b>nc-class</b>.</li> <li>3. Click <b>OK</b> twice.</li> </ol>                                                                                                                                                                                                                                                       | <p>From the top of the configuration hierarchy, enter</p> <p><b>edit firewall filter mf-classifier term network-control</b></p> <p><b>set then forwarding-class nc-class</b></p>                                                |
| Create and name the term for the best-effort traffic class.                                                                              | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Term.</li> <li>2. In the Rule name box, type a name for the best-effort term—for example, <b>best-effort-data</b>.</li> </ol>                                                                                                                                                                                                                                                                                                      | <p>Enter</p> <p><b>edit term best-effort-data</b></p>                                                                                                                                                                           |
| Create the forwarding class for the best-effort traffic class. (Because this is the last term in the filter, it has no match condition.) | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to Then.</li> <li>2. In the Forwarding class box, type the forwarding class for best effort traffic—for example, <b>be-class</b>.</li> <li>3. Click <b>OK</b> four times.</li> </ol>                                                                                                                                                                                                                                                      | <p>From the top of the configuration hierarchy, enter</p> <p><b>set then forwarding-class be-class</b></p>                                                                                                                      |

**Table 152: Configuring and Applying a Firewall Filter for a Multifield Classifier (continued)**

| 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 |
| Apply the multifield classifier firewall filter as an input filter on the customer-facing or host-facing interfaces. | <ol style="list-style-type: none"> <li>Click the Interface and Unit of each interface needing the filter—for example, <b>fe-0/0/0</b>, unit <b>0</b>.</li> <li>Click <b>Configure</b> next to Inet.</li> <li>Click <b>Configure</b> next to Filter.</li> <li>In the Input box, type the name of the previously configured filter—for example, <b>mf-classifier</b>.</li> <li>Click <b>OK</b> five times.</li> </ol> | Enter<br><br>set fe-0/0/0 unit 0 family inet filter input mf-classifier   |

### 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 153.

**Table 153: 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 480.

To assign forwarding classes to output queues for the 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 154.
- Go on to “Configuring and Applying Rewrite Rules (Required)” on page 485.

**Table 154: Assigning Forwarding Classes to Output Queues**

| <b>Task</b>                                                                   | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                             | <b>CLI Configuration Editor</b>                                                 |
|-------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Navigate to the <b>Class of service</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Class of service</b> .                                                                                                                                                                                                                                                                       | From the top of the configuration hierarchy, enter<br><br>edit class-of-service |
| Assign best-effort traffic to queue 0.                                        | <ol style="list-style-type: none"> <li>Click <b>Configure</b> next to Forwarding classes.</li> <li>Click <b>Add new entry</b> next to Queue.</li> <li>In the Queue num box, type <b>0</b>.</li> <li>In the Class name box, type the previously configured name of the best-effort class—<b>be-class</b>.</li> <li>Click <b>OK</b>.</li> </ol> | Enter<br><br>set forwarding-classes queue 0 be-class                            |
| Assign expedited forwarding traffic to queue 1.                               | <ol style="list-style-type: none"> <li>Click <b>Add new entry</b> next to Queue.</li> <li>In the Queue num box, type <b>1</b>.</li> <li>In the Class name box, type the previously configured name of the expedited forwarding class—<b>ef-class</b>.</li> <li>Click <b>OK</b>.</li> </ol>                                                    | Enter<br><br>set forwarding-classes queue 1 ef-class                            |
| Assign assured forwarding traffic to queue 2.                                 | <ol style="list-style-type: none"> <li>Click <b>Add new entry</b> next to Queue.</li> <li>In the Queue num box, type <b>2</b>.</li> <li>In the Class name box, type the previously configured name of the assured forwarding class—<b>af-class</b>.</li> <li>Click <b>OK</b>.</li> </ol>                                                      | Enter<br><br>set forwarding-classes queue 2 af-class                            |
| Assign network control traffic to queue 3.                                    | <ol style="list-style-type: none"> <li>Click <b>Add new entry</b> next to Queue.</li> <li>In the Queue num box, type <b>3</b>.</li> <li>In the Class name box, type the previously configured name of the expedited forwarding class—<b>nc-class</b>.</li> <li>Click <b>OK</b> twice.</li> </ol>                                              | Enter<br><br>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 155.

**Table 155: 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 |
| 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 156.
3. Go on to “Configuring and Applying Behavior Aggregate Classifiers (Required)” on page 490.

**Table 156: Configuring and Applying Rewrite Rules**

| Task                                                                          | J-Web Configuration Editor                                                                                                                                                                                                                                       | CLI Configuration Editor                                                                     |
|-------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|
| Navigate to the <b>Class of service</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Class of service</b> .                                                                                                                                                                                          | From the top of the configuration hierarchy, enter<br><br><code>edit class-of-service</code> |
| Configure rewrite rules for DiffServ CoS.                                     | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to Rewrite rules.</li> <li>2. Click <b>Add new entry</b> next to Dscp.</li> <li>3. In the Name box, type the name of the rewrite rules—for example, <code>rewrite-dscps</code>.</li> </ol> | Enter<br><br><code>edit rewrite-rules dscp rewrite-dscps</code>                              |



**Table 156: Configuring and Applying Rewrite Rules (continued)**

| <b>Task</b>                                           | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <b>CLI Configuration Editor</b>                                                                                                                                 |
|-------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure best-effort forwarding class rewrite rules. | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Forwarding class.</li> <li>2. In the Class name box, type the name of the previously configured best-effort forwarding class—<b>be-class</b>.</li> <li>3. Click <b>Add new entry</b> next to Loss priority.</li> <li>4. From the Loss val list, select <b>low</b>.</li> <li>5. In the Code point box, type the value of the low-priority code point for best-effort traffic—for example, 000000.</li> <li>6. Click <b>OK</b>.</li> <li>7. Click <b>Add new entry</b> next to Loss priority.</li> <li>8. From the Loss val list, select <b>high</b>.</li> <li>9. In the Code point box, type the value of the high-priority code point for best-effort traffic—for example, 000001.</li> <li>10. Click <b>OK</b> twice.</li> </ol> | <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 156: Configuring and Applying Rewrite Rules (continued)**

| <b>Task</b>                                         | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>CLI Configuration Editor</b>                                                                                                                                 |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure expedited forwarding class rewrite rules. | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Forwarding class.</li> <li>2. In the Class name box, type the name of the previously configured expedited forwarding class—<b>ef-class</b>.</li> <li>3. Click <b>Add new entry</b> next to Loss priority.</li> <li>4. From the Loss val list, select <b>low</b>.</li> <li>5. In the Code point box, type the value of the low-priority code point for expedited forwarding traffic—for example, <b>101110</b>.</li> <li>6. Click <b>OK</b>.</li> <li>7. Click <b>Add new entry</b> next to Loss priority.</li> <li>8. From the Loss val list, select <b>high</b>.</li> <li>9. In the Code point box, type the value of the high-priority code point for expedited forwarding traffic—for example, <b>101111</b>.</li> <li>10. Click <b>OK</b> twice.</li> </ol> | <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 156: Configuring and Applying Rewrite Rules (continued)**

| <b>Task</b>                                       | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>CLI Configuration Editor</b>                                                                                                                                 |
|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure assured forwarding class rewrite rules. | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Forwarding class.</li> <li>2. In the Class name box, type the name of the previously configured assured forwarding class—<b>af-class</b>.</li> <li>3. Click <b>Add new entry</b> next to Loss priority.</li> <li>4. From the Loss val list, select <b>low</b>.</li> <li>5. In the Code point box, type the value of the low-priority code point for assured forwarding traffic—for example, <b>001010</b>.</li> <li>6. Click <b>OK</b>.</li> <li>7. Click <b>Add new entry</b> next to Loss priority.</li> <li>8. From the Loss val list, select <b>high</b>.</li> <li>9. In the Code point box, type the value of the high-priority code point for assured forwarding traffic—for example, <b>001100</b>.</li> <li>10. Click <b>OK</b> twice.</li> </ol> | <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 156: Configuring and Applying Rewrite Rules (continued)**

| <b>Task</b>                                    | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>CLI Configuration Editor</b>                                                                                                                                 |
|------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure network control class rewrite rules. | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Forwarding class.</li> <li>2. In the Class name box, type the name of the previously configured network control forwarding class—<b>nc-class</b>.</li> <li>3. Click <b>Add new entry</b> next to Loss priority.</li> <li>4. From the Loss val list, select <b>low</b>.</li> <li>5. In the Code point box, type the value of the low-priority code point for network control traffic—for example, <b>110000</b>.</li> <li>6. Click <b>OK</b>.</li> <li>7. Click <b>Add new entry</b> next to Loss priority.</li> <li>8. From the Loss val list, select <b>high</b>.</li> <li>9. In the Code point box, type the value of the high-priority code point for network control traffic—for example, <b>110001</b>.</li> <li>10. Click <b>OK</b> twice.</li> </ol> | <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"> <li>1. Click <b>Add new entry</b> next to Interfaces</li> <li>2. In the Interface name box, type the name of the interface—for example, <b>fe-0/0/0</b>.</li> <li>3. In the Rewrite rules box, type the name of the previously configured rewrite rules—<b>rewrite-dscps</b>.</li> <li>4. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <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 157, to incoming packets in the four forwarding classes.

**Table 157: 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 158.
3. Go on to “Configuring RED Drop Profiles for Assured Forwarding Congestion Control (Required)” on page 494.

**Table 158: Configuring and Applying Behavior Aggregate Classifiers**

| Task                                                                          | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                     | CLI Configuration Editor                                                        |
|-------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Navigate to the <b>Class of service</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Class of service</b> .                                                                                                                                                                                                                                                                        | From the top of the configuration hierarchy, enter<br><br>edit class-of-service |
| Configure BA classifiers for DiffServ CoS.                                    | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to Classifiers.</li> <li>2. Click <b>Add new entry</b> next to Dscp.</li> <li>3. In the Name box, type the name of the BA classifier—for example, <b>ba-classifier</b>.</li> <li>4. In the Import box, type the name of the default DSCP map, <b>default</b>.</li> </ol> | Enter<br><br>edit classifiers dscp ba-classifier<br><br>set import default      |

**Table 158: Configuring and Applying Behavior Aggregate Classifiers (continued)**

| <b>Task</b>                                          | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>CLI Configuration Editor</b>                                                             |
|------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| Configure a best-effort forwarding class classifier. | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Forwarding class.</li> <li>2. In the Class name box, type the name of the previously configured best-effort forwarding class—<b>be-class</b>.</li> <li>3. Click <b>Add new entry</b> next to Loss priority.</li> <li>4. From the Loss val list, select <b>high</b>.</li> <li>5. In the Code point box, type the value of the high-priority code point for best-effort traffic—for example, <b>00001</b>.</li> <li>6. Click <b>OK</b> three times.</li> </ol>         | <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"> <li>1. Click <b>Add new entry</b> next to Forwarding class.</li> <li>2. In the Class name box, type the name of the previously configured expedited forwarding class—<b>ef-class</b>.</li> <li>3. Click <b>Add new entry</b> next to Loss priority.</li> <li>4. From the Loss val list, select <b>high</b>.</li> <li>5. In the Code point box, type the value of the high-priority code point for expedited forwarding traffic—for example, <b>101111</b>.</li> <li>6. Click <b>OK</b> three times.</li> </ol> | <p>Enter</p> <pre>set forwarding-class ef-class loss-priority high code points 101111</pre> |

**Table 158: Configuring and Applying Behavior Aggregate Classifiers (continued)**

| <b>Task</b>                                       | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | <b>CLI Configuration Editor</b>                                                             |
|---------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| Configure an assured forwarding class classifier. | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Forwarding class.</li> <li>2. In the Class name box, type the name of the previously configured assured forwarding class—<b>af-class</b>.</li> <li>3. Click <b>Add new entry</b> next to Loss priority.</li> <li>4. From the Loss val list, select <b>high</b>.</li> <li>5. In the Code point box, type the value of the high-priority code point for assured forwarding traffic—for example, <b>001100</b>.</li> <li>6. Click <b>OK</b> three times.</li> </ol>      | <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"> <li>1. Click <b>Add new entry</b> next to Forwarding class.</li> <li>2. In the Class name box, type the name of the previously configured network control forwarding class—<b>nc-class</b>.</li> <li>3. Click <b>Add new entry</b> next to Loss priority.</li> <li>4. From the Loss val list, select <b>high</b>.</li> <li>5. In the Code point box, type the value of the high-priority code point for network control traffic—for example, <b>110001</b>.</li> <li>6. Click <b>OK</b> three times.</li> </ol> | <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"> <li>1. Click <b>Add new entry</b> next to Interfaces.</li> <li>2. In the Interface name box, type the name of the interface—for example, <b>fe-0/0/0</b>.</li> <li>3. In the Classifiers box, type the name of the previously configured BA classifier—<b>ba-classifier</b>.</li> <li>4. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                            | <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 159.

**Table 159: 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 160.
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 496.
  - To apply rules to logical interfaces, see “Configuring and Applying Virtual Channels (Optional)” on page 503.
  - To use adaptive shapers to limit bandwidth for Frame Relay, see “Configuring and Applying Adaptive Shaping (Optional)” on page 507.
  - To check the configuration, see “Verifying a DiffServ Configuration” on page 508.



**Table 160: Configuring RED Drop Profiles for Assured Forwarding Congestion Control**

| <b>Task</b>                                                                   | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>CLI Configuration Editor</b>                                                                                     |
|-------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Class of service</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Class of service</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | From the top of the configuration hierarchy, enter<br><br>edit class-of-service                                     |
| Configure the lower drop probability for normal, non-PLP traffic.             | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Drop profiles.</li> <li>2. In the Profile name box, type the name of the drop profile—for example, <b>af-normal</b>.</li> <li>3. Click <b>Configure</b> next to Interpolate.</li> <li>4. Click <b>Add new entry</b> next to Drop probability.</li> <li>5. In the Value box, type a number for the first drop point—for example, <b>0</b>.</li> <li>6. Click <b>OK</b>.</li> <li>7. Click <b>Add new entry</b> next to Drop probability again.</li> <li>8. In the Value box, type a number for the next drop point—for example, <b>100</b>.</li> <li>9. Click <b>OK</b>.</li> </ol> | Enter<br><br>edit drop-profiles af-normal interpolate<br><br>set drop-probability 0<br><br>set drop-probability 100 |
| Configure a queue fill level for the lower non-PLP drop probability.          | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Fill level.</li> <li>2. In the Value box, type a number for the first fill level—for example, <b>95</b>.</li> <li>3. Click <b>OK</b>.</li> <li>4. In the Value box, type a number for the next fill level—for example, <b>100</b>.</li> <li>5. Click <b>OK</b> three times.</li> </ol>                                                                                                                                                                                                                                                                                             | Enter<br><br>set fill-level 95<br><br>set fill-level 100                                                            |

**Table 160: Configuring RED Drop Profiles for Assured Forwarding Congestion Control (continued)**

| <b>Task</b>                                                       | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <b>CLI Configuration Editor</b>                                                                                               |
|-------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| Configure the higher drop probability for PLP traffic.            | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Drop profiles.</li> <li>2. In the Profile name box, type the name of the drop profile—for example, <b>af-with-plp</b>.</li> <li>3. Click <b>Configure</b> next to Interpolate.</li> <li>4. Click <b>Add new entry</b> next to Drop probability.</li> <li>5. In the Value box, type a number for the first drop point—for example, <b>95</b>.</li> <li>6. Click <b>OK</b>.</li> <li>7. In the Value box, type a number for the next drop point—for example, <b>100</b>.</li> <li>8. Click <b>OK</b>.</li> </ol> | <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"> <li>1. Click <b>Add new entry</b> next to Fill level.</li> <li>2. In the Value box, type a number for the first fill level—for example, <b>80</b>.</li> <li>3. Click <b>OK</b>.</li> <li>4. In the Value box, type a number for the next fill level—for example, <b>95</b>.</li> <li>5. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                      | <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 161.

**Table 161: Sample Schedulers**

| <b>Scheduler</b> | <b>For CoS Traffic Type</b>  | <b>Assigned Priority</b> | <b>Allocated Portion of Queue Buffer</b> | <b>Assigned Bandwidth (Transmit Rate)</b> |
|------------------|------------------------------|--------------------------|------------------------------------------|-------------------------------------------|
| 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 162.
3. Go on to “Configuring and Applying Scheduler Maps (Optional)” on page 500.

**Table 162: Configuring Schedulers**

| <b>Task</b>                                                                   | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                              | <b>CLI Configuration Editor</b>                                                 |
|-------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Navigate to the <b>Class of service</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Class of service</b> .                                                                                                                                                                                                                                                                                                                                                                        | From the top of the configuration hierarchy, enter<br><br>edit class-of-service |
| Configure a best-effort scheduler.                                            | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Schedulers.</li> <li>2. In the Scheduler name box, type the name of the best-effort scheduler—for example, <b>be-scheduler</b>.</li> </ol>                                                                                                                                                                                                                        | Enter<br><br>edit schedulers be-scheduler                                       |
| Configure a best-effort scheduler priority and buffer size.                   | <ol style="list-style-type: none"> <li>1. In the Priority box, type <b>low</b>.</li> <li>2. Click <b>Configure</b> next to Buffer size.</li> <li>3. From the Buffer size choice list, select the basis for the buffer allocation method—for example, <b>percent</b>.</li> <li>4. In the Percent box, type the percentage of the buffer to be used by the best-effort scheduler—for example, <b>40</b>.</li> <li>5. Click <b>OK</b>.</li> </ol> | Enter<br><br><b>set priority low</b><br><br><b>set buffer-size percent 40</b>   |

**Table 162: Configuring Schedulers (continued)**

| <b>Task</b>                                                           | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>CLI Configuration Editor</b>                                  |
|-----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------|
| Configure a best-effort scheduler transmit rate.                      | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to Transmit rate.</li> <li>2. From the Transmit rate choice list, select the basis for the transmit rate method—for example, <b>percent</b>.</li> <li>3. In the Percent box, type the percentage of the bandwidth to be used by the best-effort scheduler—for example, <b>10</b>.</li> <li>4. Click <b>OK</b> twice.</li> </ol>                                                    | Enter<br><br>set transmit-rate percent 10                        |
| Configure an expedited forwarding scheduler.                          | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Schedulers.</li> <li>2. In the Scheduler name box, type the name of the expedited forwarding scheduler—for example, <b>ef-scheduler</b>.</li> </ol>                                                                                                                                                                                                                         | Enter<br><br>edit schedulers ef-scheduler                        |
| Configure an expedited forwarding scheduler priority and buffer size. | <ol style="list-style-type: none"> <li>1. In the Priority box, type <b>high</b>.</li> <li>2. Click <b>Configure</b> next to Buffer size.</li> <li>3. From the Buffer size choice list, select the basis for the buffer allocation method—for example, <b>percent</b>.</li> <li>4. In the Percent box, type the percentage of the buffer to be used by the expedited forwarding scheduler—for example, <b>10</b>.</li> <li>5. Click <b>OK</b>.</li> </ol> | Enter<br><br>set priority high<br><br>set buffer-size percent 10 |
| Configure an expedited forwarding scheduler transmit rate.            | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to Transmit rate.</li> <li>2. From the Transmit rate choice list, select the basis for the transmit rate method—for example, <b>percent</b>.</li> <li>3. In the Percent box, type the percentage of the bandwidth to be used by the expedited forwarding scheduler—for example, <b>10</b>.</li> <li>4. Click <b>OK</b> twice.</li> </ol>                                           | Enter<br><br>set transmit-rate percent 10                        |
| Configure an assured forwarding scheduler.                            | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Schedulers.</li> <li>2. In the Scheduler name box, type the name of the assured forwarding scheduler—for example, <b>af-scheduler</b>.</li> </ol>                                                                                                                                                                                                                           | Enter<br><br>edit schedulers af-scheduler                        |

