

J-series™ Services Router

Configuration Guide

Release 7.2

Juniper Networks, Inc.

1194 North Mathilda Avenue

Sunnyvale, California 94089

USA

408-745-2000

www.juniper.net

Part Number: 530-013028-01, Revision 1

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

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

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

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

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

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

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

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

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

J-series™ Services Router Configuration Guide, Release 7.2
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
13 April 2005—Revision 1.

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

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

Products made or sold by Juniper Networks (including the ERX-310, ERX-705, ERX-710, ERX-1410, ERX-1440, M5, M7i, M10, M10i, M20, M40, M40e, M160, M320, and T320 routers, T640 routing node, and the JUNOS and SDX-300 software) or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. The JUNOS software has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

SOFTWARE LICENSE

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

End User License Agreement

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

1. **The Parties.** The parties to this Agreement are Juniper Networks, Inc. and its subsidiaries (collectively "Juniper"), and the person or organization that originally purchased from Juniper or an authorized Juniper reseller the applicable license(s) for use of the Software ("Customer") (collectively, the "Parties").

2. **The Software.** In this Agreement, "Software" means the program modules and features of the Juniper or Juniper-supplied software, and updates and releases of such software, for which Customer has paid the applicable license or support fees to Juniper or an authorized Juniper reseller.

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

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

b. Customer shall use the Software on a single hardware chassis having a single processing unit, or as many chassis or processing units for which Customer has paid the applicable license fees.

c. 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

xvii

Part 1

Using the Configuration Interfaces

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

Part 2

Configuring Router Interfaces

Chapter 2	Configuring Network Interfaces	41
Chapter 3	Configuring Point-to-Point Protocol over Ethernet	77

Part 3

Configuring Routing Protocols

Chapter 4	Routing Overview	97
Chapter 5	Configuring Static Routes	127
Chapter 6	Configuring a RIP Network	139
Chapter 7	Configuring an OSPF Network	155
Chapter 8	Configuring BGP Sessions	177

Part 4

Configuring Private Communications over Public Networks with MPLS

Chapter 9	Multiprotocol Label Switching Overview	197
Chapter 10	Configuring Signaling Protocols for Traffic Engineering	213

Chapter 11	Configuring Virtual Private Networks ..	227
Chapter 12	Configuring IPSec for Secure Packet Exchange ..	251
Part 5	Managing Multicast Transmissions	
Chapter 13	Multicast Overview ..	269
Chapter 14	Configuring a Multicast Network ..	279
Part 6	Configuring Routing Policy, Firewall Filters, and Class of Service	
Chapter 15	Policy, Firewall Filter, and Class-of-Service Overview ..	291
Chapter 16	Configuring Routing Policies ..	317
Chapter 17	Configuring Firewall Filters and NAT ..	331
Chapter 18	Configuring Class of Service with DiffServ ..	371
Part 7	Index	

Table of Contents

About This Guide	xvii
Objectives	xvii
Audience.....	xviii
How to Use This Guide	xviii
Document Conventions	xix
Related Juniper Networks Documentation.....	xx
Documentation Feedback.....	xxii
Requesting Support.....	xxii

Part 1

Using the Configuration Interfaces

Chapter 1	Using J-series 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
Entering Operational Mode Commands During Configuration	33
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	Configuring Network Interfaces	41
	Network Interfaces Terms	41
	Interfaces Overview	44
	Network Interface Types	44
	Interfaces and Interface Naming	45
	Before You Begin	47
	Configuring Network Interfaces with Quick Configuration	47
	Configuring an E1 Interface with Quick Configuration	49
	Configuring a Fast Ethernet Interface with Quick Configuration	52
	Configuring a T1 Interface with Quick Configuration	53
	Configuring a T3 Interface with Quick Configuration	57
	Configuring a Serial Interface with Quick Configuration	60
	Configuring Network Interfaces with a Configuration Editor	64
	Adding a Network Interface with a Configuration Editor	64
	Adding an ATM-for-ADSL Network Interface with a Configuration Editor ..	66
	Deleting a Network Interface with a Configuration Editor	69
	Verifying Interface Configuration	70
	Verifying the Link State of All Interfaces	70
	Verifying Interface Properties	71
	Verifying ADSL Interface Properties	72
Chapter 3	Configuring Point-to-Point Protocol over Ethernet	77
	PPPoE Terms	77
	PPPoE Overview	78
	PPPoE Interfaces	79
	Fast Ethernet Interface	79
	ATM-for-ADSL Interface	79
	PPPoE Stages	80
	PPPoE Discovery Stage	80
	PPPoE Session Stage	81

Optional CHAP Authentication	81
Before You Begin	82
Configuring PPPoE with a Configuration Editor	82
Setting the Appropriate Encapsulation on the Interface (Required)	82
Configuring PPPoE Encapsulation on an Ethernet Interface	83
Configuring PPPoE Encapsulation on an ATM-for-ADSL Interface	83
Configuring a PPPoE Interface (Required)	85
Configuring CHAP (Optional)	87
Verifying a PPPoE Configuration	88
Displaying a PPPoE Configuration for an ATM-for-ADSL Interface	89
Verifying PPPoE Interfaces	90
Verifying PPPoE Sessions	91
Verifying the PPPoE Version	92
Verifying PPPoE Statistics	92

Part 3

Configuring Routing Protocols

Chapter 4

Routing Overview	97
Routing Terms	97
Routing Overview	101
Networks and Subnetworks	102
Autonomous Systems	102
Interior and Exterior Gateway Protocols	102
Routing Tables	103
Forwarding Tables	103
Dynamic and Static Routing	104
Route Advertisements	105
Route Aggregation	105
RIP Overview	107
Distance-Vector Routing Protocols	107
Maximizing Hop Count	108
RIP Packets	109
Split Horizon and Poison Reverse Efficiency Techniques	109
Limitations of Unidirectional Connectivity	110
OSPF Overview	111
Link-State Advertisements	112
Role of the Designated Router	112
Path Cost Metrics	113
Areas and Area Border Routers	113
Role of the Backbone Area	114
Stub Areas and Not-So-Stubby Areas	115
BGP Overview	116
Point-to-Point Connections	117
BGP Messages for Session Establishment	118
BGP Messages for Session Maintenance	118
IBGP and EBGP	118
Route Selection	119
Local Preference	120
AS Path	121
Origin	121
Multiple Exit Discriminator	122