**Table 162: Configuring Schedulers (continued)**

| <b>Task</b>                                                                                                                                                      | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>CLI Configuration Editor</b>                                                                                                                                                                   |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure an assured forwarding scheduler priority and buffer size.                                                                                              | <ol style="list-style-type: none"> <li>1. In the Priority box, type <b>high</b>.</li> <li>2. Click <b>Configure</b> next to Buffer size.</li> <li>3. From the Buffer size choice list, select the basis for the buffer allocation method—for example, <b>percent</b>.</li> <li>4. In the Percent box, type the percentage of the buffer to be used by the assured forwarding scheduler—for example, <b>45</b>.</li> <li>5. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                 | <p>Enter</p> <p><b>set priority high</b></p> <p><b>set buffer-size percent 45</b></p>                                                                                                             |
| Configure an assured forwarding scheduler transmit rate.                                                                                                         | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to Transmit rate.</li> <li>2. From the Transmit rate choice list, select the basis for the transmit rate method—for example, <b>percent</b>.</li> <li>3. In the Percent box, type the percentage of the bandwidth to be used by the assured forwarding scheduler—for example, <b>45</b>.</li> <li>4. Click <b>OK</b>.</li> </ol>                                                                                                                                                                                                                                                                                 | <p>Enter</p> <p><b>set transmit-rate percent 45</b></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"> <li>1. Click <b>Add new entry</b> next to Drop profile map.</li> <li>2. From the Loss priority box, select <b>Low</b>.</li> <li>3. From the Protocol box, select <b>Any</b>.</li> <li>4. In the Drop profile box, type the name of the drop profile—for example, <b>af-normal</b>.</li> <li>5. Click <b>OK</b>.</li> <li>6. Click <b>Add new entry</b> next to Drop profile map.</li> <li>7. From the Loss priority box, select <b>High</b>.</li> <li>8. From the Protocol box, select <b>Any</b>.</li> <li>9. In the Drop profile box, type the name of the drop profile—for example, <b>af-with-PLP</b>.</li> <li>10. Click <b>OK</b>.</li> </ol> | <p>Enter</p> <p><b>set drop-profile-map loss-priority low protocol any drop-profile af-normal</b></p> <p><b>set drop-profile-map loss-priority high protocol any drop-profile af-with-PLP</b></p> |

**Table 162: Configuring Schedulers (continued)**

| <b>Task</b>                                                     | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                 | <b>CLI Configuration Editor</b>                                                     |
|-----------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------|
| Configure a network control scheduler.                          | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Schedulers.</li> <li>2. In the Scheduler name box, type the name of the network control scheduler—for example, <b>nc-scheduler</b>.</li> </ol>                                                                                                                                                                                                                       | <p>Enter</p> <p><b>edit schedulers nc-scheduler</b></p>                             |
| Configure a network control scheduler priority and buffer size. | <ol style="list-style-type: none"> <li>1. In the Priority box, type <b>low</b>.</li> <li>2. Click <b>Configure</b> next to Buffer size.</li> <li>3. From the Buffer size choice list, select the basis for the buffer allocation method—for example, <b>percent</b>.</li> <li>4. In the Percent box, type the percentage of the buffer to be used by the network control scheduler—for example, <b>5</b>.</li> <li>5. Click <b>OK</b>.</li> </ol> | <p>Enter</p> <p><b>set priority low</b></p> <p><b>set buffer-size percent 5</b></p> |
| Configure a network control scheduler transmit rate.            | <ol style="list-style-type: none"> <li>1. Click <b>Configure</b> next to Transmit rate.</li> <li>2. From the Transmit rate choice list, select the basis for the transmit rate method—for example, <b>percent</b>.</li> <li>3. In the Percent box, type the percentage of the bandwidth to be used by the network control scheduler—for example, <b>5</b>.</li> <li>4. Click <b>OK</b> twice.</li> </ol>                                          | <p>Enter</p> <p><b>set transmit-rate percent 5</b></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 480 to the schedulers configured in “Configuring Schedulers (Optional)” on page 496, as shown in Table 163.

**Table 163: Sample diffserv-cos-map Scheduler Mapping**

| <b>mf-classifier Forwarding Class</b> | <b>For CoS Traffic Type</b>  | <b>diffserv-cos-map Scheduler</b> |
|---------------------------------------|------------------------------|-----------------------------------|
| 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 164.
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 503.
  - To use adaptive shapers to limit bandwidth for Frame Relay, see “Configuring and Applying Adaptive Shaping (Optional)” on page 507.
  - To check the configuration, see “Verifying a DiffServ Configuration” on page 508.

**Table 164: Configuring Scheduler Maps**

| <b>Task</b>                                                                   | <b>J-Web Configuration Editor</b>                                                                                                                                                                                 | <b>CLI Configuration Editor</b>                                                 |
|-------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Navigate to the <b>Class of service</b> level in the configuration hierarchy. | In the configuration editor hierarchy, select <b>Class of service</b> .                                                                                                                                           | From the top of the configuration hierarchy, enter<br><br>edit class-of-service |
| Configure a scheduler map for DiffServ CoS.                                   | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Scheduler maps.</li> <li>2. In the Map name box, type the name of the scheduler map—for example, <b>diffserv-cos-map</b>.</li> </ol> | Enter<br><br>edit scheduler-maps diffserv-cos-map                               |

**Table 164: Configuring Scheduler Maps (continued)**

| <b>Task</b>                                             | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                 | <b>CLI Configuration Editor</b>                                                 |
|---------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| Configure a best-effort forwarding class and scheduler. | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Forwarding class.</li> <li>2. In the Class name box, type the name of the previously configured best-effort forwarding class—<b>be-class</b>.</li> <li>3. In the Scheduler box, type the name of the previously configured best-effort scheduler—<b>be-scheduler</b>.</li> <li>4. Click <b>OK</b>.</li> </ol>        | <p>Enter</p> <p><b>set forwarding-class be-class scheduler be-scheduler</b></p> |
| Configure an expedited forwarding class and scheduler.  | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Forwarding class.</li> <li>2. In the Class name box, type the name of the previously configured expedited forwarding class—<b>ef-class</b>.</li> <li>3. In the Scheduler box, type the name of the previously configured expedited forwarding scheduler—<b>ef-scheduler</b>.</li> <li>4. Click <b>OK</b>.</li> </ol> | <p>Enter</p> <p><b>set forwarding-class ef-class scheduler ef-scheduler</b></p> |
| Configure an assured forwarding class and scheduler.    | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Forwarding class.</li> <li>2. In the Class name box, type the name of the previously configured assured forwarding class—<b>af-class</b>.</li> <li>3. In the Scheduler box, type the name of the previously configured assured forwarding scheduler—<b>af-scheduler</b>.</li> <li>4. Click <b>OK</b>.</li> </ol>     | <p>Enter</p> <p><b>set forwarding-class af-class scheduler af-scheduler</b></p> |



**Table 164: Configuring Scheduler Maps (continued)**

| <b>Task</b>                                      | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                             | <b>CLI Configuration Editor</b>                                                |
|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|
| Configure a network control class and scheduler. | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Forwarding class.</li> <li>2. In the Class name box, type the name of the previously configured network control class—<b>nc-class</b>.</li> <li>3. In the Scheduler box, type the name of the previously configured network control scheduler—<b>nc-scheduler</b>.</li> <li>4. Click <b>OK</b> twice.</li> </ol> | <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"> <li>1. Click <b>Add new entry</b> next to Interfaces.</li> <li>2. In the Interface name box, type the name of the interface—for example, <b>fe-0/0/0</b>.</li> <li>3. In the Scheduler map box, type the name of the previously configured scheduler map—<b>diffserv-cos-map</b>.</li> <li>4. Click <b>OK</b>.</li> </ol>                                  | <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 165.
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 496.

- To use adaptive shapers to limit bandwidth for Frame Relay, see “Configuring and Applying Adaptive Shaping (Optional)” on page 507.
- To check the configuration, see “Verifying a DiffServ Configuration” on page 508.

**Table 165: Configuring and Applying Virtual Channels**

| Task                                                                                                                                                                 | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | CLI Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                         |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Navigate to the <b>Class of service</b> level in the configuration hierarchy.                                                                                        | In the configuration editor hierarchy, select <b>Class of service</b> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | From the top of the configuration hierarchy, enter<br><br>edit class-of-service                                                                                                                                                                                                                                                                                                                                                  |
| Define the virtual channels <b>branch1-vc</b> , <b>branch2-vc</b> , <b>branch3-vc</b> , and the default virtual channel. You must specify a default virtual channel. | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Virtual channels.</li> <li>2. In the Channel name box, type the name of the virtual channel—for example, <b>branch1-vc</b>.</li> <li>3. Click <b>OK</b>.</li> <li>4. Create additional virtual channels for <b>branch2-vc</b>, <b>branch3-vc</b>, and <b>default-vc</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                      | Enter<br><br>set virtual-channels branch1-vc<br><br>Repeat this statement for <b>branch2-vc</b> , <b>branch3-vc</b> , and <b>default-vc</b> .                                                                                                                                                                                                                                                                                    |
| Define the virtual channel group <b>wan-vc-group</b> to include the four virtual channels, and assign each virtual channel the scheduler map <b>bestscheduler</b> .  | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Virtual channel groups.</li> <li>2. In the Group name box, type the name of the virtual channel group—<b>wan-vc-group</b>.</li> <li>3. Click <b>Add new entry</b> next to Channel.</li> <li>4. In the Channel name box, enter the name of the previously configured virtual channels—<b>branch1-vc</b>.</li> <li>5. In the Scheduler map box, enter the name of the previously configured scheduler map—<b>bestscheduler</b>.</li> <li>6. Click <b>OK</b>.</li> <li>7. Add the virtual channels <b>branch2-vc</b>, <b>branch3-vc</b>, and <b>default-vc</b>. Select the <b>Default</b> box when adding the virtual channel <b>default-vc</b>.</li> </ol> | <ol style="list-style-type: none"> <li>1. Enter<br/><br/>           set virtual-channel-groups<br/>           wan-vc-group branch1-vc<br/>           scheduler-map bestscheduler         </li> <li>2. Repeat this statement for <b>branch2-vc</b>, <b>branch3-vc</b>, and <b>default-vc</b>.</li> <li>3. Enter<br/><br/>           set virtual-channel-groups<br/>           wan-vc-group default-vc default         </li> </ol> |

**Table 165: Configuring and Applying Virtual Channels (continued)**

| <b>Task</b>                                                                                   | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>CLI Configuration Editor</b>                                                                                                                                                                                                      |
|-----------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Specify a shaping rate of 1.5 Mbps for each virtual channel within the virtual channel group. | <ol style="list-style-type: none"> <li>1. Click <b>branch1-vc</b> in the list of virtual channels.</li> <li>2. Select the <b>Shaping rate</b> box.</li> <li>3. Click <b>Configure</b>.</li> <li>4. Select <b>Absolute rate</b> from the Rate choice box..</li> <li>5. In the Absolute rate box, enter the shaping rate—<b>1.5m</b>.</li> <li>6. Add the shaping rate for the <b>branch2-vc</b> and <b>branch3-vc</b> virtual channels.</li> <li>7. Click <b>OK</b>.</li> </ol>                        | <ol style="list-style-type: none"> <li>1. Enter<br/><br/> <pre>set virtual-channel-groups wan-vc-group branch1-vc shaping-rate 1.5m</pre> </li> <li>2. Repeat this statement for <b>branch2-vc</b> and <b>branch3-vc</b>.</li> </ol> |
| Apply the virtual channel group to the logical interface t3-1/0/0.0.                          | <ol style="list-style-type: none"> <li>1. Click <b>Add new entry</b> next to Interfaces.</li> <li>2. In the Interface name box, type the name of the interface—<b>t3-1/0/0</b>.</li> <li>3. Click <b>Add new entry</b> next to Unit.</li> <li>4. In the Unit number box, type the logical interface unit number—<b>0</b>.</li> <li>5. In the Virtual channel group box, type the name of the previously configured virtual channel group—<b>wan-vc-group</b>.</li> <li>6. Click <b>OK</b>.</li> </ol> | <pre>Enter set interfaces t3-1/0/0 unit 0 virtual-channel-group wan-vc-group</pre>                                                                                                                                                   |

**Table 165: Configuring and Applying Virtual Channels (continued)**

| <b>Task</b>                                                                                                            | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create the firewall filter <b>choose-vc</b> to select the traffic that is transmitted on a particular virtual channel. | <ol style="list-style-type: none"> <li>1. Navigate to the top of the configuration hierarchy and select <b>Firewall</b>.</li> <li>2. Click <b>Add new entry</b> next to Filter.</li> <li>3. In the Filter name box, enter the name of the firewall filter—<b>choose-vc</b>.</li> <li>4. Click <b>Add new entry</b> next to Term.</li> <li>5. In the Rule name box, enter the name of the firewall term—<b>branch1</b>.</li> <li>6. Click <b>Configure</b> next to From.</li> <li>7. Click <b>Add new entry</b> next to Destination address.</li> <li>8. In the Address box, enter the IP address of the destination host—<b>192.168.10.0/24</b>.</li> <li>9. Click <b>OK</b> twice.</li> <li>10. On the firewall term page, click <b>Configure</b> next to Then.</li> <li>11. Select <b>Accept</b> from the Designation box.</li> <li>12. In the Virtual channel box, enter the name of the previously configured virtual channel—<b>branch1-vc</b>.</li> <li>13. Click <b>OK</b>.</li> <li>14. Repeat these steps for the virtual channels <b>branch2-vc</b> and <b>branch3-vc</b>.</li> </ol> | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><code>edit firewall</code></li> <li>2. Enter<br/><code>set family inet filter choose-vc term branch1 from destination 192.168.10.0/24</code></li> <li>3. Enter<br/><code>set family inet filter choose-vc term branch1 then accept</code></li> <li>4. Enter<br/><code>set family inet filter choose-vc term branch1 then virtual-channel branch1-vc</code></li> <li>5. Repeat these steps for virtual channels <b>branch2-vc</b> and <b>branch3-vc</b>.</li> </ol> |
| Apply the firewall filter <b>choose-vc</b> to output traffic on the <b>t3-1/0/0.0</b> interface.                       | <ol style="list-style-type: none"> <li>1. Navigate to the top of the configuration hierarchy and select <b>Interfaces</b>.</li> <li>2. Click <b>t3-1/0/0</b> in the list of configured interfaces.</li> <li>3. Click <b>0</b> in the list of configured logical units for the interface.</li> <li>4. Click <b>Edit</b> next to Inet.</li> <li>5. Click <b>Configure</b> next to Filter.</li> <li>6. In the Output box, enter the name of the previously configured firewall filter—<b>choose-vc</b>.</li> <li>7. 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 interfaces</code></li> <li>2. Enter<br/><code>set t3-1/0/0 unit 0 family inet filter output choose-vc</code></li> </ol>                                                                                                                                                                                                                                                                                                                                 |

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.

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 166.
- 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 496.
  - To apply rules to logical interfaces, see “Configuring and Applying Virtual Channels (Optional)” on page 503.
  - To check the configuration, see “Verifying a DiffServ Configuration” on page 508.

Table 166: Configuring and Applying an Adaptive Shaper

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

**Table 166: Configuring and Applying an Adaptive Shaper (continued)**

| <b>Task</b>                                                            | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | <b>CLI Configuration Editor</b>                                                      |
|------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| Define the adaptive shaper name and maximum transmit rate.             | <ol style="list-style-type: none"> <li>Next to Adaptive Shapers, click <b>Add new entry</b>.</li> <li>In the Adaptive shaper name box, type <b>fr-shaper</b>.</li> <li>Next to Trigger, click <b>Add new entry</b>.</li> <li>Next to Becn, select the check box.</li> <li>Next to Shaping rate, select the check box and click <b>Configure</b>.</li> <li>From the Rate choice list, select <b>Absolute rate</b>.</li> <li>In the Absolute rate box, type <b>64k</b>.</li> <li>Click <b>OK</b> three times.</li> </ol> | <p>Enter</p> <pre>set adaptive-shapers fr-shaper trigger becn shaping-rate 64k</pre> |
| Apply the adaptive shaper to the logical interface <b>t1-0/0/2.0</b> . | <ol style="list-style-type: none"> <li>Next to Interfaces, click <b>Add new entry</b>.</li> <li>In the Interface name box, type the name of the interface—<b>t1-0/0/2</b>.</li> <li>Next to Unit, click <b>Add new entry</b>.</li> <li>In the Unit number box, type the logical interface unit number—<b>0</b>.</li> <li>In the Adaptive shaper box, type the name of the adaptive shaper—<b>fr-shaper</b>.</li> <li>Click <b>OK</b>.</li> </ol>                                                                       | <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 509
- Verifying an Adaptive Shaper Configuration on page 509
- Verifying a Virtual Channel Configuration on page 510
- Verifying a Virtual Channel Group Configuration on page 510

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&gt; 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"><li>■ Each group address configured, especially the default 224.2.127.254, is listed.</li><li>■ Each port configured, especially the default 9875, is listed.</li></ul> <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&gt; show class-of-service adaptive-shaper  Adaptive shaper: fr-shaper, Index: 35320   Trigger type      Shaping rate       BECN          64000 bps  user@host&gt; 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: &lt;default&gt;, 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"><li>■ The trigger type and shaping rate are consistent with the configured adaptive shaper.</li><li>■ The adaptive shaper applied to the logical interface is displayed under Name.</li></ul>                                                                                                                                                                                                                                                                                  |

## Verifying a Virtual Channel Configuration

|                      |                                                                                                                                                |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b>       | Verify the virtual channel configuration on a logical interface. Verify the class-of-service (CoS) configuration associated with an interface. |
| <b>Action</b>        | From the CLI, enter the show class-of-service virtual-channel commands.                                                                        |
| <b>Sample Output</b> | <pre>user@host&gt; show class-of-service virtual-channel</pre>                                                                                 |
| <b>What It Means</b> | <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

|                      |                                                                                                                                                                                    |
|----------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b>       | Verify the virtual channel group configuration on a logical interface. Verify the class-of-service (CoS) configuration associated with an interface.                               |
| <b>Action</b>        | From the CLI, enter the show class-of-service virtual-channel-group commands.                                                                                                      |
| <b>Sample Output</b> | <pre>user@host&gt; 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> |
| <b>What It Means</b> | Verify that the name of the configured virtual channel group is displayed in the output.                                                                                           |