	Scaling BGP for Large Networks	122
	Route Reflectors—for Added Hierarchy	123
	Confederations—for Subdivision	125
Chapter 5	Configuring Static Routes	127
	Static Routing Overview	127
	Static Route Preferences	128
	Qualified Next Hops	128
	Control of Static Routes	128
	Route Retention	129
	Readvertisement Prevention	129
	Forced Rejection of Passive Route Traffic	129
	Default Properties	129
	Before You Begin	130
	Configuring Static Routes with Quick Configuration	130
	Configuring Static Routes with a Configuration Editor	132
	Configuring a Basic Set of Static Routes (Required)	132
	Controlling Static Route Selection (Optional)	133
	Controlling Static Routes in the Routing and Forwarding Tables (Optional)	135
	Defining Default Behavior for All Static Routes (Optional)	136
	Verifying the Static Route Configuration	137
	Displaying the Routing Table	137
Chapter 6	Configuring a RIP Network	139
	RIP Overview	139
	RIP Traffic Control with Metrics	139
	Authentication	140
	Before You Begin	140
	Configuring a RIP Network with Quick Configuration	140
	Configuring a RIP Network with a Configuration Editor	143
	Configuring a Basic RIP Network (Required)	143
	Controlling Traffic in a RIP Network (Optional)	146
	Controlling Traffic with the Incoming Metric	146
	Controlling Traffic with the Outgoing Metric	147
	Enabling Authentication for RIP Exchanges (Optional)	149
	Enabling Authentication with Plain-Text Passwords	149
	Enabling Authentication with MD5 Authentication	150
	Verifying the RIP Configuration	151
	Verifying the RIP-Enabled Interfaces	151
	Verifying the Exchange of RIP Messages	151
	Verifying Reachability of All Hosts in the RIP Network	152
Chapter 7	Configuring an OSPF Network	155
	OSPF Overview	155
	Enabling OSPF	155
	OSPF Areas	156
	Path Cost Metrics	156
	Before You Begin	156
	Configuring an OSPF Network with Quick Configuration	156

Configuring an OSPF Network with a Configuration Editor	160
Configuring the Router Identifier (Required).....	160
Configuring a Single-Area OSPF Network (Required)	161
Configuring a Multiarea OSPF Network (Optional)	162
Creating the Backbone Area.....	163
Creating Additional OSPF Areas.....	163
Configuring Area Border Routers	164
Configuring Stub and Not-So-Stubby Areas (Optional).....	165
Tuning an OSPF Network for Efficient Operation	167
Controlling Route Selection in the Forwarding Table.....	167
Controlling the Cost of Individual Network Segments	168
Enabling Authentication for OSPF Exchanges	169
Controlling Designated Router Election	170
Verifying an OSPF Configuration	171
Verifying OSPF-Enabled Interfaces	171
Verifying OSPF Neighbors	172
Verifying the Number of OSPF Routes	173
Verifying Reachability of All Hosts in an OSPF Network.....	174

Chapter 8 Configuring BGP Sessions 177

BGP Overview.....	177
BGP Peering Sessions.....	177
IBGP Full Mesh Requirement.....	178
Route Reflectors and Clusters	178
BGP Confederations	178
Before You Begin.....	179
Configuring BGP Sessions with Quick Configuration	179
Configuring BGP Sessions with a Configuration Editor	181
Configuring a Point-to-Point Peering Session (Required).....	181
Configuring BGP Within a Network (Required)	184
Configuring a Route Reflector (Optional)	185
Configuring BGP Confederations (Optional)	188
Verifying a BGP Configuration	190
Verifying BGP Neighbors	190
Verifying BGP Groups	191
Verifying BGP Summary Information	192
Verifying Reachability of All Peers in a BGP Network	193

Part 4 Configuring Private Communications over Public Networks with MPLS

Chapter 9 Multiprotocol Label Switching Overview 197

MPLS and VPN Terms	197
MPLS Overview	199
Label Switching	200
Label-Switched Paths	200
Label-Switching Routers	201
Labels	202
Label Operations.....	202
Penultimate Hop Popping	203

	LSP Establishment	203
	Static LSPs	203
	Dynamic LSPs	203
	Signaling Protocols Overview	204
	Label Distribution Protocol	204
	LDP Operation	204
	LDP Messages	204
	Resource Reservation Protocol	204
	RSVP Fundamentals	205
	Bandwidth Reservation Requirement	205
	Explicit Route Objects	205
	Constrained Shortest Path First	207
	Link Coloring	207
	VPN Overview	208
	VPN Components	208
	VPN Routing Requirements	209
	VPN Routing Information	210
	VRF Instances	210
	Route Distinguishers	210
	Route Targets to Control the VRF Table	211
	Types of VPNs	211
	Layer 2 VPNs	211
	Layer 2 Circuits	211
	Layer 3 VPNs	211
Chapter 10	Configuring Signaling Protocols for Traffic Engineering	213
	Signaling Protocol Overview	213
	LDP Signaling Protocol	214
	RSVP Signaling Protocol	214
	Before You Begin	214
	Configuring LDP and RSVP with a Configuration Editor	215
	Configuring LDP-Signaled LSPs	215
	Configuring RSVP-Signaled LSPs	217
	Verifying an MPLS Configuration	220
	Verifying an LDP-Signaled LSP	220
	Verifying LDP Neighbors	220
	Verifying LDP Sessions	221
	Verifying the Presence of LDP-Signaled LSPs	222
	Verifying Traffic Forwarding over the LDP-Signaled LSP	222
	Verifying an RSVP-Signaled LSP	223
	Verifying RSVP Neighbors	223
	Verifying RSVP Sessions	224
	Verifying the Presence of RSVP-Signaled LSPs	224
Chapter 11	Configuring Virtual Private Networks	227
	VPN Configuration Overview	227
	Sample VPN Topology	228
	Basic Layer 2 VPN Configuration	228
	Basic Layer 2 Circuit Configuration	229
	Basic Layer 3 VPN Configuration	229
	Before You Begin	230

Configuring VPNs with a Configuration Editor	230
Configuring Interfaces Participating in a VPN	231
Configuring Protocols Used by a VPN	233
Configuring MPLS for VPNs	233
Configuring a BGP Session	235
Configuring Routing Options for VPNs	236
Configuring an IGP and a Signaling Protocol	237
Configuring LDP for Signaling	237
Configuring RSVP for Signaling	239
Configuring a Layer 2 Circuit	240
Configuring a VPN Routing Instance	241
Configuring a VPN Routing Policy	243
Configuring a Routing Policy for Layer 2 VPNs	244
Configuring a Routing Policy for Layer 3 VPNs	247
Verifying a VPN Configuration	248
Pinging a Layer 2 VPN	249
Pinging a Layer 3 VPN	249
Pinging a Layer 2 Circuit	249

Chapter 12 Configuring IPsec for Secure Packet Exchange 251

IPsec Tunnel Overview	251
Security Associations	252
Securing Incoming Traffic	252
Translating Outgoing Traffic	252
Before You Begin	252
Configuring an IPsec Tunnel with Quick Configuration	252
Configuring an IPsec Tunnel with a Configuration Editor	254
Configuring IPsec Services Interfaces	255
Configuring IPsec Service Sets	256
Configuring an IPsec Stateful Firewall Filter	260
Configuring a NAT Pool	262
Verifying the IPsec Tunnel Configuration	264
Verifying IPsec Tunnel Statistics	265