## **Part 7**

# **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 |

## A

|                                                      |     |
|------------------------------------------------------|-----|
| ABM, HDLC .....                                      | 87  |
| ABRs <i>See</i> area border routers                  |     |
| access concentrator                                  |     |
| as a PPPoE server .....                              | 144 |
| naming for PPPoE .....                               | 151 |
| action modifiers .....                               | 412 |
| actions                                              |     |
| default, routing policy .....                        | 404 |
| final, routing policy .....                          | 404 |
| NAT .....                                            | 407 |
| route list match types .....                         | 426 |
| routing policy .....                                 | 402 |
| routing policy, summary of .....                     | 403 |
| stateful firewall filters .....                      | 407 |
| stateless firewall filters .....                     | 412 |
| active routes, versus passive routes .....           | 225 |
| adaptive shaping, applying CoS rules to logical      |     |
| interfaces .....                                     | 507 |
| Add button .....                                     | 8   |
| Add new entry link .....                             | 10  |
| address match conditions .....                       | 411 |
| address translation <i>See</i> NAT                   |     |
| addresses                                            |     |
| BGP external peer address (configuration             |     |
| editor) .....                                        | 279 |
| BGP internal peer address (configuration             |     |
| editor) .....                                        | 281 |
| BGP local address (Quick Configuration) .....        | 276 |
| BGP peer address (Quick Configuration) .....         | 276 |
| multicast ranges .....                               | 379 |
| physical, in data link layer .....                   | 49  |
| translating <i>See</i> NAT                           |     |
| administrative groups, for MPLS path selection ..... | 303 |
| administrative scoping .....                         | 381 |
| 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 .....                                    | 201 |
| alternate mark inversion <i>See</i> AMI encoding            |     |
| AMI (alternate mark inversion) encoding                     |     |
| E1 .....                                                    | 106 |
| overview .....                                              | 57  |
| T1 .....                                                    | 116 |
| Annex A PIMs                                                |     |
| ATM-over-ADSL interfaces .....                              | 126 |
| operating modes .....                                       | 129 |
| <i>See also</i> ATM-over-ADSL interfaces                    |     |
| standards supported .....                                   | 71  |
| Annex B PIMs                                                |     |
| ATM-over-ADSL interfaces .....                              | 126 |
| operating modes .....                                       | 129 |
| <i>See also</i> ATM-over-ADSL interfaces                    |     |
| standards supported .....                                   | 71  |
| ANSI DMT operating mode .....                               | 129 |
| anycast IPv6 addresses .....                                | 93  |
| Apply button .....                                          | 8   |
| area border routers                                         |     |
| adding interfaces .....                                     | 261 |
| area ID (configuration editor) .....                        | 261 |
| backbone area <i>See</i> backbone area                      |     |
| backbone area interface .....                               | 261 |
| description .....                                           | 209 |
| areas <i>See</i> area border routers; backbone area; NSSAs; |     |
| stub areas                                                  |     |
| ARM, HDLC .....                                             | 87  |
| AS path                                                     |     |
| description .....                                           | 217 |
| forcing by MED .....                                        | 218 |
| prepending .....                                            | 431 |
| role in route selection .....                               | 215 |
| ASs (autonomous systems)                                    |     |
| area border routers .....                                   | 209 |
| AS number (configuration editor) .....                      | 279 |
| AS number (Quick Configuration) .....                       | 276 |
| AS number, in VPNs .....                                    | 330 |
| breaking into confederations .....                          | 221 |
| description .....                                           | 198 |
| group AS number (configuration editor) .....                | 279 |
| individual AS number (configuration editor) .....           | 279 |

|                                                   |     |
|---------------------------------------------------|-----|
| LSPs through .....                                | 296 |
| sample BGP confederation .....                    | 285 |
| stub areas <i>See</i> stub areas                  |     |
| sub-AS number .....                               | 285 |
| assured forwarding .....                          | 495 |
| asymmetrical digital subscriber line (ADSL)       |     |
| <i>See</i> ATM-over-ADSL interfaces               |     |
| Asynchronous Balance Mode, HDLC .....             | 87  |
| asynchronous networks                             |     |
| data stream clocking .....                        | 78  |
| explicit clocking signal transmission .....       | 78  |
| overview .....                                    | 78  |
| Asynchronous Response Mode, HDLC .....            | 87  |
| Asynchronous Transfer Mode (ATM) interface        |     |
| <i>See</i> ATM-over-ADSL interfaces               |     |
| at-0/0/0 <i>See</i> ATM-over-ADSL interfaces      |     |
| ATM interface <i>See</i> ATM-over-ADSL interfaces |     |
| ATM NLPID encapsulation .....                     | 130 |
| ATM PPP over AAL5 LLC encapsulation .....         | 130 |
| ATM PVC encapsulation .....                       | 129 |
| ATM SNAP encapsulation .....                      | 130 |
| ATM VC multiplex encapsulation .....              | 130 |
| ATM-over-ADSL interfaces .....                    | 126 |
| adding .....                                      | 126 |
| ADSL overview .....                               | 70  |
| ADSL systems .....                                | 71  |
| ADSL topology .....                               | 72  |
| ADSL2 .....                                       | 72  |
| ADSL2+ .....                                      | 72  |
| ATM interface type .....                          | 72  |
| CHAP for PPPoA .....                              | 131 |
| CHAP for PPPoE .....                              | 152 |
| encapsulation types, logical .....                | 130 |
| encapsulation types, physical .....               | 129 |
| logical properties .....                          | 130 |
| operating modes .....                             | 129 |
| physical properties .....                         | 127 |
| PPPoE configuration .....                         | 150 |
| PPPoE encapsulation .....                         | 149 |
| PPPoE session on .....                            | 145 |
| statistics .....                                  | 141 |
| VCI .....                                         | 131 |
| verifying .....                                   | 137 |
| verifying a PPPoA configuration .....             | 141 |
| verifying a PPPoE configuration .....             | 154 |
| <i>See also</i> PPPoE; PPPoE over ATM-over-ADSL   |     |
| authentication                                    |     |
| CHAP, for PPPoE interfaces .....                  | 147 |
| OSPF, MD5 .....                                   | 266 |
| OSPF, plain-text passwords .....                  | 266 |
| RIPv2, MD5 .....                                  | 246 |
| RIPv2, plain-text passwords .....                 | 245 |
| auto operating mode .....                         | 129 |
| Auto-RP .....                                     | 382 |

## B

|                                                       |          |
|-------------------------------------------------------|----------|
| B-channels                                            |          |
| description .....                                     | 73       |
| naming convention .....                               | 161      |
| verifying .....                                       | 184      |
| B8ZS encoding .....                                   | 57       |
| BA classifiers <i>See</i> classifiers                 |          |
| backbone area                                         |          |
| area ID (configuration editor) .....                  | 258      |
| area ID (Quick Configuration) .....                   | 254      |
| area type (Quick Configuration) .....                 | 255      |
| configuring .....                                     | 256      |
| description .....                                     | 210      |
| interface .....                                       | 261      |
| backoff algorithm, collision detection .....          | 52       |
| backup connection, ISDN .....                         | 159      |
| backward-explicit congestion notification (BECN)      |          |
| bits .....                                            | 81       |
| bandwidth, for RSVP-signaled LSPs .....               | 316      |
| bandwidth-on-demand                                   |          |
| dialer interface .....                                | 171      |
| ISDN interface .....                                  | 173      |
| overview .....                                        | 171      |
| bc-0/0/0 .....                                        | 161      |
| BECN (backward-explicit congestion notification)      |          |
| bits .....                                            | 81       |
| behavior aggregate classifiers <i>See</i> classifiers |          |
| BERTs (bit error rate tests) .....                    | 77       |
| best-effort service .....                             | 413      |
| BGP (Border Gateway Protocol)                         |          |
| AS number (Quick Configuration) .....                 | 276      |
| <i>See also</i> ASs (autonomous systems), AS          |          |
| number                                                |          |
| AS path .....                                         | 217      |
| <i>See also</i> AS path                               |          |
| confederations <i>See</i> BGP confederations          |          |
| enabling (Quick Configuration) .....                  | 276      |
| export policy for CLNS .....                          | 350      |
| external .....                                        | 214      |
| <i>See also</i> EBGp                                  |          |
| external group type (configuration editor) .....      | 279      |
| external neighbor (peer) address (configuration       |          |
| editor) .....                                         | 279      |
| for CLNS VPN NLRI .....                               | 353      |
| full mesh requirement .....                           | 215, 274 |
| injecting OSPF routes into BGP .....                  | 428      |
| internal .....                                        | 214      |
| <i>See also</i> IBGP                                  |          |
| internal group type (configuration editor) .....      | 281      |
| internal neighbor (peer) address (configuration       |          |
| editor) .....                                         | 281      |
| local address (Quick Configuration) .....             | 276      |
| local preference .....                                | 216      |
| MED metric .....                                      | 218      |
| origin value .....                                    | 217      |

- overview .....212, 273
- peer address (Quick Configuration)..... 276
- peer AS number (Quick Configuration)..... 276
- peering sessions *See* BGP peers; BGP sessions
- point-to-point internal peer session (configuration editor)..... 280
- point-to-point peer session (configuration editor)..... 277
- policy to make routes less preferable..... 431
- Quick Configuration ..... 275
- requirements..... 275
- route reflectors *See* BGP route reflectors
- route selection process ..... 215
  - See also* route selection
- route-flap damping ..... 433
- router ID (Quick Configuration) ..... 276
- routing policy (configuration editor) ..... 281
  - See also* routing policies
- sample BGP peer network..... 278
- sample confederation ..... 285
- sample full mesh ..... 280
- sample route reflector ..... 282
- scaling techniques ..... 218
- session establishment ..... 214
- session maintenance ..... 214
- verifying BGP configuration ..... 288
- verifying BGP groups ..... 287
- verifying BGP peers (neighbors) ..... 286
- verifying peer reachability ..... 289
- VPNs ..... 329
- BGP confederations
  - confederation members..... 286
  - confederation number ..... 285
  - creating (configuration editor) ..... 284
  - description ..... 221, 274
  - route-flap damping ..... 433
  - sample network..... 285
  - sub-AS number ..... 285
- BGP groups
  - cluster identifier (configuration editor) ..... 283
  - confederations (configuration editor)..... 284
  - external group type (configuration editor)..... 279
  - external, creating (configuration editor)..... 279
  - group AS number (configuration editor)..... 279
  - internal group type (configuration editor) ..... 281
  - internal, creating (configuration editor) ..... 281
  - internal, creating for a route reflector (configuration editor) ..... 283
  - verifying..... 287
- BGP messages
  - to establish sessions ..... 214
  - update, to maintain sessions..... 214
- BGP page ..... 275
- BGP peers
  - directing traffic by local preference..... 216
  - external (configuration editor) ..... 277
  - internal (configuration editor)..... 280
  - internal, sample full mesh ..... 280
  - internal, sample route reflector ..... 282
  - peer address (Quick Configuration)..... 276
  - peer AS number (Quick Configuration)..... 276
  - point-to-point connections ..... 213
  - routing policy (configuration editor) ..... 281
    - See also* routing policies
  - sample peer network ..... 278
  - sessions between peers ..... 273
  - verifying ..... 286, 288
  - verifying reachability ..... 289
- BGP route reflectors
  - cluster (configuration editor) ..... 283
  - cluster identifier (configuration editor) ..... 283
  - cluster of clusters ..... 220
  - creating (configuration editor) ..... 281
  - description ..... 219, 274
  - group type (configuration editor) ..... 283
  - multiple clusters ..... 219
  - sample IBGP network ..... 282
- BGP sessions
  - configured at both ends ..... 273
  - establishment ..... 214
  - maintenance ..... 214
  - point-to-point external (configuration editor) .... 277
  - point-to-point internal (configuration editor)..... 280
  - sample peering session ..... 213
  - types ..... 274
- bipolar with 8-zero substitution (B8ZS) encoding .....57
- bit error rate tests (BERTs).....77
- bit stuffing.....60
- bit-field logical operators, stateless firewall filters..... 412
- bit-field match conditions ..... 411
- bit-field synonym match conditions ..... 411
- bootstrap router..... 382
- Border Gateway Protocol *See* BGP
- br-0/0/0..... 161
- braces, in configuration statements.....xxiv
- brackets
  - angle, in syntax descriptions.....xxiv
  - square, in configuration statements .....xxiv
- branches ..... 378
  - See also* multicast
- bridges, on LAN segments .....53
- BSR (bootstrap router)..... 382
- buttons
  - 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 *See* 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                                         |         |
| <b>C</b>                                                              |         |
| C-bit parity frame format                                             |         |
| enable or disable on T3 ports .....                                   | 120     |
| overview .....                                                        | 62      |
| cables                                                                |         |
| T1 cable length .....                                                 | 117     |
| T3 cable length .....                                                 | 120     |
| call setup, ISDN .....                                                | 75      |
| Cancel button                                                         |         |
| J-Web configuration editor .....                                      | 11      |
| Quick Configuration .....                                             | 8       |
| canceled a commit .....                                               | 31–32   |
| carrier sense multiple access with collision detection (CSMA/CD)..... | 51      |
| ccc protocol family.....                                              | 89      |
| CE (customer edge) routers .....                                      | 322     |
| description .....                                                     | 305     |
| VPN task overview.....                                                | 324     |
| VPN topology .....                                                    | 322     |
| <i>See also</i> VPNs                                                  |         |
| Challenge Handshake Authentication Protocol                           |         |
| <i>See</i> CHAP                                                       |         |
| channel number, in interface name .....                               | 48      |
| channel service unit (CSU) device.....                                | 84      |
| CHAP (Challenge Handshake Authentication Protocol)                    |         |
| E1 local identity .....                                               | 106     |
| E3 local identity .....                                               | 109     |
| enabling for PPPoA .....                                              | 131     |
| enabling for PPPoE .....                                              | 152     |
| enabling on E1 .....                                                  | 105     |
| enabling on E3.....                                                   | 108     |
| enabling on serial interfaces .....                                   | 122     |
| enabling on T1 .....                                                  | 115     |
| enabling on T3.....                                                   | 119     |
| overview .....                                                        | 83      |
| PPP links .....                                                       | 83      |
| PPPoE .....                                                           | 147     |
| serial interface local identity .....                                 | 122     |
| T1 local identity.....                                                | 116     |
| T3 local identity.....                                                | 120     |
| CHAP secret <i>See</i> CHAP, local identity                           |         |
| checksum                                                              |         |
| E1 frame .....                                                        | 106     |
| E3 frame .....                                                        | 110     |
| overview .....                                                        | 79      |
| T1 frame .....                                                        | 117     |
| T3 frame .....                                                        | 120     |
| circuit <i>See</i> Layer 2 circuits                                   |         |
| Cisco NLPID encapsulation.....                                        | 130     |
| class of service <i>See</i> CoS                                       |         |
| classful addressing.....                                              | 89      |
| classifiers                                                           |         |
| applying BA classifiers.....                                          | 490–491 |
| default BA classifiers.....                                           | 419     |
| description .....                                                     | 416     |
| sample BA classification.....                                         | 420     |
| sample BA classifier assignments.....                                 | 491     |
| sample, for firewall filter .....                                     | 481     |
| clear system commit command .....                                     | 32      |
| CLI configuration editor                                              |         |
| activating a configuration .....                                      | 31      |
| BGP .....                                                             | 277     |
| CLNS .....                                                            | 347     |
| command summary.....                                                  | 5       |
| committing files .....                                                | 30      |
| confirming a configuration.....                                       | 31      |
| exiting.....                                                          | 22      |
| IPSec tunnels .....                                                   | 360     |
| ISDN connections.....                                                 | 162     |
| managing files .....                                                  | 34      |
| modifying a configuration.....                                        | 25      |
| MPLS traffic engineering .....                                        | 311     |
| network interfaces.....                                               | 124     |
| OSPF .....                                                            | 255     |
| PPPoE .....                                                           | 147     |
| PPPoE over ATM-over-ADSL.....                                         | 147     |
| RIP.....                                                              | 239     |
| saving files .....                                                    | 37      |
| starting .....                                                        | 22      |
| static routes .....                                                   | 228     |
| using show commands with .....                                        | 34      |
| verifying a configuration .....                                       | 30      |
| VPNs .....                                                            | 324     |
| clickable configuration.....                                          | 8       |
| committing.....                                                       | 12      |
| discarding changes .....                                              | 11      |
| viewing and editing.....                                              | 8       |
| <i>See also</i> J-Web configuration editor                            |         |
| CLNS (Connectionless Network Service) VPNs                            |         |
| BGP export policy.....                                                | 350     |
| BGP, to carry CLNS VPN NLRI.....                                      | 353     |
| displaying configurations.....                                        | 354     |
| ES-IS.....                                                            | 349     |
| IS-IS .....                                                           | 350     |
| linking hosts .....                                                   | 345     |
| overview .....                                                        | 346     |
| requirements.....                                                     | 347     |
| static routes (without IS-IS).....                                    | 352     |
| verifying configuration .....                                         | 354     |
| VPN routing instance .....                                            | 348     |
| clock rate, serial interface                                          |         |
| DTE default reduction .....                                           | 67      |
| values .....                                                          | 123     |

- clocking
  - data stream clocking .....78
  - E1 ..... 105
  - E3 ..... 108
  - explicit clocking signal transmission .....78
  - overview .....77
  - serial interface ..... 123
  - serial interface, inverting the transmit clock .. 67, 123
  - serial interface, modes .....66
  - T1.....115
  - T3.....119
- clusters *See* BGP route reflectors
- collision detection
  - backoff algorithm.....52
  - overview .....52
- coloring, link, for MPLS path selection ..... 303
- combined stations, HDLC .....87
- 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
- Compressed Real-Time Transport Protocol *See* CRTP
- confederations *See* 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
    - See also* 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..... 7
  - 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;<br>configuration history; J-Web configuration<br>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 .....                                                                                                   | 81    |
| with DiffServ assured forwarding .....                                                                                                | 494   |
| connection process                                                                                                                    |       |
| ISDN BRI interfaces .....                                                                                                             | 75    |
| LCP, for PPP .....                                                                                                                    | 82    |
| serial interfaces .....                                                                                                               | 65    |
| Connectionless Network Service <i>See</i> CLNS                                                                                        |       |
| connectivity                                                                                                                          |       |
| bidirectional (BGP) .....                                                                                                             | 212   |
| bidirectional (OSPF) .....                                                                                                            | 207   |
| unidirectional (RIP) .....                                                                                                            | 206   |
| Constrained Shortest Path First <i>See</i> CSPF                                                                                       |       |
| conventions                                                                                                                           |       |
| for interface names .....                                                                                                             | 46    |
| how to use this guide .....                                                                                                           | xxii  |
| notice icons .....                                                                                                                    | xxiii |
| text and syntax .....                                                                                                                 | xxiii |
| copy command .....                                                                                                                    | 27    |
| CoS (class of service)                                                                                                                |       |
| adaptive shaping for rules .....                                                                                                      | 507   |
| assigning forwarding classes to output queues .....                                                                                   | 484   |
| BA classifiers .....                                                                                                                  | 490   |
| configuration tasks .....                                                                                                             | 478   |
| default BA classifiers .....                                                                                                          | 419   |
| default forwarding class queue assignments .....                                                                                      | 417   |
| default scheduler settings .....                                                                                                      | 418   |
| DiffServ benefits .....                                                                                                               | 413   |
| <i>See also</i> DiffServ                                                                                                              |       |
| DSCP rewrites .....                                                                                                                   | 420   |
| DSCPs .....                                                                                                                           | 414   |
| <i>See also</i> DSCPs                                                                                                                 |       |
| firewall filter for a multifield classifier .....                                                                                     | 480   |
| JUNOS components .....                                                                                                                | 416   |
| JUNOS implementation .....                                                                                                            | 415   |
| policer for firewall filter .....                                                                                                     | 479   |
| preparation .....                                                                                                                     | 478   |
| RED drop profiles .....                                                                                                               | 494   |
| rewrite rules .....                                                                                                                   | 485   |
| sample BA classification .....                                                                                                        | 420   |
| scheduler maps .....                                                                                                                  | 500   |
| schedulers .....                                                                                                                      | 496   |
| uses .....                                                                                                                            | 477   |
| verifying adaptive shaper configuration .....                                                                                         | 509   |
| verifying multicast session announcements .....                                                                                       | 509   |

|                                                                                   |      |
|-----------------------------------------------------------------------------------|------|
| verifying virtual channel configuration .....                                     | 510  |
| verifying virtual channel group configuration .....                               | 510  |
| virtual channels for rules .....                                                  | 503  |
| cost, of a network path <i>See</i> path cost metrics                              |      |
| CPE device, Services Router as, with PPPoE .....                                  | 143  |
| <i>See also</i> PPPoE                                                             |      |
| CRC (cyclic redundancy check) .....                                               | 79   |
| CRTP (Compressed Real-Time Transport Protocol)                                    |      |
| E1 interfaces .....                                                               | 133  |
| overview .....                                                                    | 99   |
| T1 interfaces .....                                                               | 133  |
| CSMA/CD (carrier sense multiple access with collision<br>detection) .....         | 51   |
| CSPF (Constrained Shortest Path First)                                            |      |
| constraints .....                                                                 | 303  |
| disabling .....                                                                   | 316  |
| link coloring .....                                                               | 303  |
| rules .....                                                                       | 303  |
| CSPF algorithm <i>See</i> CSPF                                                    |      |
| CSU (channel service unit) device .....                                           | 84   |
| curly braces, in configuration statements .....                                   | xxiv |
| customer edge routers <i>See</i> CE routers                                       |      |
| customer premises equipment (CPE) device, Services<br>Router as, with PPPoE ..... | 143  |
| <i>See also</i> PPPoE                                                             |      |
| customer support .....                                                            | xxvi |
| contacting JTAC .....                                                             | xxvi |
| cyclic redundancy check (CRC) .....                                               | 79   |

## D

|                                                |     |
|------------------------------------------------|-----|
| D-channel                                      |     |
| description .....                              | 73  |
| naming convention .....                        | 161 |
| verifying .....                                | 186 |
| D4 framing .....                               | 58  |
| data communications equipment <i>See</i> DCE   |     |
| data inversion                                 |     |
| E1 .....                                       | 106 |
| T1 .....                                       | 116 |
| data link layer                                |     |
| error notification .....                       | 49  |
| flow control .....                             | 50  |
| frame sequencing .....                         | 49  |
| MAC addresses .....                            | 49  |
| network topology .....                         | 49  |
| physical addressing .....                      | 49  |
| purpose .....                                  | 49  |
| sublayers .....                                | 50  |
| data service unit (DSU) device .....           | 84  |
| data stream clocking .....                     | 78  |
| data terminal equipment <i>See</i> DTE         |     |
| data-link connection identifiers (DLCIs) ..... | 81  |
| Database Information page .....                | 16  |
| dc-0/0/0 .....                                 | 161 |



- DCE (data communications equipment)
  - serial connection process ..... 65
  - serial device ..... 64
- DCE clocking mode ..... 66
- DDR *See* dial-on-demand routing, ISDN
- DE (discard eligibility) bits
  - BEcn bits ..... 81
  - FECN bits ..... 81
- deactivate command ..... 29
- deactivating configuration statements or identifiers ..... 29
- default gateway, static routing ..... 227
- defaults
  - BA classifiers ..... 419
  - CoS forwarding class assignments ..... 418
  - junos-algs-outbound group, stateful firewall
    - filters ..... 406
    - routing policy actions ..... 404
    - setting for static routes ..... 232
- Delete button ..... 8
- delete command ..... 26
- Delete Configuration Below This Point radio button .... 11
- Delete link ..... 10
- deleting
  - current rescue configuration (CLI configuration editor) ..... 33
  - current rescue configuration (J-Web) ..... 22
  - network interfaces ..... 134
- denial-of-service attacks, preventing ..... 453
- dense routing mode, caution for use ..... 380
  - See also* multicast routing modes
- designated router, OSPF
  - controlling election ..... 267
  - description ..... 208
- destination prefix lengths ..... 91
- Deutsche Telekom UR-2 operating mode ..... 129
- diagnosis
  - BERT ..... 77
  - displaying CLNS VPN configurations ..... 354
  - displaying firewall filter configurations ..... 465
  - displaying firewall filter statistics ..... 472
  - displaying static routes in the routing table ..... 233
  - LDP neighbors ..... 316
  - LDP sessions ..... 317
  - LDP-signaled LSP ..... 318
  - PPP magic numbers ..... 84
  - RSVP neighbors ..... 319
  - RSVP sessions ..... 319
  - RSVP-signaled LSP ..... 320
  - traffic forwarding over LDP-signaled LSPs ..... 318
  - verifying adaptive shaper configuration ..... 509
  - verifying B-channels ..... 184
  - verifying BGP configuration ..... 288
  - verifying BGP groups ..... 287
  - verifying BGP peer reachability ..... 289
  - verifying BGP peers (neighbors) ..... 286, 509
  - verifying D-channel ..... 186
  - verifying dialer interfaces ..... 187
  - verifying firewall filter actions ..... 473
  - verifying firewall filter DoS protection ..... 474
  - verifying firewall filter flood protection ..... 474
  - verifying firewall filter handles fragments ..... 475
  - verifying firewall filters with packet logs ..... 471
  - verifying IPSec tunnel operation ..... 371
  - verifying ISDN interfaces ..... 183
  - verifying ISDN status ..... 183
  - verifying MPLS traffic engineering ..... 316
  - verifying multicast IGMP versions ..... 392
  - verifying multicast SAP and SDP configuration .. 392
  - verifying OSPF host reachability ..... 271
  - verifying OSPF neighbors ..... 269
  - verifying OSPF routes ..... 270
  - verifying OSPF-enabled interfaces ..... 268
  - verifying PIM mode and interface configuration .. 393
  - verifying PIM RPF routing table ..... 394
  - verifying PIM RPs ..... 393
  - verifying PPPoA for ATM-over-ADSL
    - configuration ..... 141
  - verifying PPPoE interfaces ..... 155
  - verifying PPPoE over ATM-over-ADSL
    - configuration ..... 154
  - verifying PPPoE sessions ..... 156
  - verifying PPPoE statistics ..... 157
  - verifying PPPoE version information ..... 157
  - verifying RIP host reachability ..... 248
  - verifying RIP message exchange ..... 247
  - verifying RIP-enabled interfaces ..... 247
  - verifying stateful firewall filters ..... 470
  - verifying virtual channel configuration ..... 510
  - verifying VPN connectivity ..... 342
- dial-on-demand connectivity, ISDN
  - dialer (dial-on-demand) filter, configuring ..... 169
  - dialer filter, applying ..... 170
  - overview ..... 169
- dial-on-demand filter *See* dialer filter
- dial-on-demand routing, ISDN
  - dialer (dial-on-demand) filter, applying ..... 175
  - dialer (dial-on-demand) filter, configuring ..... 174
  - overview ..... 174
- dialer filter
  - applying, for dial-on-demand connectivity ..... 170
  - applying, for dial-on-demand routing ..... 175
  - for dial-on-demand connectivity ..... 169
  - for dial-on-demand routing ..... 174
- dialer interface
  - applying dialer (dial-on-demand) filter ..... 170
  - bandwidth-on-demand ..... 171
  - dial-on-demand routing ..... 174
  - dialer profiles ..... 181
  - dialer watch ..... 176
  - verifying ..... 187

|                                                    |         |  |
|----------------------------------------------------|---------|--|
| dialer interface, ISDN                             |         |  |
| adding                                             | 165     |  |
| limitations                                        | 161     |  |
| naming convention                                  | 161     |  |
| restrictions                                       | 161     |  |
| dialer profiles, ISDN                              | 181     |  |
| dialer watch                                       |         |  |
| adding dialer watch interface                      | 176     |  |
| ISDN interface                                     | 179     |  |
| overview                                           | 176     |  |
| Differentiated Services <i>See</i> DiffServ        |         |  |
| DiffServ (Differentiated Services)                 |         |  |
| assigning forwarding classes to output queues      | 484     |  |
| assured forwarding                                 | 494     |  |
| BA classifiers                                     | 490     |  |
| benefits for CoS                                   | 413     |  |
| code points                                        | 414     |  |
| <i>See also</i> DSCPs                              |         |  |
| configuration tasks                                | 478     |  |
| default BA classifiers                             | 419     |  |
| default forwarding class queue assignments         | 417     |  |
| default scheduler settings                         | 418     |  |
| DSCP rewrites                                      | 420     |  |
| firewall filter for a multifield classifier        | 480     |  |
| forwarding service classes                         | 414     |  |
| interoperability                                   | 414     |  |
| JUNOS implementation                               | 415     |  |
| policer for firewall filter                        | 479     |  |
| preparation                                        | 478     |  |
| RED drop profiles                                  | 494     |  |
| rewrite rules                                      | 485     |  |
| sample BA classification                           | 420     |  |
| scheduler maps                                     | 500     |  |
| schedulers                                         | 496     |  |
| uses                                               | 477     |  |
| virtual channels for rules                         | 503     |  |
| Discard All Changes radio button                   | 11      |  |
| Discard button                                     | 11      |  |
| Discard Changes Below This Point radio button      | 11      |  |
| discard eligibility bits <i>See</i> DE bits        |         |  |
| discard interface                                  | 97      |  |
| discard rule                                       |         |  |
| firewall filters                                   | 405     |  |
| stateful firewall filters                          | 406     |  |
| stateless firewall filters                         | 408     |  |
| discarding configuration changes                   | 11      |  |
| discovery packets, PPPoE                           | 85, 146 |  |
| Distance Vector Multicast Routing Protocol         | 381     |  |
| distance-vector routing protocols                  | 203     |  |
| <i>See also</i> RIP                                |         |  |
| dl0                                                | 161     |  |
| DLCIs (data-link connection identifiers)           | 81      |  |
| documentation set                                  |         |  |
| comments on                                        | xxvi    |  |
| domains                                            |         |  |
| broadcast domains                                  | 54      |  |
| collision domains                                  | 53      |  |
| DoS (denial-of-service) attacks, preventing        | 453     |  |
| dotted decimal notation                            | 90      |  |
| downloading, configuration files (J-Web)           | 20      |  |
| downstream interfaces                              | 378     |  |
| <i>See also</i> multicast                          |         |  |
| DS1 ports <i>See</i> T1 ports                      |         |  |
| DS1 signals                                        |         |  |
| E1 and T1                                          | 56      |  |
| <i>See also</i> E1 interfaces; T1 interfaces       |         |  |
| multiplexing into DS2 signal                       | 59      |  |
| DS2 signals                                        |         |  |
| bit stuffing                                       | 60      |  |
| frame format                                       | 60      |  |
| DS3 ports <i>See</i> T3 ports                      |         |  |
| DS3 signals                                        |         |  |
| DS3 C-bit parity frame format                      | 62      |  |
| M13 frame format                                   | 61      |  |
| dsc interface                                      | 97      |  |
| DSCPs (DiffServ code points)                       |         |  |
| corresponding forwarding service classes           | 414     |  |
| default forwarding class queue assignments         | 417     |  |
| description                                        | 414     |  |
| replacing with rewrite rules                       | 486     |  |
| rewrites                                           | 420     |  |
| sample BA classification                           | 420     |  |
| DSL access multiplexer (DSLAM) connection          |         |  |
| <i>See</i> DSLAM connection                        |         |  |
| DSLAM connection                                   |         |  |
| ATM-over-ADSL interface for                        | 126     |  |
| PPPoE over ATM-over-ADSL topology                  | 145     |  |
| DSU (data service unit) device                     | 84      |  |
| DTE (data terminal equipment)                      |         |  |
| default clock rate reduction                       | 67      |  |
| serial connection process                          | 65      |  |
| serial device                                      | 64      |  |
| DTE clocking mode                                  | 66      |  |
| DVMRP (Distance Vector Multicast Routing Protocol) | 381     |  |
| dynamic LSPs                                       | 299     |  |
| dynamic routing                                    | 200     |  |
| <b>E</b>                                           |         |  |
| E1 interfaces                                      | 55      |  |
| AMI encoding                                       | 57      |  |
| data stream                                        | 56      |  |
| encoding                                           | 57      |  |
| framing                                            | 58      |  |
| HDB3 encoding                                      | 57      |  |
| loopback                                           | 59      |  |
| overview                                           | 56      |  |
| signals                                            | 56      |  |
| <i>See also</i> E1 ports                           |         |  |

- E1 ports .....55
    - adding CRTP ..... 133
    - CHAP ..... 105
    - clocking ..... 105
    - configuring ..... 103
    - data inversion ..... 106
    - encapsulation type ..... 105
    - fractional, channel number ..... 48
    - frame checksum ..... 106
    - framing ..... 106
    - logical interfaces ..... 105
    - MTU ..... 105
    - overview ..... 56
    - time slots ..... 106
    - See also* E1 interfaces
  - E3 interfaces .....59
    - bit stuffing ..... 60
    - data stream ..... 59
    - DS3 framing ..... 60
    - multiplexing on ..... 60
    - overview ..... 59
    - See also* E3 ports
  - E3 ports .....59
    - CHAP ..... 108
    - clocking ..... 108
    - configuring ..... 106
    - encapsulation type ..... 108
    - frame checksum ..... 110
    - logical interfaces ..... 108
    - MTU ..... 108
    - overview ..... 59
    - See also* E3 interfaces
  - EBGP (external BGP)
    - description ..... 214
    - route-flap damping ..... 433
    - sample network ..... 280
  - edit command .....24
  - Edit Configuration page ..... 9
  - Edit Configuration Text page .....14
  - Edit link .....10
  - EGPs (exterior gateway protocols) ..... 198
  - egress router *See* LSPs; outbound router
  - EIA-232 .....68
  - EIA-422 .....69
  - EIA-449 .....69
  - EIA-530 .....68
  - encapsulation type ..... 105
    - ATM-over-ADSL logical interfaces ..... 130
    - ATM-over-ADSL physical interfaces ..... 129
    - E1 ..... 105
    - E3 ..... 108
    - Frame Relay ..... 80
    - HDLC ..... 86
    - overview ..... 80
    - PPP ..... 82
    - PPPoE ..... 143
    - PPPoE for Ethernet ..... 148
    - PPPoE, over ATM for ADSL ..... 149
    - PPPoE, overview ..... 85
    - serial interfaces ..... 122
    - T1 ..... 115
    - T3 ..... 119
    - See also* packet encapsulation
  - encoding
    - AMI .....57
    - B8ZS .....57
    - HDB3 .....57
  - End System-to-Intermediate System *See* ES-IS
  - EROs (Explicit Route Objects)
    - loose hops ..... 302
    - strict hops ..... 302
  - error notification, in the data link layer .....49
  - ES-IS (End System-to-Intermediate System)
    - for a PE router in a CLNS island ..... 349
    - overview ..... 346
  - ESF (extended superframe) framing .....58
  - Ethernet interfaces .....51
    - access control .....51
    - broadcast domains .....54
    - collision detection .....52
    - collision domains .....53
    - CSMA/CD .....51
    - frame format .....54
    - overview .....51
    - See also* Fast Ethernet ports
  - Ethernet over ATM encapsulation ..... 129
  - Ethernet over LLC encapsulation ..... 130
  - Ethernet ports *See* Fast Ethernet ports
  - ETSI operating mode ..... 129
  - EU-64 addresses .....50
  - exact route list match type ..... 426
  - exit command
    - to leave configuration mode .....23
    - to navigate the configuration hierarchy .....24
  - exit configuration-mode command .....23
  - explicit clocking signal transmission .....78
  - Explicit Route Objects *See* EROs
  - export routing policy, for Layer 2 VPNs ..... 339
  - export statement, for routing policies ..... 404
  - extended superframe (ESF) framing .....58
  - exterior gateway protocols ..... 198
  - external BGP *See* EBGp
- F**
- failover connection, ISDN ..... 159
  - Fast Ethernet ports .....51
    - CHAP for PPPoA ..... 131
    - CHAP for PPPoE ..... 152
    - configuring ..... 111
    - logical interfaces ..... 113

|                                                            |       |                                                         |          |
|------------------------------------------------------------|-------|---------------------------------------------------------|----------|
| MTU .....                                                  | 113   | forwarding states, multicast notation .....             | 379      |
| overview .....                                             | 51    | forwarding table                                        |          |
| PPPoE configuration .....                                  | 150   | controlling OSPF routes in .....                        | 264      |
| PPPoE encapsulation .....                                  | 148   | controlling static routes in .....                      | 224, 231 |
| PPPoE session on .....                                     | 145   | description .....                                       | 199      |
| See also Ethernet interfaces                               |       | MED to determine routes in .....                        | 218      |
| FCS (frame check sequence)                                 |       | FPC (PIM slot on a Services Router) See PIMs            |          |
| checksums .....                                            | 79    | frame check sequence See FCS                            |          |
| CRCs .....                                                 | 79    | Frame Relay encapsulation                               |          |
| overview .....                                             | 79    | congestion control .....                                | 81       |
| two-dimensional parity .....                               | 79    | DLCIs .....                                             | 81       |
| fe-0/0/0                                                   |       | overview .....                                          | 80       |
| disabling PIM on .....                                     | 389   | PVCs .....                                              | 81       |
| management interface .....                                 | 98    | SVCs .....                                              | 81       |
| FEAC C-bit condition indicators .....                      | 64    | virtual circuits .....                                  | 81       |
| FECN (forward-explicit congestion notification) bits ..... | 81    | Frame Relay network, typical .....                      | 80       |
| file management                                            |       | frames                                                  |          |
| configuration files (CLI configuration editor) .....       | 34    | DS2 M-frame format .....                                | 60       |
| configuration files (J-Web) .....                          | 15    | DS3 C-bit parity frame format .....                     | 62       |
| firewall filters                                           |       | DS3 M13 frame format .....                              | 61       |
| applying CoS rules to logical interfaces .....             | 503   | Ethernet frame format .....                             | 54       |
| displaying configurations .....                            | 465   | sequencing, data link layer .....                       | 49       |
| displaying statistics .....                                | 472   | framing                                                 |          |
| multifield classifier filter terms .....                   | 481   | E1 .....                                                | 106      |
| overview .....                                             | 404   | T1 .....                                                | 116      |
| policer for .....                                          | 479   | T3 .....                                                | 120      |
| sample classifier terms .....                              | 481   | FRF.15 and FRF.16 .....                                 | 99       |
| stateful firewall filters .....                            | 405   | from statement, routing policy match conditions .....   | 400      |
| See also stateful firewall filters                         |       | full mesh requirement                                   |          |
| stateless firewall filters .....                           | 405   | description .....                                       | 215      |
| See also stateless firewall filters                        |       | fulfilling with confederations .....                    | 221      |
| term number caution .....                                  | 406   | fulfilling with route reflectors .....                  | 219      |
| verifying configuration .....                              | 465   | sample network .....                                    | 280      |
| verifying flood protection .....                           | 474   | fxp0 interface (not supported) .....                    | 95       |
| verifying fragment handling .....                          | 475   |                                                         |          |
| verifying packet logging .....                             | 471   | <b>G</b>                                                |          |
| Firewall/NAT application page .....                        | 440   | *,G notation, for multicast forwarding states .....     | 379      |
| Firewall/NAT page .....                                    | 439   | gateway, local and remote, for IPSec service sets ..... | 363      |
| field summary .....                                        | 441   | global unicast IPv6 addresses .....                     | 93       |
| flap damping .....                                         | 433   | glossary                                                |          |
| flooding, preventing .....                                 | 453   | CLNS .....                                              | 345      |
| flow control                                               |       | configuration .....                                     | 3        |
| actions in routing policies .....                          | 403   | CoS .....                                               | 397      |
| data link layer .....                                      | 50    | firewall filters .....                                  | 397      |
| font conventions .....                                     | xxiii | interfaces .....                                        | 42       |
| forward-explicit congestion notification (FECN) bits ..... | 81    | ISDN .....                                              | 159      |
| forwarding classes                                         |       | MPLS .....                                              | 293      |
| assigning to output queues .....                           | 485   | multicast .....                                         | 375      |
| default queue assignments .....                            | 417   | ports .....                                             | 42       |
| description .....                                          | 416   | PPPoE .....                                             | 143      |
| mapping to schedulers .....                                | 501   | routing .....                                           | 193      |
| policy to group source and destination prefixes .....      | 430   | routing policies .....                                  | 397      |
| sample BA classification .....                             | 420   | VPNs .....                                              | 293      |
| sample mappings .....                                      | 501   | gr-0/0/0 interface .....                                | 95       |
| forwarding policy options .....                            | 416   | gre interface .....                                     | 95       |