Part 5 Managing Multicast Transmissions

Chapter 13 Multicast Overview 269

Multicast Terms	269
Multicast Architecture	272
Upstream and Downstream Interfaces	272
Subnetwork Leaves and Branches	272
Multicast IP Address Ranges	273
Notation for Multicast Forwarding States	273
Dense and Sparse Routing Modes	274
Strategies for Preventing Routing Loops	274
Reverse-Path Forwarding for Loop Prevention	274
Shortest-Path Tree for Loop Prevention	275
Administrative Scoping for Loop Prevention	275
Multicast Protocol Building Blocks	275

Chapter 14	Configuring a Multicast Network	279
	Before You Begin.....	280
	Configuring a Multicast Network with a Configuration Editor.....	280
	Configuring SAP and SDP (Optional)	280
	Configuring IGMP (Required).....	281
	Configuring the PIM Static RP (Optional)	282
	Configuring a PIM RPF Routing Table (Optional)	284
	Verifying a Multicast Configuration.....	285
	Verifying SAP and SDP Addresses and Ports.....	286
	Verifying the IGMP Version.....	286
	Verifying the PIM Mode and Interface Configuration	287
	Verifying the PIM RP Configuration	287
	Verifying the RPF Routing Table Configuration	288
 Part 6	 Configuring Routing Policy, Firewall Filters, and Class of Service	
 Chapter 15	 Policy, Firewall Filter, and Class-of-Service Overview	 291
	Policy, Firewall Filter, and CoS Terms	291
	Routing Policy Overview	293
	Routing Policy Components.....	293
	Routing Policy Terms.....	294
	Routing Policy Match Conditions.....	294
	Routing Policy Actions	296
	Default and Final Actions.....	298
	Applying Routing Policies	298
	Firewall Filter Overview	298
	Stateful and Stateless Firewall Filters.....	299
	Process for Configuring a Stateful Firewall Filter and NAT	300
	Summary of Stateful Firewall Filter and NAT Match Conditions and Actions.....	300
	Planning a Stateless Firewall Filter	302
	Stateless Firewall Filter Match Conditions, Actions, and Action Modifiers.....	303
	Class-of-Service Overview.....	307
	Benefits of DiffServ CoS	307
	DSCPs and Forwarding Service Classes	308
	JUNOS CoS Functions.....	309
	How Forwarding Classes and Schedulers Work.....	311
	Default Forwarding Class Queue Assignments.....	311
	Default Scheduler Settings.....	312
	Default Behavior Aggregate (BA) Classifiers	313
	DSCP Rewrites.....	314
	Sample BA Classification	314
 Chapter 16	 Configuring Routing Policies	 317
	Before You Begin.....	318
	Configuring a Routing Policy with a Configuration Editor	318

	Configuring the Policy Name (Required)	319
	Configuring a Policy Term (Required)	319
	Rejecting Known Invalid Routes (Optional)	320
	Injecting OSPF Routes into the BGP Routing Table (Optional)	322
	Grouping Source and Destination Prefixes in a Forwarding Class (Optional)	324
	Configuring a Policy to Prepend the AS Path (Optional)	325
	Configuring Damping Parameters (Optional)	327
Chapter 17	Configuring Firewall Filters and NAT	331
	Before You Begin	332
	Configuring a Stateful Firewall Filter with Quick Configuration	332
	Configuring a Stateful Firewall Filter with a Configuration Editor	336
	Configuring a Stateless Firewall Filter with a Configuration Editor	342
	Stateless Firewall Filter Strategies	343
	Strategy for a Typical Stateless Firewall Filter	343
	Strategy for Handling Packet Fragments	343
	Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources	344
	Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods	347
	Configuring a Routing Engine Firewall Filter to Handle Fragments	353
	Applying a Stateless Firewall Filter to an Interface	358
	Verifying Firewall Filter Configuration	359
	Displaying Firewall Filter Configurations	359
	Verifying a Stateful Firewall Filter	364
	Displaying Firewall Filter Logs	365
	Displaying Firewall Filter Statistics	366
	Verifying a Services, Protocols, and Trusted Sources Firewall Filter	367
	Verifying a TCP and ICMP Flood Firewall Filter	368
	Verifying a Firewall Filter That Handles Fragments	369
Chapter 18	Configuring Class of Service with DiffServ	371
	Before You Begin	372
	Configuring CoS with DiffServ with a Configuration Editor	372
	Configuring a Policer for a Firewall Filter (Required)	373
	Configuring and Applying a Firewall Filter for a Multifield Classifier (Required)	374
	Assigning Forwarding Classes to Output Queues (Required)	378
	Configuring and Applying Rewrite Rules (Required)	379
	Configuring and Applying Behavior Aggregate Classifiers (Required)	384
	Configuring RED Drop Profiles for Assured Forwarding Congestion Control (Required)	388
	Configuring Schedulers (Optional)	390
	Configuring and Applying Scheduler Maps (Optional)	394
	Configuring and Applying Virtual Channels (Optional)	397
	Configuring and Applying Adaptive Shaping (Optional)	401
	Verifying a DiffServ Configuration	402
	Verifying Multicast Session Announcements	402
	Verifying an Adaptive Shaper Configuration	403

Part 7

Index

Index..... 407

About This Guide

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

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

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 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 7.2 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 xx.

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>
“Configuring Network Interfaces”	■ <i>JUNOS Network Interfaces and Class of Service Configuration Guide</i>
“Configuring Point-to-Point Protocol over Ethernet”	■ <i>JUNOS Network and Services Interfaces Command Reference</i>
“Routing Overview”	■ <i>JUNOS Routing Protocols Configuration Guide</i>
“Configuring Static Routes”	■ <i>JUNOS Protocols, Class of Service, and System Basics 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 Configuration Guide</i>
“Configuring Virtual Private Networks”	■ <i>JUNOS VPNs Configuration Guide</i>
“Configuring IPSec for Secure Packet Exchange”	■ <i>JUNOS System Basics Configuration Guide</i>
	■ <i>JUNOS Services Interfaces Configuration Guide</i>
	■ <i>JUNOS Network and Services Interfaces Command Reference</i>
“Multicast Overview”	<i>JUNOS Multicast Protocols Configuration Guide</i>
“Configuring a Multicast Network”	
“Policy, Firewall Filter, and Class-of-Service Overview”	<i>JUNOS Policy Framework Configuration Guide</i>
“Configuring Routing Policies”	
“Configuring Firewall Filters and NAT”	■ <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>

Chapter in a J-series Guide	Corresponding JUNOS Software Manual
"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 Protocols, Class of Service, and System Basics 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 Protocols, Class of Service, and System Basics Command Reference</i> ■ <i>JUNOS Network and Services Interfaces Command Reference</i>
"Monitoring Real-Time Performance"	<i>JUNOS Network and Services Interfaces Command Reference</i>

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can send your comments to techpubs-comments@juniper.net, or fill out the documentation feedback form at <http://www.juniper.net/techpubs/docbug/docbugreport.html>. If you are using e-mail, be sure to include the following information with your comments:

- Document name
- Document part number
- Page number
- Software release version

Requesting Support

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

Part 1

Using the Configuration Interfaces

- Using J-series Configuration Tools on page 3

Chapter 1

Using J-series Configuration Tools

Use J-series configuration tools to configure all services on a J-series Services 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 J-series 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

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)

Juniper NETWORKS

ROUTER - J4300

Logged in as: **regress**

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

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

Configuration > View and Edit > Edit Configuration

Configuration

([Expand all](#) | [Hide all](#))

- + [groups](#)
- + [system](#)
- + [interfaces](#)
- + [routing-options](#)
- + [protocols](#)
- + [policy-options](#)
- + [firewall](#)
- + [services](#)

Access [Configure](#)

Accounting options [Configure](#)

Applications [Configure](#)

Chassis [Configure](#)

Class of service [Configure](#)

Firewall [Edit](#) [Delete](#)

Forwarding options [Configure](#)

Interfaces [Edit](#) [Delete](#)

Policy options [Edit](#) [Delete](#)

Protocols [Edit](#) [Delete](#)

Routing instances [Configure](#)

Routing options [Edit](#) [Delete](#)

Security [Configure](#)

OK Cancel Refresh Commit... Discard...