- groups
  - BGP *See* BGP groups
  - default junos-algs-outbound group, for stateful firewall filters ..... 406
  - OSPF areas ..... 258
  - RIP routers ..... 239
- H**
- handling packet fragments ..... 461
- HDB3 encoding ..... 57
- HDLC (High-Level Data Link Control)
  - encapsulation ..... 86
  - HDLC operational modes ..... 87
  - HDLC stations ..... 86
- hierarchy, configuration ..... 24
- high-density bipolar 3 code (HDB3) encoding ..... 57
- High-Level Data Link Control *See* HDLC
- history *See* configuration history
- hold time, to maintain a session ..... 214
- hop count, maximizing ..... 204
  - See also* RIP
- host reachability
  - verifying for a RIP network ..... 248
  - verifying for an OSPF network ..... 271
- hostname
  - for PPPoA CHAP ..... 132
  - for PPPoE CHAP ..... 153
- how to use this guide ..... xxii
- I**
- IBGP (internal BGP)
  - description ..... 214
  - full mesh (configuration editor) ..... 280
  - full mesh requirement ..... 274
  - sample network ..... 280
  - sample route reflector ..... 282
- ICMP (Internet Control Message Protocol), policers... 455
- identifier link ..... 10
- identifiers, configuration
  - adding or modifying ..... 26
  - deactivating ..... 29
  - deleting ..... 26
  - inserting ..... 28
  - renaming ..... 27
- IGMP (Internet Group Management Protocol)
  - IGMPv1 ..... 382
  - IGMPv2 ..... 382
  - IGMPv3 ..... 382
  - setting the version ..... 387
  - verifying the version ..... 392
- IGPs (interior gateway protocols) ..... 331
  - overview ..... 198
  - VPNs ..... 331
  - See also* OSPF; RIP
- IKE (Internet Key Exchange)
  - description ..... 358
  - preshared key (configuration editor) ..... 365
  - preshared key (Quick Configuration) ..... 360
- import routing policy, for Layer 2 VPNs ..... 338
- import statement, for routing policies ..... 404
- inbound router, in an LSP ..... 297
- incoming metric (RIP)
  - description ..... 236
  - modifying ..... 243
- inet protocol family ..... 88
- inet routing table ..... 390
- inet6 protocol family ..... 88
- ingress router *See* inbound router; LSPs
- injecting routes ..... 429
- insert command ..... 28
- inserting configuration identifiers ..... 28
- Integrated Services Digital Network *See* ISDN
- interface naming conventions ..... 46
- interfaces ..... 41
  - ATM-over-ADSL interfaces ..... 70
  - clocking ..... 77
  - data link layer ..... 49
  - E1 interfaces ..... 55
  - E3 interfaces ..... 59
  - Ethernet interfaces ..... 51
  - FCS ..... 79
  - IPv4 addressing ..... 89
  - IPv6 addressing ..... 92
  - ISDN interfaces ..... 73
  - logical properties ..... 87
  - overview ..... 41
  - physical encapsulation ..... 80
    - See also* encapsulation types
  - physical properties ..... 76
  - protocol families ..... 88
  - serial interfaces ..... 64
  - special interfaces ..... 95
  - T1 interfaces ..... 55
  - T3 interfaces ..... 59
  - VLANs ..... 94
  - See also* ATM-over-ADSL interfaces; ISDN
    - interfaces; loopback interfaces; management
    - interfaces; network interfaces; services
    - interfaces; special interfaces; ports
- Interfaces page ..... 102
  - for E1 ..... 104
  - for E3 ..... 107
  - for Fast Ethernet ..... 112
  - for serial interfaces ..... 121
  - for T1 ..... 114
  - for T3 (DS3) ..... 118
- interior gateway protocols *See* IGPs
- Intermediate System-to-Intermediate System *See* IS-IS
- internal BGP *See* IBGP

|                                                         |     |                                                |          |
|---------------------------------------------------------|-----|------------------------------------------------|----------|
| Internet Control Message Protocol policers.....         | 455 | ISDN BRI interfaces.....                       | 73       |
| Internet Group Management Protocol <i>See</i> IGMP      |     | adding an interface .....                      | 162      |
| Internet Key Exchange <i>See</i> IKE                    |     | B-channel interface .....                      | 161      |
| Internet routing, with BGP .....                        | 273 | bandwidth-on-demand .....                      | 171      |
| invalid configuration, replacing                        |     | call setup .....                               | 75       |
| with J-Web .....                                        | 21  | connection initialization.....                 | 75       |
| with the CLI .....                                      | 33  | D-channel interface.....                       | 161      |
| invalid routes, rejecting.....                          | 428 | dial-on-demand connectivity.....               | 169      |
| inverting the transmit clock .....                      | 123 | dial-on-demand routing .....                   | 174      |
| IP Security <i>See</i> IPSec                            |     | dial-on-demand routing, with OSPF .....        | 180      |
| ip-0/0/0 interface.....                                 | 96  | dialer interface .....                         | 161      |
| ip-ip interface .....                                   | 96  | dialer interface, adding.....                  | 165      |
| IPSec (IP Security)                                     |     | dialer profiles .....                          | 181      |
| IKE <i>See</i> IKE                                      |     | dialer watch .....                             | 176      |
| security associations.....                              | 358 | ISDN channels .....                            | 73       |
| tunnels <i>See</i> IPSec tunnels                        |     | naming conventions.....                        | 161      |
| verifying tunnels .....                                 | 371 | NT1 devices.....                               | 74       |
| IPSec security associations.....                        | 358 | overview .....                                 | 73       |
| <i>See also</i> IKE                                     |     | PIMs supported .....                           | 161      |
| IPSec tunnels                                           |     | requirements.....                              | 162      |
| IKE key (configuration editor).....                     | 365 | S/T interfaces .....                           | 74, 161  |
| IKE key (Quick Configuration) .....                     | 360 | secondary (backup) connection.....             | 168      |
| IPSec rule (configuration editor) .....                 | 366 | session establishment .....                    | 75       |
| local endpoint (Quick Configuration) .....              | 360 | switch types supported.....                    | 164      |
| NAT pools (configuration editor) .....                  | 368 | typical network .....                          | 73       |
| outgoing traffic filters.....                           | 358 | U interface .....                              | 75, 161  |
| overview .....                                          | 357 | verifying B-channels.....                      | 184      |
| private addresses (Quick Configuration).....            | 360 | verifying D-channel .....                      | 186      |
| Quick Configuration .....                               | 358 | verifying dialer interfaces .....              | 187      |
| remote endpoint (Quick Configuration) .....             | 360 | verifying ISDN interfaces.....                 | 183      |
| requirements.....                                       | 358 | verifying ISDN status .....                    | 183      |
| services interfaces (configuration editor).....         | 361 | <i>See also</i> ISDN connections               |          |
| services sets (configuration editor) .....              | 362 | ISDN connections.....                          | 159      |
| stateful firewall filter rules(configuration editor) .. | 366 | adding an interface .....                      | 162      |
| verifying.....                                          | 371 | bandwidth-on-demand .....                      | 171      |
| IPSec Tunnels page .....                                | 359 | configuring.....                               | 159      |
| field summary .....                                     | 360 | dial-on-demand connectivity.....               | 169      |
| IPv4 addressing                                         |     | dial-on-demand routing .....                   | 174      |
| classful addressing.....                                | 89  | dial-on-demand routing, with OSPF .....        | 180      |
| dotted decimal notation.....                            | 90  | dialer (dial-on-demand) filter, applying ..... | 170, 175 |
| MAC-48 address format .....                             | 50  | dialer (dial-on-demand) filter, configuring..  | 169, 174 |
| overview .....                                          | 89  | dialer interface .....                         | 161      |
| subnets.....                                            | 90  | dialer interface for bandwidth-on-demand.....  | 171      |
| VLSMs .....                                             | 91  | dialer interface, adding.....                  | 165      |
| IPv6 addressing                                         |     | dialer profiles .....                          | 181      |
| address format.....                                     | 92  | dialer watch .....                             | 176      |
| address scope.....                                      | 93  | dialer watch interface .....                   | 176      |
| address structure .....                                 | 93  | interface naming conventions .....             | 161      |
| address types .....                                     | 93  | ISDN interface for bandwidth-on-demand .....   | 173      |
| overview .....                                          | 92  | ISDN interface for dialer watch .....          | 179      |
| IPv6 support .....                                      | 193 | ISDN interface types.....                      | 161      |
| IS-IS (Intermediate System-to-Intermediate System)      |     | overview .....                                 | 160      |
| for CLNS route exchange.....                            | 350 | requirements.....                              | 162      |
| overview .....                                          | 346 | secondary (backup) connection.....             | 168      |
|                                                         |     | switch types supported.....                    | 164      |

- verifying B-channels ..... 184
  - verifying D-channel ..... 186
  - verifying dialer interfaces ..... 187
  - verifying ISDN interfaces ..... 183
  - verifying ISDN status ..... 183
  - See also* ISDN BRI interfaces
  - ISO protocol family ..... 88
  - ITU Annex B non-UR-2 operating mode ..... 129
  - ITU Annex B UR-2 operating mode ..... 129
  - ITU DMT BIS operating mode ..... 129
  - ITU DMT operating mode ..... 129
- J**
- J-series
- BGP routing ..... 273
  - CLNS VPNs ..... 345
  - configuration tools ..... 3
  - CoS overview ..... 413
  - CoS with DiffServ ..... 477
  - firewall filter overview ..... 404
  - firewall filters ..... 437
  - interfaces overview ..... 41
  - IPSec tunnels ..... 357
  - ISDN connections ..... 159
  - MPLS for VPNs overview ..... 293
  - MPLS traffic engineering ..... 309
  - multicast ..... 385
  - multicast overview ..... 375
  - NAT ..... 437
  - network interfaces ..... 101
  - OSPF routing ..... 251
  - PPPoE ..... 143
  - release notes, URL ..... xxi
  - RIP routing ..... 235
  - routing policies ..... 423
  - routing policy overview ..... 399
  - routing protocols overview ..... 193
  - static routing ..... 223
  - VPNs ..... 321
- J-Web configuration editor
- BGP ..... 277
  - clickable configuration, committing ..... 12
  - clickable configuration, discarding changes ..... 11
  - clickable configuration, editing ..... 8
  - CLNS ..... 347
  - committing a text file, with caution ..... 13
  - configuration text, viewing ..... 12
  - editing a text file, with caution ..... 13
  - IPSec tunnels ..... 360
  - ISDN connections ..... 162
  - managing files ..... 15
  - MPLS traffic engineering ..... 311
  - network interfaces ..... 124
  - OSPF ..... 255
  - PPPoE ..... 147
  - PPPoE over ATM-over-ADSL ..... 147
  - RIP ..... 239
  - static routes ..... 228
  - uploading a file ..... 14
  - VPNs ..... 324
- 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 ..... 416
  - CoS functions ..... 415
  - DiffServ implementation ..... 415
  - ISDN connections ..... 159
  - release notes, URL ..... xxi
- junos-als-outbound group, for stateful firewall filters ..... 406
- K**
- keepalive interval, for LDP-signaled LSPs ..... 313
  - keepalive messages, for session hold time ..... 214
- L**
- Label Distribution Protocol *See* LDP
- label switching ..... 296
- label-switched paths *See* LSPs
- label-switching routers (LSRs) ..... 297
- labels, MPLS ..... 298
- label operations ..... 298
  - PHP ..... 299
- LANs
- bridges on LAN segments ..... 53
  - collision domains ..... 53
  - repeaters on LAN segments ..... 53
  - topology ..... 94
- Layer 2 circuits
- AS number ..... 330
  - basic, description ..... 323
  - encapsulation ..... 326
  - IGPs ..... 331
  - MPLS ..... 327
  - neighbor address ..... 334
  - participating interfaces ..... 325
  - signaling protocols ..... 331
  - task overview ..... 324
  - verifying PE router connections ..... 343
  - verifying PE router interfaces ..... 343
  - virtual circuit ID ..... 334

|                                                 |     |                                                               |          |
|-------------------------------------------------|-----|---------------------------------------------------------------|----------|
| Layer 2 VPNs                                    |     | link-state advertisements <i>See</i> LSAs                     |          |
| AS number                                       | 330 | lo0 interface functions                                       | 98       |
| basic, description                              | 322 | lo0.1 6385 interface                                          | 96       |
| BGP                                             | 329 | load command                                                  | 35       |
| encapsulation                                   | 326 | load merge command                                            | 35       |
| export routing policies                         | 339 | load override command                                         | 35       |
| IGPs                                            | 331 | load patch command                                            | 35       |
| import routing policies                         | 338 | load replace command                                          | 35       |
| MPLS                                            | 327 | loading a configuration file                                  |          |
| overview                                        | 307 | CLI configuration editor                                      | 34       |
| participating interfaces                        | 325 | downloading (J-Web)                                           | 20       |
| routing instance                                | 335 | rollback (J-Web)                                              | 21       |
| signaling protocols                             | 331 | rollback command                                              | 32       |
| task overview                                   | 324 | uploading (J-Web)                                             | 14       |
| verifying PE router connections                 | 343 | without specifying full hierarchy                             | 35       |
| verifying PE router interfaces                  | 343 | local preference                                              |          |
| Layer 3 VPNs                                    |     | description                                                   | 216      |
| AS number                                       | 330 | high value preferred                                          | 217      |
| basic, description                              | 323 | role in route selection                                       | 215      |
| BGP                                             | 329 | local tunnel endpoint, IPSec                                  | 360      |
| IGPs                                            | 331 | locked configuration                                          | 23       |
| overview                                        | 307 | logical interfaces                                            |          |
| participating interfaces                        | 325 | adaptive shaping for                                          | 507      |
| route target                                    | 335 | adding (configuration editor)                                 | 126      |
| routing instance                                | 335 | ATM-over-ADSL                                                 | 130      |
| routing policies                                | 341 | CoS rules for                                                 | 503, 507 |
| signaling protocols                             | 331 | E1                                                            | 105      |
| task overview                                   | 324 | E3                                                            | 108      |
| verifying PE router connections                 | 343 | Fast Ethernet                                                 | 113      |
| LCP (Link Control Protocol), connection process | 82  | inside services interface, IPSec                              | 361      |
| LDP (Label Distribution Protocol)               |     | outside services interface, IPSec                             | 361      |
| and OSPF for VPNs                               | 331 | serial                                                        | 122      |
| LDP-signaled LSPs                               | 311 | T1                                                            | 115      |
| messages                                        | 300 | T3                                                            | 119      |
| operation                                       | 300 | virtual channels for                                          | 503      |
| overview                                        | 310 | logical units                                                 |          |
| requirements                                    | 310 | adding (configuration editor)                                 | 126      |
| verifying LSPs                                  | 318 | E1 interface                                                  | 105      |
| verifying neighbors                             | 316 | E3 interface                                                  | 108      |
| verifying sessions                              | 317 | Fast Ethernet interface                                       | 113      |
| verifying traffic forwarding                    | 318 | number in interface name                                      | 48       |
| LDP neighbors, verifying                        | 316 | pp0 interface                                                 | 150      |
| LDP-signaled LSP <i>See</i> LDP                 |     | PPPoE encapsulation                                           | 148      |
| leaves                                          | 378 | PPPoE over ATM-over-ADSL encapsulation                        | 149      |
| <i>See also</i> multicast                       |     | serial interface                                              | 122      |
| line buildout                                   |     | T1 interface                                                  | 115      |
| T1                                              | 117 | T3 interface                                                  | 119      |
| T3                                              | 120 | long buildout <i>See</i> line buildout                        |          |
| line speed, serial interface                    | 123 | longer route list match type                                  | 426      |
| line timing                                     | 66  | loop clocking mode                                            | 66       |
| link coloring, for MPLS path selection          | 303 | loopback address, for PE routers in VPNs                      | 331      |
| link services                                   | 99  | loopback interfaces                                           |          |
| <i>See also</i> ls-0/0/0                        |     | applying stateless firewall filters to (configuration editor) | 464      |
| link states, verifying                          | 135 | functions                                                     | 98       |
| link-local unicast IPv6 addresses               | 93  |                                                               |          |



- loopback signals, E1 and T1 .....59
- loose hops, RSVP ..... 302
- loss priority, CoS ..... 416
- ls-0/0/0
  - adding CRTP ..... 133
  - interface description.....96
- LSAs (link-state advertisements)
  - description ..... 208
  - three-way handshake ..... 208
- lsi interface.....96
- LSPs (label-switched paths)
  - bandwidth ..... 316
  - description ..... 296
  - disabling CSPF ..... 316
  - dynamic LSPs..... 299
  - for RSVP in a VPN ..... 328
  - keepalive interval for LDP link ..... 313
  - label operations..... 298
  - label switching ..... 296
  - labels ..... 298
  - LDP ..... 300
  - LDP-signaled LSPs ..... 311
  - LSR types ..... 297
  - overview ..... 309
  - PHP ..... 299
  - RSVP ..... 300
  - RSVP-signaled LSPs..... 313
  - static LSPs ..... 299
  - verifying LDP-signaled LSPs ..... 316
  - verifying RSVP-signaled LSPs ..... 319
- LSRs (label-switching routers) ..... 297
- lt-0/0/0 interface .....96
- M**
- M13 frame format .....61
- MAC (media access control) *See* MAC addresses
- MAC addresses
  - EUI-64 addresses .....50
  - MAC-48 address format .....50
  - overview .....50
  - physical addressing .....49
- MAC-48 addresses .....50
- magic numbers, PPP .....84
- management interfaces
  - disabling PIM on ..... 389
  - overview .....98
- managing files *See* file management
- manuals
  - comments on .....xxvi
- mapping, CoS forwarding classes to schedulers..... 501
- match conditions
  - routing policy ..... 400
  - routing policy, summary of..... 400
  - stateful firewall filter and NAT..... 407
  - stateless firewall filters ..... 409
  - stateless firewall filters, summary of..... 410
  - match types ..... 426
  - maximum hop count, RIP..... 204
  - maximum transmission unit *See* MTU
  - MED (multiple exit discriminator)
    - description ..... 218
    - role in route selection ..... 215
  - media access control *See* MAC addresses
  - media types supported .....45
  - merging a configuration file .....35
    - example .....37
  - messages, LDP ..... 300
  - metrics *See* path cost metrics
  - MF classifier..... 480
  - MLFR (Multilink Frame Relay).....99
  - MLFR FRF.15 and FRF.16.....99
  - mlfr-end-to-end protocol family.....89
  - mlfr-uni-nni protocol family .....89
  - MLPPP (Multilink Point-to-Point Protocol) .....99
  - mlppp protocol family .....89
  - MPLS (Multiprotocol Label Switching) ..... 304
    - dynamic LSPs..... 299
    - label operations..... 298
    - label switching ..... 296
    - labels ..... 298
    - Layer 2 VPNs and Layer 2 circuits ..... 327
    - LDP ..... 300
    - LSP for RSVP in a VPN ..... 328
    - LSPs ..... 296
    - LSR types ..... 297
    - overview ..... 293
    - PHP ..... 299
    - RSVP ..... 300
    - static LSPs..... 299
    - traffic engineering *See* MPLS traffic engineering
    - verifying ..... 316
    - See also* VPNs
  - MPLS protocol family .....88
  - MPLS traffic engineering
    - LDP signaling ..... 310
    - LDP-signaled LSPs ..... 311
    - overview ..... 309
    - requirements..... 310
    - RSVP signaling..... 310
    - RSVP-signaled LSPs..... 313
    - signaling protocols overview ..... 300
    - verifying LDP neighbors..... 316
    - verifying LDP sessions ..... 317
    - verifying LDP-signaled LSPs..... 318
    - verifying RSVP neighbors ..... 319
    - verifying RSVP sessions ..... 319
    - verifying RSVP-signaled LSPs ..... 320
    - verifying traffic forwarding over LDP-signaled LSPs..... 318
  - MSDP (Multicast Source Discovery Protocol) ..... 383

|                                                   |          |
|---------------------------------------------------|----------|
| mt-0/0/0 interface.....                           | 96       |
| MTU (maximum transmission unit)                   |          |
| E1 .....                                          | 105      |
| E3 .....                                          | 108      |
| Fast Ethernet.....                                | 113      |
| T1.....                                           | 115      |
| T3.....                                           | 119, 122 |
| mtun interface .....                              | 96       |
| multiarea network, OSPF.....                      | 258      |
| multicast                                         |          |
| administrative scoping .....                      | 381      |
| architecture .....                                | 378      |
| Auto-RP.....                                      | 382      |
| BSR.....                                          | 382      |
| downstream interface .....                        | 378      |
| DVMRP .....                                       | 381      |
| forwarding state notation .....                   | 379      |
| *,G notation.....                                 | 379      |
| IGMP <i>See</i> IGMP                              |          |
| IP address ranges.....                            | 379      |
| MSDP.....                                         | 383      |
| network elements .....                            | 379      |
| overview .....                                    | 375      |
| PGM .....                                         | 383      |
| PIM dense mode <i>See</i> PIM                     |          |
| PIM source-specific multicast (SSM) .....         | 382      |
| PIM sparse mode <i>See</i> PIM                    |          |
| preparation.....                                  | 386      |
| preventing routing loops .....                    | 380      |
| protocols .....                                   | 381      |
| reverse-path forwarding (RPF) .....               | 380      |
| routing modes <i>See</i> multicast routing modes  |          |
| S,G notation.....                                 | 379      |
| SAP and SDP <i>See</i> SAP; SDP                   |          |
| session announcements.....                        | 386      |
| shortest-path tree (SPT) .....                    | 381      |
| static RP.....                                    | 388      |
| <i>See also</i> RP                                |          |
| subnetwork leaves and branches .....              | 378      |
| upstream interface .....                          | 378      |
| verifying IGMP versions .....                     | 392      |
| verifying PIM mode and interface configuration .. | 393      |
| verifying PIM RPF routing table.....              | 394      |
| verifying PIM RPs.....                            | 393      |
| verifying SAP and SDP configuration.....          | 392      |
| multicast IPv6 addresses .....                    | 93       |
| multicast routing modes                           |          |
| dense mode.....                                   | 380      |
| dense mode, caution for use.....                  | 380      |
| sparse mode .....                                 | 380      |
| Multicast Source Discovery Protocol .....         | 383      |
| multifield classifier.....                        | 480      |
| Multilink Frame Relay (MLFR).....                 | 99       |
| Multilink Frame Relay Forum .....                 | 99       |
| Multilink Point-to-Point Protocol (MLPPP) .....   | 99       |

|                                               |     |
|-----------------------------------------------|-----|
| multilink services                            |     |
| CRTP .....                                    | 99  |
| MLFR.....                                     | 99  |
| MLFR FRF.15 and FRF.16.....                   | 99  |
| MLPPP.....                                    | 99  |
| multiple exit discriminator <i>See</i> MED    |     |
| multiple push label operation .....           | 299 |
| Multiprotocol Label Switching <i>See</i> MPLS |     |

## N

|                                                     |          |
|-----------------------------------------------------|----------|
| names, of network interfaces .....                  | 47       |
| NAPT .....                                          | 405      |
| NAT (Network Address Translation)                   |          |
| actions .....                                       | 407      |
| applying to an interface (configuration editor) ... | 447      |
| configuration editor .....                          | 442, 444 |
| description .....                                   | 404      |
| enabling (Quick Configuration) .....                | 441      |
| match conditions .....                              | 407      |
| pools for IPSec tunnels (configuration editor) ...  | 368      |
| preparation.....                                    | 438      |
| Quick Configuration .....                           | 438      |
| sample rules .....                                  | 443      |
| verifying.....                                      | 470      |
| NCPs (Network Control Protocols) .....              | 83       |
| neighbors <i>See</i> BGP peers; OSPF neighbors; RIP |          |
| neighbors                                           |          |
| Network Address Port Translation (NAPT) .....       | 405      |
| Network Address Translation <i>See</i> NAT          |          |
| Network Control Protocols (NCPs) .....              | 83       |
| network interfaces                                  |          |
| adding.....                                         | 124      |
| ATM-over-ADSL configuration .....                   | 126      |
| ATM-over-ADSL interfaces.....                       | 70       |
| clocking .....                                      | 77       |
| deleting.....                                       | 134      |
| DS3 configuration .....                             | 117      |
| E1 configuration .....                              | 103      |
| E1 interfaces .....                                 | 55       |
| E3 configuration .....                              | 106      |
| E3 interfaces .....                                 | 59       |
| enabling PIM on .....                               | 389      |
| enabling RIP on.....                                | 238      |
| Ethernet interfaces .....                           | 51       |
| Fast Ethernet configuration .....                   | 111      |
| FCS .....                                           | 79       |
| IPv4 addressing.....                                | 89       |
| IPv6 addressing.....                                | 92       |
| ISDN interfaces .....                               | 73       |
| logical properties .....                            | 87       |
| media types.....                                    | 45       |
| multicast, upstream and downstream.....             | 378      |
| names .....                                         | 47       |
| naming conventions.....                             | 46       |
| output, understanding.....                          | 48       |

- physical encapsulation ..... 80
    - See also* encapsulation types
  - physical properties ..... 76
  - preparation ..... 101
  - protocol families ..... 88
  - sample name ..... 48
  - serial configuration ..... 120
  - serial interfaces ..... 64
  - supported ..... 44
  - T1 configuration ..... 113
  - T1 interfaces ..... 55
  - T3 configuration ..... 117
  - T3 interfaces ..... 59
  - verifying ATM-over-ADSL properties ..... 137
  - verifying link states ..... 135
  - verifying PIM on ..... 393
  - verifying properties ..... 136
  - verifying RIP message exchange ..... 247
  - verifying RIP on ..... 247
  - VLANs ..... 94
  - VPN configuration ..... 325
  - network layer reachability information (NLRI), BGP, for
    - CLNS ..... 353
  - network service access points *See* NSAPs
  - networks ..... 322
    - description ..... 198
    - designated router *See* designated router, OSPF
    - IPv4 subnets ..... 90
    - path cost metrics *See* path cost metrics
    - PPPoE session on an ATM-over-ADSL loop ..... 146
    - PPPoE session on an Ethernet loop ..... 145
    - sample BGP AS path ..... 217
    - sample BGP confederation ..... 285
    - sample BGP confederations ..... 222
    - sample BGP external and internal links ..... 280
    - sample BGP local preference use ..... 216
    - sample BGP MED use ..... 218
    - sample BGP peer network ..... 278
    - sample BGP peer session ..... 213
    - sample BGP route reflector (one cluster) ... 219, 282
    - sample BGP route reflectors (cluster of clusters) .. 221
    - sample BGP route reflectors (multiple clusters) .. 220
    - sample distance-vector routing ..... 204
    - sample LSP topology ..... 297
    - sample multiarea OSPF routing ..... 210
    - sample OSPF backbone area ..... 211
    - sample OSPF multiarea network ..... 258
    - sample OSPF network with stubs and NSSAs .... 212
    - sample OSPF single-area network ..... 257
    - sample OSPF stub areas and NSSAs ..... 262
    - sample OSPF topology ..... 270
    - sample poison reverse routing ..... 206
    - sample RIP network with incoming metric ..... 242
    - sample RIP network with outgoing metric ..... 244
    - sample RIP topology ..... 239
    - sample route advertisement ..... 201
    - sample route aggregation ..... 202
    - sample routing topology ..... 199
    - sample RSVP topology ..... 302
    - sample split horizon routing ..... 205
    - sample static route, preferred path ..... 230
    - sample stub network for static routes ..... 228
    - sample unidirectional routing ..... 207
    - sample VPN topology ..... 322
    - static routing ..... 200
    - trusted ..... 404
    - untrusted ..... 404
    - See also* VPNs
  - next hop
    - address for static routes ..... 227
    - defining for static routes ..... 229
    - qualified, defining for static routes ..... 231
    - qualified, for static routes ..... 224
    - service set, for IPsec tunnels ..... 362
  - NLRI (network layer reachability information), BGP, for
    - CLNS ..... 353
  - non-UR-2 operating mode ..... 129
  - Normal Response Mode, HDLC ..... 87
  - not-so-stubby areas *See* NSSAs
  - notice icons ..... xxiii
  - NRM, HDLC ..... 87
  - NSAPs (network service access points)
    - overview ..... 346
    - samples ..... 352
  - NSSAs (not-so-stubby areas)
    - area ID (configuration editor) ..... 260
    - area ID (Quick Configuration) ..... 254
    - area type (Quick Configuration) ..... 255
    - creating (configuration editor) ..... 261
    - description ..... 211
    - example ..... 212
    - sample topology ..... 262
  - NT1 devices ..... 74
  - numeric range match conditions ..... 410
- O**
- OK button
    - J-Web configuration editor ..... 11
    - Quick Configuration ..... 8
  - Open Shortest Path First protocol *See* OSPF
  - Open Systems Interconnection (OSI) networks, CLNS
    - VPNs ..... 345
  - operational mode, entering during configuration ..... 34
  - origin, of BGP route ..... 217
  - orlonger route list match type ..... 427
  - OSI (Open Systems Interconnection) networks, CLNS
    - VPNs ..... 345
  - OSPF (Open Shortest Path First)
    - and LDP for VPNs ..... 332
    - and RSVP for VPNs ..... 333

|                                                                          |          |
|--------------------------------------------------------------------------|----------|
| area border routers <i>See</i> area border routers                       |          |
| area type (Quick Configuration).....                                     | 255      |
| areas.....                                                               | 209, 252 |
| <i>See also</i> area border routers; backbone area;<br>NSSAs; stub areas |          |
| authenticating exchanges (OSPFv2 only).....                              | 265      |
| backbone area <i>See</i> backbone area                                   |          |
| controlling designated router election .....                             | 267      |
| controlling route cost .....                                             | 264      |
| designated router <i>See</i> designated router, OSPF                     |          |
| designating OSPF interfaces (configuration<br>editor).....               | 258, 260 |
| designating OSPF interfaces (Quick<br>Configuration) .....               | 255      |
| dial-on-demand routing support, ISDN .....                               | 180      |
| enabling (Quick Configuration) .....                                     | 254      |
| enabling, description .....                                              | 251      |
| ensuring efficient operation.....                                        | 263      |
| injecting OSPF routes into BGP .....                                     | 428      |
| ISDN dial-on-demand routing support.....                                 | 180      |
| LSAs .....                                                               | 208      |
| multiarea network (configuration editor).....                            | 258      |
| NSSAs <i>See</i> NSSAs                                                   |          |
| overview .....                                                           | 207, 251 |
| path cost metrics <i>See</i> path cost metrics                           |          |
| Quick Configuration .....                                                | 253      |
| requirements.....                                                        | 253      |
| route preferences.....                                                   | 264      |
| router ID (configuration editor) .....                                   | 256      |
| router ID (Quick Configuration) .....                                    | 254      |
| sample multiarea network .....                                           | 258      |
| sample network topology .....                                            | 270      |
| sample NSSAs.....                                                        | 262      |
| sample single-area network .....                                         | 257      |
| sample stub areas .....                                                  | 262      |
| single-area network (configuration editor) .....                         | 256      |
| stub areas <i>See</i> stub areas                                         |          |
| supported versions .....                                                 | 208      |
| three-way handshake .....                                                | 208      |
| tuning an OSPF network .....                                             | 263      |
| verifying host reachability.....                                         | 271      |
| verifying neighbors .....                                                | 269      |
| verifying RIP-enabled interfaces .....                                   | 268      |
| verifying routes .....                                                   | 270      |
| OSPF interfaces                                                          |          |
| enabling .....                                                           | 255      |
| enabling (configuration editor).....                                     | 258, 260 |
| enabling, for area border routers .....                                  | 261      |
| verifying.....                                                           | 268      |
| OSPF neighbors, verifying.....                                           | 269      |
| OSPF page .....                                                          | 253      |
| field summary .....                                                      | 254      |
| outbound router, in an LSP.....                                          | 297      |
| outgoing metric (RIP)                                                    |          |
| description .....                                                        | 236      |
| modifying .....                                                          | 245      |
| output queues                                                            |          |
| assigning forwarding classes .....                                       | 485      |
| sample assignments.....                                                  | 484      |
| overriding a configuration file.....                                     | 35       |
| example .....                                                            | 36       |
| <b>P</b>                                                                 |          |
| P routers <i>See</i> provider routers                                    |          |
| packet encapsulation                                                     |          |
| E1 interfaces .....                                                      | 105      |
| E3 interfaces .....                                                      | 108      |
| Layer 2 circuits.....                                                    | 326      |
| Layer 2 VPNs.....                                                        | 326      |
| serial interfaces .....                                                  | 122      |
| T1 interfaces .....                                                      | 115      |
| T3 interfaces .....                                                      | 119      |
| packets                                                                  |          |
| applying CoS scheduling rules .....                                      | 503      |
| handling packet fragments.....                                           | 449      |
| handling packet fragments (configuration<br>editor).....                 | 461      |
| PADI .....                                                               | 85       |
| PADO .....                                                               | 86       |
| PADR .....                                                               | 86       |
| PADS .....                                                               | 86       |
| PADT.....                                                                | 86       |
| PPPoE discovery .....                                                    | 85, 146  |
| RIP, description .....                                                   | 205      |
| PADI packets .....                                                       | 85       |
| PADO packets.....                                                        | 86       |
| PADR packets .....                                                       | 86       |
| PADS packets .....                                                       | 86       |
| PADT packets.....                                                        | 86       |
| parentheses, in syntax descriptions .....                                | xxiv     |
| passive routes, rejection, in static routing .....                       | 225      |
| password                                                                 |          |
| for OSPFv2 authentication .....                                          | 266      |
| for RIPv2 authentication .....                                           | 245      |
| patching a configuration file .....                                      | 35       |
| path cost metrics                                                        |          |
| for OSPF routes, description .....                                       | 209, 252 |
| for OSPF routes, modifying .....                                         | 264      |
| for RIP routes, description .....                                        | 235      |
| for RIP routes, modifying .....                                          | 242      |
| path selection <i>See</i> traffic engineering database                   |          |
| path-vector protocol <i>See</i> BGP                                      |          |
| pd-0/0/0 interface .....                                                 | 97       |
| PE (provider edge) routers.....                                          | 322      |
| description .....                                                        | 305      |
| ES-IS for a CLNS island .....                                            | 349      |
| route distinguishers.....                                                | 335      |
| verifying Layer 2 circuit connections .....                              | 343      |
| verifying Layer 2 circuit interfaces.....                                | 343      |
| verifying Layer 2 VPN connections .....                                  | 343      |

- verifying Layer 2 VPN interfaces ..... 343
  - verifying Layer 3 VPN connections ..... 343
  - VPN task overview ..... 324
  - VPN topology ..... 322
  - See also* VPNs
- pe-0/0/0 interface ..... 97
- peering sessions *See* BGP peers; BGP sessions
- penultimate hop popping (PHP) ..... 299
- penultimate router, in an LSP ..... 297
- permanent routes, adding ..... 223
- permanent virtual circuits (PVCs) ..... 81
- PGM (Pragmatic General Multicast) ..... 383
- PHP (penultimate hop popping) ..... 299
- Physical Interface Module *See* PIMs
- physical interface properties
  - BERT ..... 77
  - encapsulation ..... 80
  - FCS ..... 79
  - interface clocking ..... 77
  - key properties ..... 76
- PIC (PIM on a Services Router) *See* PIMs
- PIM (Protocol Independent Multicast)
  - dense mode ..... 382
  - disabling on the network management
    - interface ..... 388
  - RPF routing table group ..... 390
  - source-specific multicast (SSM) ..... 382
  - sparse mode ..... 382
  - static RP router ..... 388
  - supported versions ..... 385
  - verifying the mode ..... 393
  - verifying the RP ..... 393
- pimd interface ..... 97
- pime interface ..... 97
- PIMs (Physical Interface Modules)
  - output about, understanding ..... 48
  - PIM number, always 0 ..... 47
  - PIM slot number ..... 47
- ping
  - verifying link states ..... 135
  - VPN connection ..... 342
- ping command ..... 475
  - explanation ..... 475
- Ping Host page, output for BGP ..... 289
- ping mpls l2circuit interface command ..... 343
- ping mpls l2circuit virtual-circuit command ..... 343
- ping mpls l2vpn instance ..... 343
- ping mpls l2vpn interface command ..... 343
- ping mpls l3vpn command ..... 343
- ping trusted-nw-trusted-host ..... 470
  - explanation ..... 471
- ping untrusted-nw-untrusted-host command ..... 470
  - explanation ..... 471
- plesiochronous networks ..... 78
- 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 ..... 205
  - polarity, signal ..... 66
  - policers
    - description ..... 417
    - for firewall filter ..... 479
    - for stateless firewall filters ..... 455
  - policy *See* routing policies
  - pop label operation ..... 298
  - 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;
      - ISDN interfaces; loopback interfaces;
      - management interfaces; network
      - interfaces; services interfaces; special
      - interfaces
    - number in interface name ..... 48
    - T1 *See* T1 ports
    - T3 *See* T3 ports
- pp0
  - creating ..... 150
  - enabling CHAP ..... 152
  - information about ..... 155
  - interface description ..... 97
  - logical Ethernet interface on ..... 151
- PPP encapsulation
  - CHAP authentication ..... 83
  - CSU/DSU devices ..... 84
  - LCP connection process ..... 82
  - magic numbers ..... 84
  - NCPs ..... 83
  - overview ..... 82
- PPP over ATM *See* PPPoA
- PPP over ATM AAL5 multiplex encapsulation ..... 130
- PPP over Ethernet *See* PPPoE
- PPPoA (Point-to-Point Protocol over ATM)
  - CHAP ..... 131
  - logical encapsulation ..... 130
  - physical encapsulation ..... 129
  - verifying ATM-over-ADSL configuration ..... 141
- PPPoE (Point-to-Point Protocol over Ethernet) ..... 145
  - CHAP ..... 147, 152
  - client and server ..... 144
  - creating the pp0 interface ..... 150
  - discovery packets ..... 85, 146
  - encapsulation on an Ethernet interface ..... 85, 148
  - interfaces ..... 145
  - overview ..... 144
  - preparation ..... 147
  - sample topology ..... 145

|                                                         |         |
|---------------------------------------------------------|---------|
| service type .....                                      | 151     |
| session limit .....                                     | 86, 146 |
| session overview .....                                  | 86, 146 |
| session reconnection time .....                         | 151     |
| verifying interfaces .....                              | 155     |
| verifying sessions .....                                | 156     |
| verifying statistics .....                              | 157     |
| verifying version information .....                     | 157     |
| <i>See also</i> PPPoE over ATM-over-ADSL                |         |
| PPPoE Active Discovery Initiation (PADI) packets .....  | 85      |
| PPPoE Active Discovery Offer (PADO) packets .....       | 86      |
| PPPoE Active Discovery Request (PADR) packets .....     | 86      |
| PPPoE Active Discovery Session-Confirmation (PADS)      |         |
| packets .....                                           | 86      |
| PPPoE Active Discovery Termination (PADT) packets ..... | 86      |
| PPPoE encapsulation <i>See</i> PPPoE                    |         |
| PPPoE over ATM LLC encapsulation .....                  | 130     |
| PPPoE over ATM-over-ADSL .....                          | 145     |
| CHAP .....                                              | 152     |
| creating the pp0 interface .....                        | 150     |
| encapsulation .....                                     | 149     |
| preparation .....                                       | 147     |
| sample topology .....                                   | 145     |
| verifying configuration .....                           | 154     |
| <i>See also</i> PPPoE                                   |         |
| PPPoEoA <i>See</i> PPPoE over ATM-over-ADSL             |         |
| Pragmatic General Multicast .....                       | 383     |
| preferences                                             |         |
| for OSPF routes .....                                   | 264     |
| for static routes .....                                 | 224     |
| setting for static routes .....                         | 231     |
| prefix-length-range match type .....                    | 427     |
| primary stations, HDLC .....                            | 86      |
| propagation, suppressing .....                          | 433     |
| properties, verifying                                   |         |
| for ATM-over-ADSL network interfaces .....              | 137     |
| for network interfaces .....                            | 136     |
| protocol families                                       |         |
| ccc .....                                               | 89      |
| common protocol suites .....                            | 88      |
| inet .....                                              | 88      |
| inet6 .....                                             | 88      |
| ISO .....                                               | 88      |
| mlfr-end-to-end .....                                   | 89      |
| mlfr-uni-nni .....                                      | 89      |
| mlppp .....                                             | 89      |
| MPLS .....                                              | 88      |
| overview .....                                          | 88      |
| tcc .....                                               | 89      |
| tnp .....                                               | 89      |
| vpls .....                                              | 89      |
| Protocol Independent Multicast <i>See</i> PIM           |         |
| protocols                                               |         |
| Auto-RP .....                                           | 382     |
| BGP <i>See</i> BGP                                      |         |
| distance vector <i>See</i> RIP                          |         |
| DVMRP .....                                             | 381     |
| EGPs .....                                              | 198     |
| EIA-530 .....                                           | 68      |
| IGMP <i>See</i> IGMP                                    |         |
| IGPs .....                                              | 198     |
| IPSec <i>See</i> IPSec                                  |         |
| LDP <i>See</i> LDP                                      |         |
| MPLS <i>See</i> MPLS                                    |         |
| MSDP .....                                              | 383     |
| multicast <i>See</i> multicast                          |         |
| NAT <i>See</i> NAT                                      |         |
| OSPF <i>See</i> OSPF                                    |         |
| overview .....                                          | 193     |
| path vector <i>See</i> BGP                              |         |
| PGM .....                                               | 383     |
| PIM dense mode <i>See</i> PIM                           |         |
| PIM source-specific multicast (SSM) .....               | 382     |
| PIM sparse mode <i>See</i> PIM                          |         |
| PPPoE <i>See</i> PPPoE                                  |         |
| RIP <i>See</i> RIP                                      |         |
| RS-232 .....                                            | 68      |
| RS-422/449 .....                                        | 69      |
| RSVP <i>See</i> RSVP                                    |         |
| SAP and SDP <i>See</i> SAP; SDP                         |         |
| serial .....                                            | 67      |
| V.35 .....                                              | 69      |
| X.21 .....                                              | 70      |
| provider edge routers <i>See</i> PE routers             |         |
| provider routers .....                                  | 322     |
| description .....                                       | 305     |
| VPN task overview .....                                 | 324     |
| VPN topology .....                                      | 322     |
| <i>See also</i> VPNs                                    |         |
| push label operation .....                              | 298     |
| PVCs (permanent virtual circuits) .....                 | 81      |