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 or edit 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
Edit	Displays information for a configuration option that has already been configured, allowing you to edit a statement.
Configure	Displays information for a configuration option that has not been configured, allowing you to include a statement.
Add new entry	Displays fields and drop-down menus for a statement identifier, allowing you to add a new identifier to a statement.
<i>identifier</i>	Displays fields and drop-down menus 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.

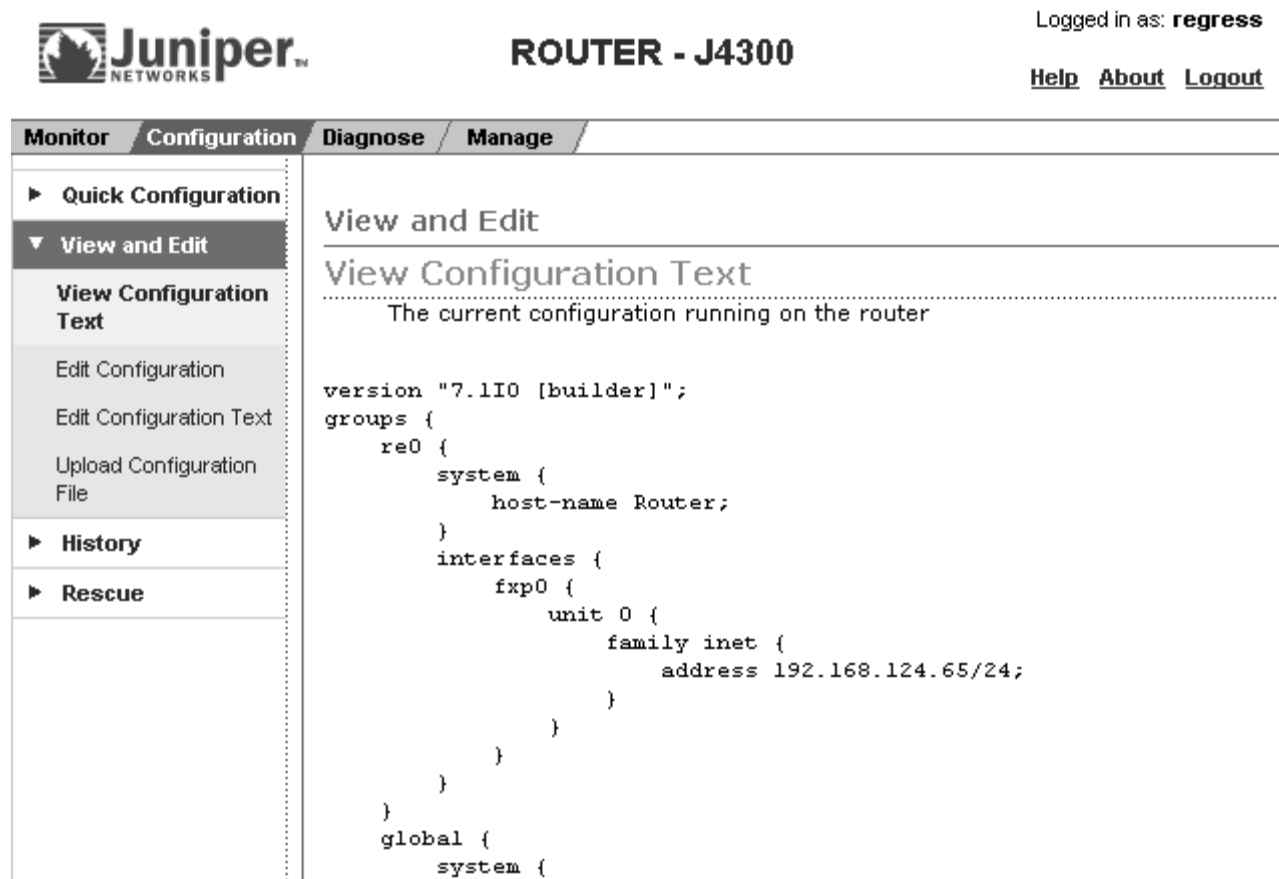
3. To display all the edits applied to the running configuration, click **Refresh**.

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 "7.1I10 [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"/>	<u>Current</u>	2005-01-18 16:12:46 PST	root	cli			<u>Download</u>
<input type="checkbox"/>	<u>1</u>	2005-01-18 15:01:13 PST	root	cli			<u>Download</u> <u>Rollback</u>

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.l.e87-ds4JDbSz6A0hcbs2goG"; ## SECRET-DATA	
[edit]	[edit]
	chassis { alarm { ethernet { link-down yellow; } } }
[edit interfaces fe-0/0/0 unit 0 family inet]	[edit interfaces fe-0/0/0 unit 0 family inet]
service { input { service-set jweb-wan-sfw-service-set; } output { service-set jweb-wan-sfw-service-set; } }	
[edit interfaces fe-0/0/0 unit 0 family inet]	[edit interfaces fe-0/0/0 unit 0 family inet]
	address 192.168.124.75/24;

Downloading a Configuration File

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

1. Select **Configuration > History**.

A list of the current and previous 49 configurations is displayed as Configuration History in the main pane (see Figure 6). Table 14 summarizes the Configuration History display.

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

3. Select the options your Web browser provides that allow you to save the configuration file to a target directory on your local system.

The file is saved as an ASCII file.

Loading a Previous Configuration File

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

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

1. Select **Configuration > History**.

A list of the current and previous 49 configurations is displayed as Configuration History in the main pane (see Figure 6). Table 14 summarizes the Configuration History display.

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

The main pane displays the results of the rollback operation.