## Q

|                                     |     |
|-------------------------------------|-----|
| queuing rules, CoS .....            | 503 |
| Quick Configuration                 |     |
| BGP page .....                      | 275 |
| buttons .....                       | 8   |
| E1 Interfaces page .....            | 104 |
| E3 Interfaces page .....            | 107 |
| Fast Ethernet Interfaces page ..... | 112 |
| Interfaces page .....               | 102 |
| IPSec Tunnels page .....            | 359 |
| network interfaces .....            | 102 |
| OSPF page .....                     | 253 |
| overview .....                      | 7   |
| RIP page .....                      | 237 |
| serial Interfaces page .....        | 121 |
| Static Routes page .....            | 226 |
| Summary page .....                  | 7   |
| T1 Interfaces page .....            | 114 |

- T3 (DS3) Interfaces page ..... 118
- R**
- radio buttons
    - Delete Configuration Below This Point ..... 11
    - Discard All Changes ..... 11
    - Discard Changes Below This Point ..... 11
  - RADIUS authentication, of PPP sessions ..... 147
  - random early detection *See* RED drop profiles
  - reactivate command ..... 29
  - RED (random early detection) drop profiles ..... 495
    - samples ..... 494
  - redistributing routes ..... 429
  - Refresh button ..... 11
  - rejecting invalid routes ..... 428
  - relative option ..... 35
  - release notes, URL ..... xxi
  - Remote Authentication Dial-In User Service (RADIUS)
    - authentication, of PPP sessions ..... 147
  - remote tunnel endpoint, IPsec ..... 360
  - rename command ..... 27
  - renaming configuration identifiers ..... 27
  - repeaters, on LAN segments ..... 53
  - 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 *See* RSVP
  - reset button, for return to factory configuration
    - See* CONFIG button
  - Resource Reservation Protocol *See* RSVP
  - reverse-path forwarding *See* RPF
  - rewrite rules
    - description ..... 417
    - replacing DSCPs ..... 486
    - sample rules ..... 486
    - when applied ..... 420
  - RIB *See* routing table
  - RIP (Routing Information Protocol)
    - authentication (RIPv2 only) ..... 236
    - authentication (RIPv2 only), configuring ..... 245
    - basic network (configuration editor) ..... 239
    - designating RIP interfaces ..... 238
    - distance vector protocol ..... 203
    - efficiency techniques ..... 205
    - enabling (Quick Configuration) ..... 238
    - maximum hop count ..... 204
    - overview ..... 203, 235
    - packets ..... 205
    - path cost metrics *See* path cost metrics
    - poison reverse technique ..... 205
    - Quick Configuration ..... 236
    - requirements ..... 236
    - routing policy (configuration editor) ..... 239
    - sample network with incoming metric ..... 242
    - sample network with outgoing metric ..... 244
    - sample topology ..... 239
    - split horizon technique ..... 205
    - supported versions ..... 203
    - traffic control with metrics *See* path cost metrics
    - traffic control with metrics, configuring ..... 242
    - unidirectional limitations ..... 206
    - verifying host reachability ..... 248
    - verifying RIP message exchange ..... 247
    - verifying RIP-enabled interfaces ..... 247
  - RIP neighbors, verifying ..... 247
  - RIP page ..... 237
    - field summary ..... 238
  - 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 ..... 217
    - BGP, update messages ..... 214
    - description ..... 201
    - external, EBGP ..... 214
    - internal, IBGP ..... 215
    - LSAs ..... 208
    - stub areas and NSSAs, to control ..... 211
  - route aggregation ..... 201
  - route distinguishers
    - description ..... 306
    - formats for ..... 335
  - route injection ..... 428
  - route list match types ..... 426
  - route manipulation actions, routing policies ..... 403
  - route redistribution ..... 428
  - route reflectors *See* BGP route reflectors
  - route selection
    - BGP process ..... 215
    - BGP, determining by AS path ..... 217
    - BGP, determining by local preference ..... 216
    - BGP, determining by MED metric ..... 218
    - BGP, lowest origin value preferred ..... 217
    - static routes, defining ..... 229
  - route target, in a VPN routing instance ..... 335
  - route targets, VPN ..... 307

|                                                                                    |          |
|------------------------------------------------------------------------------------|----------|
| route-flap damping .....                                                           | 433      |
| router <i>See</i> Services Router                                                  |          |
| routing.....                                                                       | 193      |
| advertisements .....                                                               | 201      |
| aggregation .....                                                                  | 201      |
| BGP <i>See</i> BGP                                                                 |          |
| configuring PPPoE .....                                                            | 143      |
| configuring VPNs .....                                                             | 321      |
| dynamic.....                                                                       | 200      |
| filtering and classifying routes .....                                             | 397      |
| filtering routes with policies.....                                                | 423      |
| filtering traffic through a firewall.....                                          | 437      |
| forwarding tables .....                                                            | 199      |
| from one source to many destinations .....                                         | 385      |
| in multiple ASs with BGP .....                                                     | 273      |
| in one AS with OSPF .....                                                          | 251      |
| in one AS with RIP.....                                                            | 235      |
| MPLS for VPNs.....                                                                 | 293      |
| MPLS traffic engineering .....                                                     | 309      |
| multicast <i>See</i> multicast                                                     |          |
| neighbors <i>See</i> BGP peers; OSPF neighbors; RIP<br>neighbors                   |          |
| OSPF <i>See</i> OSPF                                                               |          |
| overriding default packet forwarding with CoS ..                                   | 477      |
| protecting local IP addresses with NAT .....                                       | 437      |
| protocol overview.....                                                             | 193      |
| RIP <i>See</i> RIP                                                                 |          |
| RIP statistics .....                                                               | 247      |
| routing tables .....                                                               | 199      |
| static <i>See</i> static routing                                                   |          |
| through IPsec tunnels .....                                                        | 357      |
| VPNs .....                                                                         | 321      |
| <i>See also</i> protocols; routing policies; routing<br>solutions                  |          |
| Routing Engine                                                                     |          |
| handling packet fragments for (configuration<br>editor).....                       | 459      |
| protecting against DoS attacks (configuration<br>editor).....                      | 453      |
| protecting against untrusted services and<br>protocols (configuration editor)..... | 450      |
| routing information base <i>See</i> routing table                                  |          |
| Routing Information Protocol <i>See</i> RIP                                        |          |
| routing instance                                                                   |          |
| for CLNS static routes (with IS-IS) .....                                          | 348      |
| for CLNS static routes (without IS-IS) .....                                       | 352      |
| VPN configuration .....                                                            | 335      |
| VPN route target .....                                                             | 335      |
| VRF instances.....                                                                 | 306      |
| VRF table.....                                                                     | 335      |
| routing policies                                                                   |          |
| actions .....                                                                      | 402      |
| applying.....                                                                      | 404      |
| BGP export, for CLNS.....                                                          | 350      |
| BGP routing policy (configuration editor).....                                     | 281      |
| components.....                                                                    | 399      |
| configuration tasks .....                                                          | 424      |
| default actions .....                                                              | 404      |
| export statement .....                                                             | 404      |
| final actions.....                                                                 | 404      |
| forwarding class with source and destination ....                                  | 430      |
| grouping source and destination prefixes .....                                     | 430      |
| import statement .....                                                             | 404      |
| injecting routes from one protocol into another ..                                 | 428      |
| Layer 2 VPN export policy.....                                                     | 339      |
| Layer 2 VPN import policy .....                                                    | 338      |
| Layer 3 VPNs.....                                                                  | 341      |
| making BGP routes less preferable .....                                            | 431      |
| match conditions .....                                                             | 400      |
| overview .....                                                                     | 399      |
| policy name.....                                                                   | 425      |
| preparation.....                                                                   | 424      |
| prepending AS paths .....                                                          | 431      |
| reducing update messages with flap damping ...                                     | 433      |
| rejecting invalid routes .....                                                     | 426      |
| RIP routing policy (configuration editor) .....                                    | 239      |
| route redistribution .....                                                         | 428      |
| route-flap damping .....                                                           | 433      |
| terms .....                                                                        | 400      |
| terms, creating.....                                                               | 425      |
| VPN configuration .....                                                            | 337      |
| routing protocols <i>See</i> protocols                                             |          |
| routing solutions                                                                  |          |
| BGP confederations, for scaling problems.....                                      | 284      |
| BGP route reflectors, for scaling problems .....                                   | 281      |
| BGP scaling techniques.....                                                        | 218      |
| controlling designated router election .....                                       | 267      |
| controlling OSPF route cost .....                                                  | 264      |
| controlling OSPF route selection.....                                              | 264      |
| controlling RIP traffic with the incoming metric ..                                | 242      |
| controlling RIP traffic with the outgoing metric ..                                | 243      |
| CoS with DiffServ .....                                                            | 413, 477 |
| designated router, to reduce flooding.....                                         | 208      |
| directing BGP traffic by local preference .....                                    | 216      |
| filtering unwanted services and protocols.....                                     | 450      |
| firewall filters and NAT .....                                                     | 404, 437 |
| handling packet fragments.....                                                     | 449      |
| handling packet fragments (configuration<br>editor).....                           | 459      |
| making BGP routes less preferable .....                                            | 431      |
| MPLS traffic engineering .....                                                     | 309      |
| multicast administrative scoping .....                                             | 381      |
| multicast reverse-path forwarding (RPF) .....                                      | 380      |
| multicast shortest-path tree (SPT).....                                            | 381      |
| NSSAs, to control route advertisement .....                                        | 211      |
| path cost metrics, for packet flow control <i>See</i> path<br>cost metrics         |          |
| point-to-point sessions over Ethernet.....                                         | 143      |
| poison reverse, for traffic reduction .....                                        | 205      |
| preventing multicast routing loops .....                                           | 380      |



- protecting against DoS attacks ..... 453
- reducing update messages with flap damping ... 433
- rejecting invalid routes ..... 426
- routing policies ..... 399, 423
- securing OSPF routing (OSPFv2 only) ..... 265
- split horizon, for traffic reduction ..... 205
- static route control techniques ..... 224
- stub areas, to control route advertisement ..... 211
- VPNs ..... 321
- routing table
  - controlling static routes in ..... 224, 231
  - description ..... 199
  - displaying static routes in ..... 233
  - RPF group, for multicast ..... 390
  - sample distance-vector routing ..... 204
  - updates, limitations in RIP ..... 206
  - verifying for RPF ..... 394
  - verifying LDP-signaled LSPs ..... 318
  - verifying OSPF routes ..... 270
  - verifying RSVP-signaled LSPs ..... 320
- RP (rendezvous point)
  - static ..... 388
  - verifying ..... 393
- RPF (reverse-path forwarding)
  - description ..... 380
  - routing table group ..... 390
  - verifying the routing table ..... 394
- RS-232 ..... 68
- RS-422/449 ..... 69
- RS-530 ..... 68
- RSVP (Resource Reservation Protocol)
  - and OSPF for VPNs ..... 333
  - bandwidth reservation ..... 301
  - CSPF ..... 303
  - disabling CSPF ..... 316
  - EROs ..... 301
  - fundamentals ..... 301
  - link coloring ..... 303
  - overview ..... 310
  - requirements ..... 310
  - RSVP-signaled LSPs ..... 313
  - verifying LSPs ..... 320
  - verifying neighbors ..... 319
  - verifying sessions ..... 319
  - verifying the routing table on the entry router ... 320
- RSVP neighbors, verifying ..... 319
- RSVP-signaled LSP *See* RSVP
- run command ..... 34

## S

- S,G notation, for multicast forwarding states ..... 379
- S/T interfaces
  - overview ..... 74
  - PIMs ..... 161

- samples ..... 141, 154
  - CLNS VPN configuration ..... 354
  - firewall filter configurations ..... 466
  - PPPoA for ATM-over-ADSL configuration ..... 141
  - PPPoE over ATM-over-ADSL configuration ..... 154
  - See also* networks; topology
- SAP (Session Announcement Protocol)
  - description ..... 383
  - session announcements ..... 386
  - verifying ..... 392
- saving, configuration files ..... 37
- scaling BGP *See* BGP confederations; BGP route reflectors
- schedulers
  - assigning resources ..... 497
  - default settings ..... 418
  - description ..... 417
  - mapping to forwarding classes ..... 501
  - sample mappings ..... 501
  - sample schedulers ..... 497
- scheduling a commit ..... 31
- scope, IPv6 addresses
  - global unicast ..... 93
  - link-local unicast ..... 93
  - multicast types ..... 93
  - site-local unicast ..... 93
- scoping, administrative ..... 381
- SDP (Session Discovery Protocol)
  - description ..... 383
  - session announcements ..... 386
  - verifying ..... 392
- secondary stations, HDLC ..... 87
- secret, CHAP *See* CHAP, local identity
- security
  - IPSec tunnels ..... 357
  - MD5 authentication for OSPF ..... 266
  - MD5 authentication for RIPv2 ..... 246
  - password authentication for OSPFv2 ..... 266
  - password authentication for RIPv2 ..... 245
- security association *See* IPSec security associations
- serial interfaces ..... 64
  - clocking modes ..... 66
  - connection process ..... 65
  - DTE default clock rate reduction ..... 67
  - EIA-530 ..... 68
  - inverting the transmit clock ..... 67
  - line protocols ..... 67
  - RS-232 ..... 68
  - RS-422/449 ..... 69
  - signal polarity ..... 66
  - transmission signals ..... 65
  - V.35 ..... 69
  - X.21 ..... 70
  - See also* serial ports
- serial numbers, in MAC addresses ..... 50

|                                                                     |     |                                                           |              |
|---------------------------------------------------------------------|-----|-----------------------------------------------------------|--------------|
| serial ports .....                                                  | 64  | sessions .....                                            |              |
| CHAP .....                                                          | 122 | announcements, multicast .....                            | 386          |
| clock rate .....                                                    | 123 | BGP session establishment .....                           | 214          |
| clocking .....                                                      | 123 | BGP session maintenance .....                             | 214          |
| clocking, inverting the transmit clock .....                        | 123 | ISDN session establishment .....                          | 75           |
| configuring .....                                                   | 120 | LDP, verifying .....                                      | 317          |
| encapsulation type .....                                            | 122 | limit on PPPoE sessions .....                             | 146          |
| line speed .....                                                    | 123 | PPPoE .....                                               | 86, 146      |
| logical interfaces .....                                            | 122 | PPPoE, reconnection time .....                            | 151          |
| <i>See also</i> serial interfaces                                   |     | RSVP, verifying .....                                     | 319          |
| service classes, corresponding DSCPs .....                          | 414 | shortest path first algorithm .....                       | 207          |
| service sets .....                                                  |     | shortest-path tree .....                                  | 381          |
| for IPsec tunnels .....                                             | 362 | show access command .....                                 | 141          |
| for NAT rules .....                                                 | 447 | show bgp group command .....                              | 287          |
| for stateful firewall filters .....                                 | 447 | explanation .....                                         | 288          |
| service types, naming for PPPoE .....                               | 151 | show bgp neighbor command .....                           | 286          |
| services interfaces .....                                           |     | explanation .....                                         | 287          |
| applying a NAT rule to (configuration editor) .....                 | 447 | show bgp summary command .....                            | 288          |
| applying a stateful firewall filter to (configuration editor) ..... | 447 | explanation .....                                         | 289          |
| CRTP .....                                                          | 99  | show chassis hardware command .....                       | 48           |
| for IPsec tunnels .....                                             | 361 | show class-of-service adaptive-shaper command .....       | 509          |
| MLFR .....                                                          | 99  | show class-of-service interface command .....             | 509–510      |
| MLFR FRF.15 and FRF.16 .....                                        | 99  | show class-of-service virtual-channel command .....       | 510          |
| MLPPP .....                                                         | 99  | show class-of-service virtual-channel-group command ..... | 510          |
| overview .....                                                      | 99  | show cli history command .....                            | 34           |
| Services Router .....                                               |     | show command .....                                        | 25, 354      |
| as a PPPoE client .....                                             | 144 | show firewall command .....                               | 466          |
| BGP routing .....                                                   | 273 | show firewall filter protect-RE command .....             | 472          |
| CLNS VPNs .....                                                     | 345 | show firewall log command .....                           | 471          |
| configuration tools .....                                           | 3   | explanation .....                                         | 472          |
| CoS overview .....                                                  | 413 | show igmp interface command .....                         | 392          |
| CoS with DiffServ .....                                             | 477 | explanation .....                                         | 393          |
| CPE, with PPPoE .....                                               | 143 | show interfaces bc-0/0/4 extensive command .....          | 184          |
| <i>See also</i> PPPoE                                               |     | show interfaces br-6/0/0 extensive command .....          | 183          |
| firewall filter overview .....                                      | 404 | show interfaces command .....                             | 141, 154     |
| firewall filters .....                                              | 437 | show interfaces dc-0/0/4 extensive command .....          | 186          |
| interfaces overview .....                                           | 41  | show interfaces detail command .....                      | 136          |
| IPsec tunnels .....                                                 | 357 | show interfaces dl0 extensive command .....               | 187          |
| ISDN connections .....                                              | 159 | show interfaces extensive command .....                   | 137          |
| MPLS for VPNs overview .....                                        | 293 | explanation, for ATM-over-ADSL interfaces .....           | 140          |
| MPLS traffic engineering .....                                      | 309 | explanation, for ISDN interfaces .....                    | 184, 186–187 |
| multicast .....                                                     | 385 | show interfaces lo0 command .....                         | 465          |
| multicast overview .....                                            | 375 | show interfaces ppo command .....                         | 155          |
| NAT .....                                                           | 437 | show isdn status command .....                            | 183          |
| network interfaces .....                                            | 101 | show ldp neighbor command .....                           | 316          |
| OSPF routing .....                                                  | 251 | explanation .....                                         | 317          |
| PPPoE .....                                                         | 143 | show ldp session detail command .....                     | 317          |
| RIP routing .....                                                   | 235 | explanation .....                                         | 318          |
| routing policies .....                                              | 423 | show multicast rpf command .....                          | 394          |
| routing policy overview .....                                       | 399 | explanation .....                                         | 394          |
| routing protocols overview .....                                    | 193 | show ospf interface command .....                         | 268          |
| static routing .....                                                | 223 | explanation .....                                         | 269          |
| VPNs .....                                                          | 321 | show ospf neighbor command .....                          | 269          |
| Session Announcement Protocol <i>See</i> SAP; SDP                   |     | explanation .....                                         | 269          |

- show ospf route command ..... 270
  - results ..... 271
- show pim interface command ..... 393
  - explanation ..... 393
- show pim rps command ..... 393
  - explanation ..... 394
- show pppoe interfaces command ..... 156
- show pppoe statistics command ..... 157
- show pppoe version command ..... 157
- show rip neighbor command ..... 247
  - explanation ..... 247
- show rip statistics command ..... 247
- show route summary command ..... 473, 475
  - explanation ..... 474, 476
- show route table inet.3 command ..... 318, 320
  - explanation ..... 318, 320
- show route terse command ..... 233
  - explanation ..... 234
- show rsvp neighbor command ..... 319
  - explanation ..... 319
- show rsvp session detail command ..... 320
  - explanation ..... 320
- show sap listen command ..... 392, 509
  - explanation ..... 392, 509
- show services command ..... 466
- show services ipsec-vpn ipsec statistics command ... 371
  - explanation ..... 371
- show system reboot command ..... 34
- signaling protocols ..... 309
  - overview ..... 300
  - VPNs ..... 331
  - See also* LDP; MPLS traffic engineering; RSVP
- signals
  - DS1 ..... 56
  - E1 loopback (control) ..... 59
  - explicit clocking signal transmission ..... 78
  - multiplexing DS1 into DS2 signal ..... 59
  - serial polarity ..... 66
  - serial transmission ..... 65
  - T1 loopback (control) ..... 59
  - V.35 ..... 69
  - X.21 ..... 70
- single-area network, OSPF ..... 256
- site-local unicast IPv6 addresses ..... 93
- source-specific multicast ..... 382
- sp-0/0/0
  - for IPsec tunnels (configuration editor) ..... 361
  - interface description ..... 97
  - no stateful firewall filters ..... 406
- sparse mode *See* multicast routing modes
- special interfaces
  - CRTP ..... 99
  - dsc interface ..... 97
  - IPv4 addressing ..... 89
  - IPv6 addressing ..... 92
  - logical properties ..... 87
  - loopback interface ..... 98
  - management interface ..... 98
  - MLFR ..... 99
  - MLFR FRF.15 and FRF.16 ..... 99
  - MLPPP ..... 99
  - names ..... 47
  - naming conventions ..... 46
  - output, understanding ..... 48
  - overview ..... 95
  - physical properties ..... 76
  - protocol families ..... 88
  - services interfaces ..... 99
  - summary ..... 95
- SPF (shortest path first) algorithm ..... 207
- split horizon technique ..... 205
- SPT (shortest-path tree) ..... 381
- ssh command ..... 473
  - explanation ..... 474
- stateful firewall filter rules
  - for IPsec tunnels (configuration editor) ..... 366
- stateful firewall filters
  - actions ..... 407
  - applying to an interface (configuration editor) ... 447
  - automatic discard rule ..... 405
  - configuration editor ..... 442, 444
  - configuration overview ..... 406
  - description ..... 405
  - do not apply to sp-0/0/0 ..... 406
  - enabling (Quick Configuration) ..... 441
  - junos-algs-outbound default group ..... 406
  - match conditions ..... 407
  - preparation ..... 438
  - Quick Configuration ..... 438
  - sample rules ..... 443
  - untrusted network ..... 406
  - verifying ..... 470
  - verifying actions ..... 473
- stateless firewall filters
  - actions and action modifiers ..... 412
  - applying to an interface (configuration editor) ... 464
  - automatic discard rule ..... 405, 408
  - bit-field logical operators ..... 412
  - description ..... 405
  - handling packet fragments ..... 449
  - handling packet fragments (configuration editor) ..... 459
  - match conditions ..... 409
  - planning ..... 408, 449
  - policers for ..... 455
  - preparation ..... 438
  - protecting the Routing Engine against ICMP floods (configuration editor) ..... 453
  - protecting the Routing Engine against TCP floods (configuration editor) ..... 453

|                                                                                       |       |
|---------------------------------------------------------------------------------------|-------|
| protecting the Routing Engine against untrusted protocols (configuration editor)..... | 450   |
| protecting the Routing Engine against untrusted services (configuration editor).....  | 450   |
| sample terms, to filter fragments.....                                                | 460   |
| sample terms, to filter services and protocols.....                                   | 450   |
| sample terms, to protect against DoS attacks.....                                     | 454   |
| typical, planning.....                                                                | 449   |
| statements                                                                            |       |
| adding or modifying.....                                                              | 26    |
| copying.....                                                                          | 27    |
| deactivating.....                                                                     | 29    |
| deleting.....                                                                         | 26    |
| replacing.....                                                                        | 35    |
| static LSPs.....                                                                      | 299   |
| static routes                                                                         |       |
| CLNS VPNs (with IS-IS).....                                                           | 348   |
| CLNS VPNs (without IS-IS).....                                                        | 352   |
| configuring basic routes (configuration editor).....                                  | 228   |
| controlling.....                                                                      | 224   |
| controlling in routing and forwarding tables.....                                     | 231   |
| default properties.....                                                               | 225   |
| default properties, setting.....                                                      | 232   |
| defining route selection.....                                                         | 229   |
| preferences.....                                                                      | 224   |
| preventing readvertisement.....                                                       | 225   |
| qualified next hops.....                                                              | 224   |
| Quick Configuration.....                                                              | 226   |
| rejecting passive traffic.....                                                        | 225   |
| requirements.....                                                                     | 226   |
| route retention.....                                                                  | 225   |
| sample preferred path.....                                                            | 230   |
| sample stub network.....                                                              | 228   |
| verifying.....                                                                        | 233   |
| Static Routes page.....                                                               | 226   |
| field summary.....                                                                    | 227   |
| static routing                                                                        |       |
| default gateway.....                                                                  | 227   |
| description.....                                                                      | 200   |
| overview.....                                                                         | 223   |
| <i>See also</i> static routes                                                         |       |
| static RP router.....                                                                 | 388   |
| <i>See also</i> RP                                                                    |       |
| statistics                                                                            |       |
| ATM-over-ADSL interfaces.....                                                         | 141   |
| firewall filters.....                                                                 | 472   |
| IPSec tunnels.....                                                                    | 371   |
| PPPoE.....                                                                            | 157   |
| RIP.....                                                                              | 247   |
| status command.....                                                                   | 22    |
| status, link states, verifying.....                                                   | 135   |
| strict hops, RSVP.....                                                                | 302   |
| stub areas                                                                            |       |
| area ID (configuration editor).....                                                   | 260   |
| area ID (Quick Configuration).....                                                    | 254   |
| area type (Quick Configuration).....                                                  | 255   |
| controlling OSPF route cost.....                                                      | 265   |
| creating (configuration editor).....                                                  | 261   |
| description.....                                                                      | 211   |
| example.....                                                                          | 212   |
| sample topology.....                                                                  | 262   |
| sub-ASs, BGP.....                                                                     | 221   |
| subautonomous systems, BGP.....                                                       | 221   |
| subnet masks.....                                                                     | 91    |
| subnets <i>See</i> subnetworks                                                        |       |
| subnetworks                                                                           |       |
| description.....                                                                      | 198   |
| IPv4 subnets.....                                                                     | 90    |
| multicast leaves and branches.....                                                    | 378   |
| route aggregation.....                                                                | 202   |
| Summary Quick Configuration page.....                                                 | 7     |
| superframe framing.....                                                               | 58    |
| support, technical <i>See</i> technical support                                       |       |
| SVCs (switched virtual circuits).....                                                 | 81    |
| swap and push label operation.....                                                    | 299   |
| swap label operation.....                                                             | 298   |
| switch types, ISDN.....                                                               | 164   |
| switched virtual circuits (SVCs).....                                                 | 81    |
| switches, on LAN segments.....                                                        | 53    |
| synchronous networks.....                                                             | 77    |
| syntax conventions.....                                                               | xxiii |
| system clock <i>See</i> clocking                                                      |       |
| <b>T</b>                                                                              |       |
| T1 interfaces.....                                                                    | 55    |
| AMI encoding.....                                                                     | 57    |
| B8ZS encoding.....                                                                    | 57    |
| D4 framing.....                                                                       | 58    |
| data stream.....                                                                      | 56    |
| encoding.....                                                                         | 57    |
| ESF framing.....                                                                      | 58    |
| framing.....                                                                          | 58    |
| loopback.....                                                                         | 59    |
| overview.....                                                                         | 56    |
| signals.....                                                                          | 56    |
| superframe framing.....                                                               | 58    |
| <i>See also</i> T1 ports                                                              |       |
| T1 ports.....                                                                         | 55    |
| adding CRTP.....                                                                      | 133   |
| cable length.....                                                                     | 117   |
| CHAP.....                                                                             | 115   |
| clocking.....                                                                         | 115   |
| configuring.....                                                                      | 113   |
| data inversion.....                                                                   | 116   |
| encapsulation type.....                                                               | 115   |
| fractional, channel number.....                                                       | 48    |
| frame checksum.....                                                                   | 117   |
| framing.....                                                                          | 116   |
| logical interfaces.....                                                               | 115   |
| MTU.....                                                                              | 115   |

- overview .....56
- time slots.....116
- See also* T1 interfaces
- T3 interfaces .....59
  - bit stuffing .....60
  - data stream .....59
  - DS3 framing .....60
  - multiplexing on .....60
  - overview .....59
  - See also* T3 ports
- T3 ports .....59
  - C-bit parity .....120
  - cable length .....120
  - CHAP .....119
  - clocking .....119
  - configuring.....117
  - encapsulation type .....119
  - frame checksum .....120
  - framing.....120
  - logical interfaces.....119
  - MTU .....119, 122
  - overview .....59
  - See also* T3 interfaces
- tap interface.....97
- tcc protocol family .....89
- TCP policers .....455
- technical support
  - contacting JTAC .....xxvi
- TED *See* traffic engineering database
- telnet command .....474
  - explanation .....475
- terminology
  - CLNS .....345
  - configuration.....3
  - CoS .....397
  - firewall filters .....397
  - interfaces.....42
  - ISDN.....159
  - MPLS .....293
  - multicast .....375
  - ports.....42
  - PPPoE .....143
  - routing .....193
  - routing policies .....397
  - VPNs .....293
- terms
  - firewall filter, for multifold classifier.....481
  - in a routing policy .....400
  - in a routing policy, creating .....425
- three-way handshake .....208
- through route list match type .....427
- time slots
  - E1 .....106
  - number in interface name .....48
  - T1.....116
- tnp protocol family .....89
- to statement, routing policy match conditions .....400
- top command .....25
- topology
  - data link layer.....49
  - IPv4 subnets .....91
  - PPPoE session on an ATM-over-ADSL loop .....146
  - PPPoE session on an Ethernet loop.....145
  - sample ATM-over-ADSL.....72
  - sample BGP AS path.....217
  - sample BGP confederation .....285
  - sample BGP confederations.....222
  - sample BGP external and internal links .....280
  - sample BGP local preference use .....216
  - sample BGP MED use .....218
  - sample BGP peer network.....278
  - sample BGP peer session.....213
  - sample BGP route reflector (one cluster) .....219, 282
  - sample BGP route reflectors (cluster of clusters) ..221
  - sample BGP route reflectors (multiple clusters) ..220
  - sample distance-vector routing .....204
  - sample Frame Relay network.....80
  - sample ISDN network .....73
  - sample LAN .....94
  - sample LSP network.....297
  - sample multiarea OSPF routing.....210
  - sample OSPF backbone area.....211
  - sample OSPF multiarea network .....258
  - sample OSPF network.....270
  - sample OSPF network with stubs and NSSAs ....212
  - sample OSPF single-area network .....257
  - sample OSPF stub areas and NSSAs.....262
  - sample poison reverse routing .....206
  - sample RIP network.....239
  - sample RIP network with incoming metric .....242
  - sample RIP network with outgoing metric .....244
  - sample route advertisement .....201
  - sample route aggregation .....202
  - sample router network .....199
  - sample RSVP-signaled LSP .....302
  - sample split horizon routing .....205
  - sample static route .....200
  - sample static route, preferred path .....230
  - sample stub network for static routes .....228
  - sample unidirectional routing.....207
  - sample VLAN .....95
  - sample VPN.....322
- topology database, OSPF .....251
- Traceroute page
  - results for OSPF.....272
  - results for RIP .....249
- traceroute source bypass-routing gateway
  - command.....318
  - explanation .....319

|                                                             |     |
|-------------------------------------------------------------|-----|
| traffic                                                     |     |
| controlling with incoming RIP metric                        | 242 |
| controlling with outgoing RIP metric                        | 243 |
| outgoing, securing                                          | 358 |
| traffic engineering <i>See</i> MPLS traffic engineering     |     |
| traffic engineering database                                |     |
| CSPF constraints on path selection                          | 303 |
| CSPF rules for path selection                               | 303 |
| link coloring for CSPF path selection                       | 303 |
| transit interfaces                                          |     |
| LDP-signaled LSPs for                                       | 311 |
| RSVP-signaled LSPs for                                      | 313 |
| transit routers, in an LSP                                  | 297 |
| transmit clock source <i>See</i> clocking                   |     |
| trusted networks, firewall filter protection                | 404 |
| tunnels, through a public network <i>See</i> IPSec tunnels; |     |
| VPNs                                                        |     |
| two-dimensional parity                                      | 79  |
| types of interfaces                                         | 47  |

## U

|                                                |     |
|------------------------------------------------|-----|
| U interface                                    |     |
| overview                                       | 75  |
| PIMs                                           | 161 |
| unicast IPv6 addresses                         | 93  |
| untrusted networks, firewall filter actions on | 404 |
| up command                                     | 24  |
| uploading a configuration file                 | 14  |
| upstream interfaces                            | 378 |
| <i>See also</i> multicast                      |     |
| upto route list match type                     | 427 |
| UR-2 operating mode                            | 129 |
| URLs                                           |     |
| release notes                                  | xxi |

## V

|                                      |     |
|--------------------------------------|-----|
| V.35                                 | 69  |
| variable-length subnet masks (VLSMs) | 91  |
| VCI (virtual channel identifier)     |     |
| ATM-over-ADSL interfaces             | 131 |
| PPPoE over ATM-over-ADSL interfaces  | 150 |
| verification                         |     |
| adaptive shaping                     | 509 |
| ATM-over-ADSL interface properties   | 137 |
| B-channels                           | 184 |
| BGP configuration                    | 288 |
| BGP groups                           | 287 |
| BGP peer reachability                | 289 |
| BGP peers (neighbors)                | 286 |
| CLNS VPNs                            | 354 |
| configuration syntax                 | 30  |
| D-channel                            | 186 |
| dialer interfaces                    | 187 |
| firewall filter actions              | 473 |
| firewall filter flood protection     | 474 |

|                                                 |     |
|-------------------------------------------------|-----|
| firewall filter handles fragments               | 475 |
| firewall filter operation                       | 471 |
| firewall filters                                | 465 |
| firewall statistics                             | 472 |
| IGMP version                                    | 392 |
| IPSec tunnel operation                          | 371 |
| ISDN interfaces                                 | 183 |
| ISDN status                                     | 183 |
| LDP neighbors                                   | 316 |
| LDP sessions                                    | 317 |
| LDP-signaled LSP                                | 318 |
| MPLS traffic engineering                        | 316 |
| multicast SAP and SDP                           | 392 |
| multicast session announcements                 | 509 |
| network interfaces                              | 135 |
| OSPF host reachability                          | 271 |
| OSPF neighbors                                  | 269 |
| OSPF routes                                     | 270 |
| OSPF-enabled interfaces                         | 268 |
| PIM mode and interface configuration            | 393 |
| PIM RP address                                  | 393 |
| PIM RPF routing table                           | 394 |
| PPPoA for ATM-over-ADSL configuration           | 141 |
| PPPoE interfaces                                | 155 |
| PPPoE over ATM-over-ADSL configuration          | 154 |
| PPPoE sessions                                  | 156 |
| PPPoE statistics                                | 157 |
| PPPoE version                                   | 157 |
| RIP host reachability                           | 248 |
| RIP message exchange                            | 247 |
| RIP-enabled interfaces                          | 247 |
| RSVP neighbors                                  | 319 |
| RSVP sessions                                   | 319 |
| RSVP-signaled LSP                               | 320 |
| stateful firewall filters                       | 470 |
| static routes in the routing table              | 233 |
| traffic forwarding over LDP-signaled LSPs       | 318 |
| virtual channel                                 | 510 |
| virtual channel group                           | 510 |
| VPNs                                            | 342 |
| version                                         |     |
| OSPF, supported                                 | 208 |
| PPPoE, verifying                                | 157 |
| RIP, supported                                  | 203 |
| View Configuration Text page                    | 13  |
| virtual channel identifier <i>See</i> VCI       |     |
| virtual channels, applying CoS rules to logical |     |
| interfaces                                      | 503 |
| virtual circuit ID, for Layer 2 circuits        | 334 |
| virtual circuits                                |     |
| DLCIs                                           | 81  |
| overview                                        | 81  |
| PVCs                                            | 81  |
| SVCs                                            | 81  |
| virtual LANs <i>See</i> VLANs                   |     |

- virtual link, through the backbone area ..... 210
  - virtual path identifier (VPI), PPPoE over ATM-over-ADSL  
  interfaces ..... 150
  - virtual private networks *See* VPNs
  - VLANs (virtual LANs) ..... 94
    - LAN comparison ..... 94
    - overview ..... 94
    - topology ..... 95
  - VLSMs (variable-length subnet masks) ..... 91
  - VPI, PPPoE over ATM-over-ADSL interfaces ..... 150
  - vpls protocol family ..... 89
  - VPN routing and forwarding (VRF) instances ..... 306
  - VPN routing and forwarding table *See* VRF table
  - VPNs (virtual private networks) ..... 321
    - AS number ..... 330
    - basic Layer 2 circuit description ..... 323
    - basic Layer 2 VPN description ..... 322
    - basic Layer 3 VPN description ..... 323
    - BGP ..... 329
    - CLNS *See* CLNS
    - components ..... 304
    - configuration overview ..... 321
    - configuration task overview ..... 324
    - IGPs ..... 331
    - Layer 2 circuit configuration ..... 334
    - LSP for RSVP ..... 328
    - MPLS ..... 327
    - overview ..... 293, 304
    - participating interfaces ..... 325
    - preparation ..... 324
    - protocols for ..... 327
    - route distinguishers ..... 306, 335
    - route target ..... 335
    - route targets ..... 307
    - routing information ..... 306
    - routing instance ..... 335
      - routing instance, for CLNS static routes (with  
  IS-IS) ..... 348
      - routing instance, for CLNS static routes (without  
  IS-IS) ..... 352
    - routing policies ..... 337
    - routing requirements ..... 305
    - sample topology ..... 322
    - signaling protocols ..... 331
    - tunneling process ..... 305
    - types ..... 307
    - verifying connectivity ..... 342
    - VRF instances ..... 306
    - VRF table *See* VRF table
    - See also* Layer 2 circuits; Layer 2 VPNs; Layer 3  
  VPNs; MPLS
  - VRF (VPN routing and forwarding) table ..... 335
    - route targets ..... 307
    - VRF instances ..... 306
  - VRF instances ..... 306
- X**
- X.21 ..... 70