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

Setting, Viewing, or Deleting the Rescue Configuration

If someone inadvertently commits a configuration that denies management access to 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 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
- Entering Operational Mode Commands During Configuration on page 33

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</code> operational mode command. (For details, see the <i>JUNOS Protocols, Class of Service, and System Basics 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 Protocols, Class of Service, and System Basics 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
```

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.

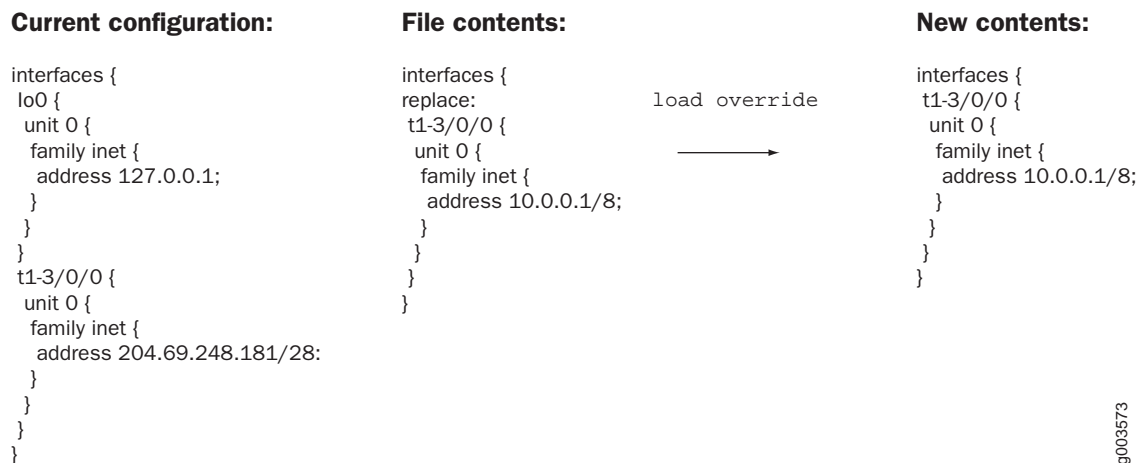
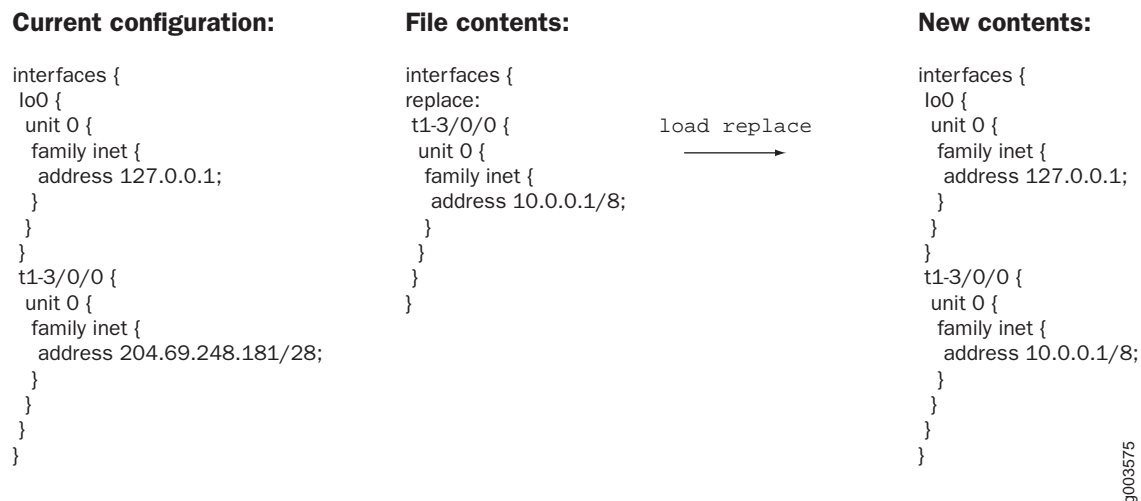
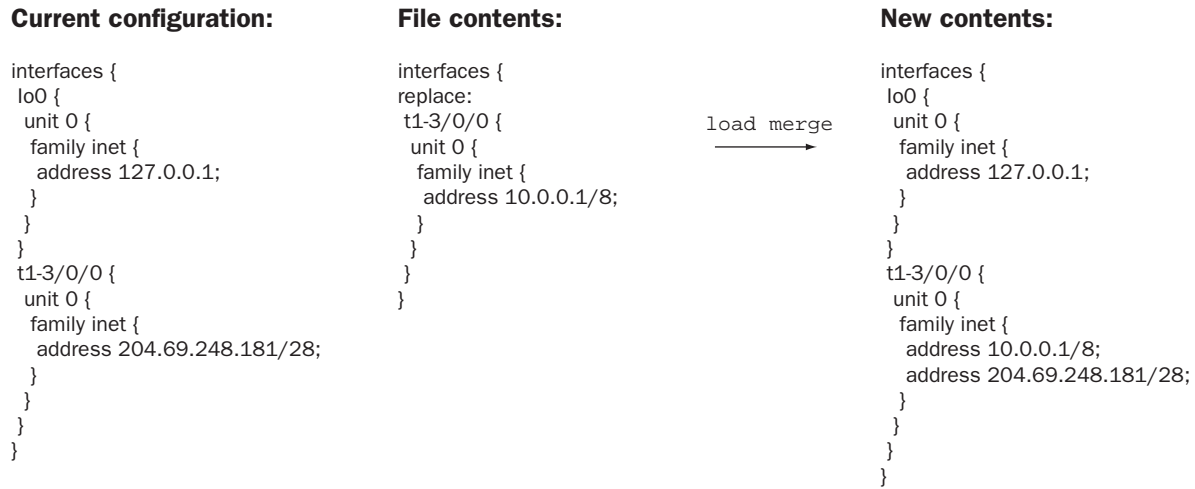
Figure 8: Loading a Configuration with the Override Operation**Figure 9: Loading a Configuration with the Replace Operation**

Figure 10: Loading a Configuration with the Merge Operation

9003574

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

- Configuring Network Interfaces on page 41
- Configuring Point-to-Point Protocol over Ethernet on page 77

Chapter 2

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.

In addition to network interfaces, Services Routers use permanent interfaces for internal communication, such as the services interfaces that provide additional features for regulating and manipulating traffic. For information about one of these interfaces, see the *J-series Services Router Getting Started Guide*.

This chapter includes the following topics. For more information about interfaces, see the *JUNOS Network Interfaces and Class of Service Configuration Guide*.

- Network Interfaces Terms on page 41
- Interfaces Overview on page 44
- Before You Begin on page 47
- Configuring Network Interfaces with Quick Configuration on page 47
- Configuring Network Interfaces with a Configuration Editor on page 64
- Verifying Interface Configuration on page 70

Network Interfaces Terms

To understand Services Router network 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.

Table 17: Network Interfaces Terms (Continued)

Term	Definition
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 9 Mbps and upstream (customer-to-provider) rates of up to 640 Kbps, depending on the implementation.
Annex A	ITU-T Standard G.992.1 that defines how ADSL works over plain old telephone lines (POTS).
Annex B	ITU-T Standard G.992.1 that defines how ADSL works over Integrated Services Digital Network (ISDN).
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 a T3 interface that allows a Services Router to connect with a channel service unit (CSU) with proprietary multiplexing at the remote end of the line. Subrating a T3 interface reduces the maximum allowable peak rate by limiting the payload encapsulated by the High-level Data Link Control protocol (HDLC).
data-link connection identifier (DLCI)	Identifier for a Frame Relay virtual connection, also called a logical interface.
data service unit (DSU)	Unit that connects a data terminal equipment (DTE) device—in this case, a Services Router—to a digital telephone line.
data terminal equipment (DTE)	RS-232 interface that a Services Router uses to exchange information with a serial device.
DS1	Digital signal 1, another name for a T1 interface.
DS3 interface	Digital signal 3, another name for a T3 interface.
data inversion	Transmission of all data bits in the data stream so that zeros are transmitted as ones and ones are transmitted as zeros. Data inversion is normally used only in alternate mark inversion (AMI) mode to guarantee ones density in the transmitted stream.
E1 interface	Physical WAN interface for transmitting signals in European digital transmission (E1) format. The E1 signal format carries information at a rate of 2.048 Mbps and can carry 32 channels of 64 Kbps each.
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.

Table 17: Network Interfaces Terms (Continued)

Term	Definition
Flexible PIM Concentrator (FPC)	Logical identifier for a Physical Interface Module (PIM) installed on a Services Router.
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 end-point 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 using inverse multiplexing. It is often used in conjunction with MLPPP.
Multilink Point-to-Point Protocol (MLPPP)	Protocol that allows you to bundle multiple PPP links into a single logical unit. It can be used to better utilize bandwidth and also has the advantages of reduced latency and improved fault tolerance.
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 T3 (DS3) WAN interface (J6300 model only) ■ ADSL WAN interface (optional) either Annex A to support DSL over POTS or Annex B which supports DSL over ISDN. (J4300 and J6300) ■ Two serial interfaces

Table 17: Network Interfaces Terms (Continued)

Term	Definition
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 the following cable types: <ul style="list-style-type: none"> ■ V.35—Serial cable with a 34-pin connector for speeds up to 8 Mbps ■ RS-232—(EIA-232) Standard serial cable with a 25-pin (DB-25) connector for speeds up to 110.5 Kbps ■ RS-422/449—(EIA-449) Serial cable with a 37-pin (DB-37) connector, for RS-422 and RS-423 interfaces ■ X.21—Standard serial cable, popular in Europe, with a 15-pin (DB-15) connector ■ RS-530—(EIA-530) Serial cable with a 25-pin connector for higher speeds than RS-232 <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.

Interfaces Overview

This section contains the following topics:

- Network Interface Types on page 44
- Interfaces and Interface Naming on page 45

Network Interface Types

J-series Services Routers support the following network interface types: E1, Fast Ethernet, serial, T1, and T3, as well as asynchronous transfer mode (ATM) over ADSL interfaces.



NOTE: J4300 and J6300 Services Routers with asymmetrical digital subscriber line (ADSL) Physical Interface Modules (PIMs) can use PPP over Ethernet (PPPoE) and PPP over ATM for ADSL (PPPoA) to connect through DSL lines only, not for direct ATM connections.

T3 interfaces, which are also called DS3 interfaces, are supported on J6300 Services Routers only.

ADSL interfaces are optional and support the following standards:

- ANSI T1.413 Issue II
- ETSI TS 101 388 V1.3.1
- ITU G.992.1

ADSL is available only on J4300 and J6300 Services Routers.

Interfaces and 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. Each interface has a unique name that 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-FPC / PIM / port

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

type-FPC / PIM / port : channel

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

type-FPC / PIM / port : <channel> . logical unit

For example, **e1-5/0/0** is the E1 interface on port 0 of FPC 5, **e1-5/0/0:15** is channel 15 on that interface, and **e1-5/0/0:15.0** is logical unit 0 on that channel.

The parts of an interface name are explained in Table 18.

Table 18: Interface Name Information

Interface Name Part	Meaning	Possible Values
<i>type</i>	Type of network medium that can connect to this interface.	<ul style="list-style-type: none"> ■ at—Asynchronous Transfer Mode (ATM) for ADSL WAN interface. ■ dsc—Virtual interface that discards packets. ■ e1—E1 WAN interface. ■ fe—Fast Ethernet LAN interface. ■ gr, gre—Generic routing encapsulation (GRE) interface for tunnel services. This interface is internally generated and not configurable. ■ ip, ipip—IP-over-IP interface. This interface is internally generated and not configurable. ■ ls—Link services interface. ■ lsi—Label-switched interface. This interface is internally generated and is not configurable. ■ lo—Loopback interface. This interface is internally generated and also configurable. ■ mtun—Multicast GRE interface. This interface is internally generated and not configurable. ■ pd, pimd—Protocol Independent Multicast (PIM) decapsulator interface. This interface is internally generated and not configurable. ■ pe, pime—PIM encapsulator interface. This interface is internally generated and not configurable. ■ se—Serial interface (including RS-232, RS-422/449, RS-530, V.35, and X.21 interfaces). ■ sp—Services interface. ■ t1—T1 (also called DS1) WAN interface. ■ t3—T3 (also called DS3) WAN interface. ■ tap—This interface is internally generated and not configurable.
<i>FPC</i>	The number of the Flexible PIM concentrator (FPC) on which the physical interface is located.	<ul style="list-style-type: none"> ■ On a J2300 router, always 0. ■ On a J4300 or J6300 router, a value from 0 through 6.
<i>PIM</i>	The number of the PIM on which the physical interface is located. For Services Router interfaces, the FPC and PIM are the same physical unit, so the PIM has no number of its own.	Always 0.

Table 18: Interface Name Information (Continued)

Interface Name Part	Meaning	Possible Values
<i>port</i>	The number of the port on a PIM on which the physical interface is located.	Either 0 or 1.
<i>channel</i>	The 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>logical unit</i>	The number of the logical unit created on the physical interface.	A value from 0 through 16384.

Before You Begin

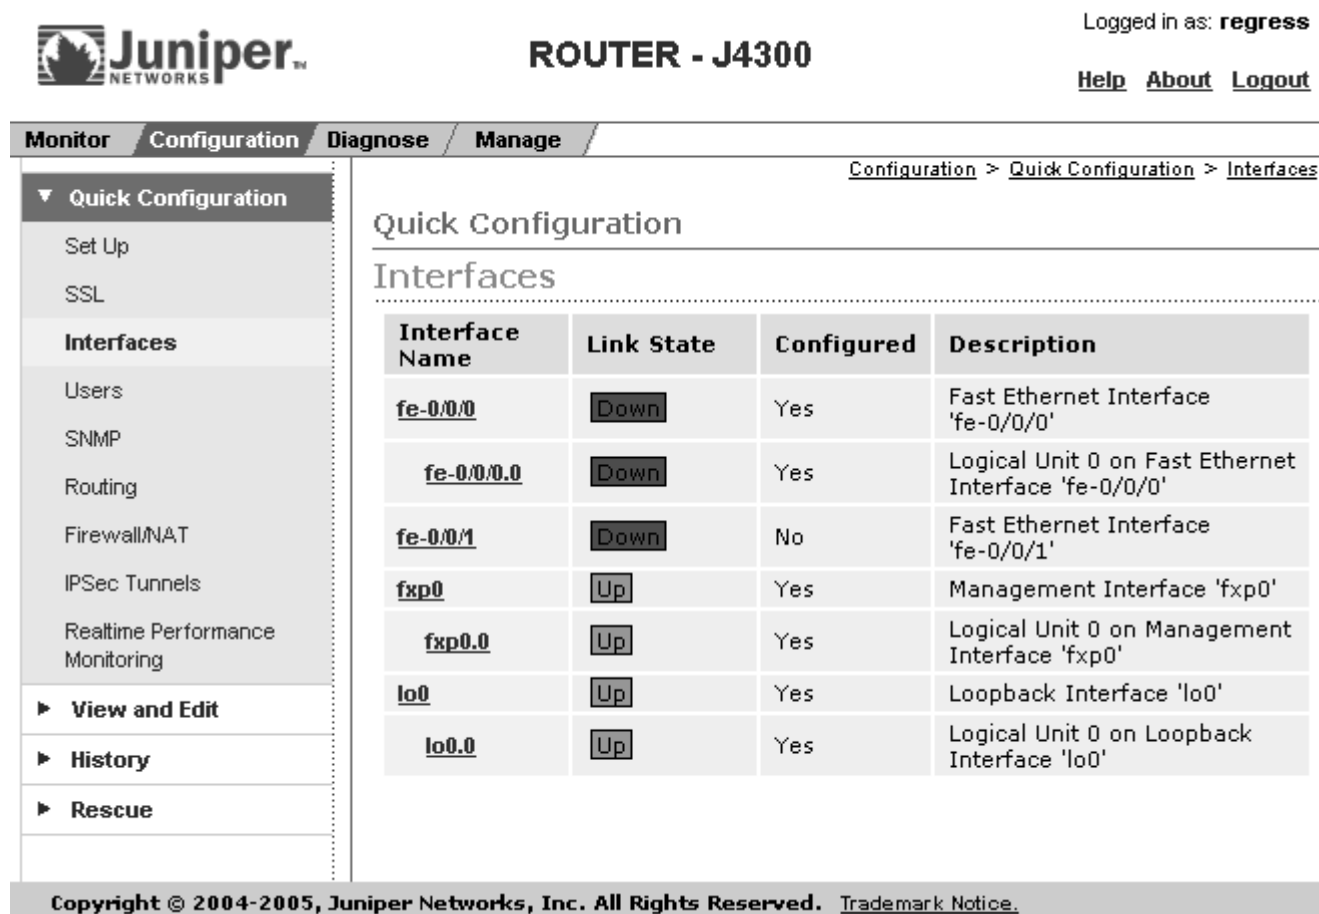
Before you configure network interfaces, you need to perform the following tasks:

- 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 44.

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 11.

Configuring Network Interfaces with Quick Configuration

The Quick Configuration page allows you to configure network interfaces on a Services Router, as shown in Figure 11.

Figure 11: Quick Configuration Interfaces Page


Router: **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 11. 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 49
 - Configuring a Fast Ethernet Interface with Quick Configuration on page 52
 - Configuring a T1 Interface with Quick Configuration on page 53

- Configuring a T3 Interface with Quick Configuration on page 57
- Configuring a Serial Interface with Quick Configuration on page 60

Configuring an E1 Interface with Quick Configuration

To configure properties on an E1 interface:

1. From the Quick Configuration page, as shown in Figure 11, select the E1 interface you want to configure.

The properties you can configure on an E1 interface are displayed, as shown in Figure 12.

Figure 12: E1 Interfaces Quick Configuration Page

The screenshot shows the Juniper Networks configuration interface for a Router - J6300. The user is logged in as 'regress'. The navigation menu on the left includes 'Monitor', 'Configuration', 'Diagnose', and 'Manage'. Under 'Configuration', the 'Quick Configuration' section is expanded, showing options like 'Set Up', 'SSL', 'Interfaces', 'Users', 'SNMP', 'Routing', 'Firewall/NAT', 'IPSec Tunnels', and 'Realtime Performance Monitoring'. The 'Interfaces' section is selected, and the 'Quick Configuration' page is displayed. The breadcrumb trail is 'Configuration > Quick Configuration > Interfaces'. The page title is 'Quick Configuration' and the interface is 'Physical Interface: 'e1-1/0/0''. The 'Logical Interfaces' section shows 'No logical interfaces configured.' with an 'Add...' button. The 'Physical Interface Description' field is empty. The 'Encapsulation' section has a dropdown menu and an 'Enable CHAP' checkbox. The 'CHAP Local Identity' section has a 'Use System Host Name' checkbox and fields for 'Local Name', '* CHAP Peer Identity', and '* CHAP Secret'.

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: 'e1-1/0/0'

Logical Interfaces

No logical interfaces configured.

[Add...](#)

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 19.
3. Click one of the following buttons:
 - To apply the configuration and stay in 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 70.

Table 19: 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.
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.
Encapsulation		
Encapsulation	Specifies the encapsulation type for traffic on the interface.	<p>From the drop-down 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.

Table 19: E1 Quick Configuration Summary (Continued)

Field	Function	Your Action
CHAP Local Identity (available if CHAP is enabled)		
Use System Host Name	Specifies that the E1 interface uses the Services Router's system hostname in CHAP challenge and response packets.	<ul style="list-style-type: none"> ■ To enable, select the check box (the default). ■ To disable, clear the check box.
Local Name	If Use System Host Name is disabled, specifies the local name for CHAP to use.	Type a local name for this E1 interface.
CHAP Peer Identity	Identifies the client or peer with which the Services Router communicates on this E1 interface.	Type the CHAP client name.
CHAP Secret	Specifies the secret password for CHAP authentication, known to both sides of the connection.	Type a password that is known to the other side of the connection. Use a combination of letters and numbers that is difficult for others to guess.
E1 Options		
MTU	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 drop-down 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
Framing Mode	Specifies the framing mode for the E1 line.	<p>From the drop-down list, select one of the following:</p> <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.	<p>Type numeric values from 2 through 32. Separate discontinuous entries with commas, and use hyphens to indicate ranges. For example:</p> <p>2,4,7–9</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.	Select 16 or 32 . The default checksum is 16 .

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 11, select the interface you want to configure.

The properties you can configure on a Fast Ethernet interface are displayed, as shown in Figure 13.

Figure 13: Fast Ethernet Interfaces Quick Configuration Page

The screenshot shows the Juniper J4300 router configuration interface. The top navigation bar includes the Juniper logo, the router model "ROUTER - J4300", and the user "regress". Below the navigation bar are tabs for Monitor, Configuration, Diagnose, and Manage. The left sidebar lists various configuration options under "Quick Configuration", with "Interfaces" selected. The main content area shows the "Quick Configuration" page for "Interfaces", with a breadcrumb trail "Configuration > Quick Configuration > Interfaces". The page displays a table of "Logical Interfaces" with columns for Logical Interface Name, Link State, Configured, and Description. A single entry is shown for "fe-0/0/0.0" with a "Down" link state and "Yes" configured. Below the table are "Add..." and "Delete" buttons. A "Physical Interface Description" field is also present. At the bottom are "OK", "Cancel", and "Apply" buttons. The footer contains copyright information for 2004-2005 Juniper Networks, Inc.

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: '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 20.
3. Click one of the following buttons:
 - To apply the configuration and stay in 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 70.

Table 20: 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.
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.

Configuring a T1 Interface with Quick Configuration

To configure properties on a T1 interface:

1. From the Quick Configuration page, as shown in Figure 11, select the interface you want to configure.

The properties you can configure on a T1 interface are displayed, as shown in Figure 14.

Figure 14: T1 Interfaces Quick Configuration Page

The screenshot shows the Juniper J6300 Router configuration interface. At the top, the Juniper logo is on the left, "ROUTER - J6300" is in the center, and "Logged in as: regress" is on the right. Below the logo is a navigation bar with "Monitor", "Configuration", "Diagnose", and "Manage". The "Configuration" tab is active. On the left is a sidebar menu with "Quick Configuration" (expanded), "Set Up", "SSL", "Interfaces" (selected), "Users", "SNMP", "Routing", "Firewall/NAT", "IPSec Tunnels", "Realtime Performance Monitoring", "View and Edit", "History", and "Rescue". The main content area shows the "Quick Configuration" page for "Interfaces". The breadcrumb trail is "Configuration > Quick Configuration > Interfaces". The page title is "Quick Configuration". Below it, "Interfaces" is followed by "Physical Interface: 't1-6/0/1'". The "Logical Interfaces" section states "No logical interfaces configured." with an "Add..." button. The "Physical Interface Description" section has a text input field. The "Encapsulation" section has a dropdown menu for "Encapsulation" and a checkbox for "Enable CHAP". The "CHAP Local Identity" section has a checked checkbox for "Use System Host Name" and input fields for "Local Name", "* CHAP Peer Identity", and "* CHAP Secret".

2. Enter information into the Quick Configuration page, as described in Table 21.
3. Click one of the following buttons:
 - To apply the configuration and stay in the Quick Configuration page, click **Apply**.
 - To apply the configuration and return to the main configuration page, click **OK**.
 - To cancel your entries and return to the main page, click **Cancel**.
4. To verify that the T1 interface is configured correctly, see “Verifying Interface Configuration” on page 70.

Table 21: 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.
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.
Encapsulation		
Encapsulation	Specifies the encapsulation type for traffic on the interface.	<p>From the drop-down list, select the encapsulation for this T1 interface:</p> <ul style="list-style-type: none"> ■ PPP ■ Frame Relay ■ Cisco HDLC
Enable CHAP	Enables or disables CHAP authentication on a T1 interface with PPP encapsulation only.	<ul style="list-style-type: none"> ■ To enable CHAP, select the check box. ■ To disable CHAP, clear the check box.
CHAP Local Identity (available if CHAP is enabled)		
Use System Host Name	Specifies that the T1 interface uses the Services Router's system hostname in CHAP challenge and response packets.	<ul style="list-style-type: none"> ■ To enable, select the check box (the default). ■ To disable, clear the check box.
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.

Table 21: T1 Quick Configuration Summary (Continued)

Field	Function	Your Action
T1 Options		
MTU	Specifies the maximum transmission unit size for the T1 interface.	Type a value between 256 and 9192 bytes. The default MTU for T1 interfaces is 1504 .
Clocking	Specifies the transmit clock source for the T1 line.	From the drop-down list, select one of the following: <ul style="list-style-type: none"> ■ internal—Services Router's own system clock (the default) ■ external—Clock received from the T1 interface
Framing Mode	Specifies the framing mode for the T1 line.	From the drop-down 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 drop-down 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 drop-down 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 21: 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	Specifies the T1 cable length, in feet.	<p>From the drop-down list, select one of the following cable lengths:</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)

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 11, select the interface you want to configure.

The properties you can configure on a T3 interface are displayed, as shown in Figure 15.

Figure 15: 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 22.
3. Click one of the following buttons:
 - To apply the configuration and stay in 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 70.

Table 22: 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.
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.
Encapsulation		
Encapsulation	Specifies the encapsulation type for traffic on the interface.	<p>From the drop-down list, select the encapsulation for this T3 interface:</p> <ul style="list-style-type: none"> ■ PPP ■ Frame Relay ■ Cisco HDLC
Enable CHAP	Enables or disables CHAP authentication on a T3 interface with PPP encapsulation only.	<ul style="list-style-type: none"> ■ To enable CHAP, select the check box. ■ To disable CHAP, clear the check box.
CHAP Local Identity (available if CHAP is enabled)		
Use System Host Name	Specifies that the T3 interface uses the Services Router's system hostname in CHAP challenge and response packets.	<ul style="list-style-type: none"> ■ To enable, select the check box (the default). ■ To disable, clear the check box.
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.

Table 22: T3 Quick Configuration Summary (Continued)

Field	Function	Your Action
T3 Options		
MTU	Specifies the maximum transmission unit size for the T3 interface.	Type a value between 256 and 9192 bytes. The default MTU for T3 interfaces is 4474 .
Clocking	Specifies the transmit clock source for the T3 line.	From the drop-down list, select one of the following: <ul style="list-style-type: none"> ■ internal—Services Router's own system clock (the default) ■ external—Clock received from the T3 interface
C-Bit Parity	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 enable, select the check box. ■ To disable, clear the check box.
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 .
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.

Configuring a Serial Interface with Quick Configuration

To configure properties on a serial interface:

1. From the Quick Configuration page, as shown in Figure 11, select the interface you want to configure.

The properties you can configure on a serial interface are displayed, as shown in Figure 16.

Figure 16: Serial 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: 'se-5/0/0'**

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 23.
3. Click one of the following buttons:
 - To apply the configuration and stay in 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 70.

Table 23: Serial Quick Configuration Summary

Field	Function	Your Action
Logical Interfaces		
Add logical interfaces	Defines one or more logical units that you connect to this physical serial interface. You must define at least one logical unit for a serial interface. You can define multiple units if the encapsulation type is Frame Relay.	Click Add .
Logical Interface Description	(Optional) Describes the logical interface.	Type a text description of the logical interface to more clearly identify it in monitoring displays.
IPv4 Addresses and Prefixes	Specifies one or more IPv4 addresses for the interface.	<ol style="list-style-type: none"> Type one or more IPv4 addresses and prefixes in dotted decimal notation. For example: 10.10.10.10/24 Click Add.
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.
Encapsulation		
Encapsulation	Specifies the encapsulation type for traffic on the interface.	<p>From the drop-down list, select the encapsulation for this serial interface:</p> <ul style="list-style-type: none"> ■ PPP ■ Frame Relay ■ Cisco HDLC
Enable CHAP	Enables or disables CHAP authentication on a serial interface with PPP encapsulation only.	<ul style="list-style-type: none"> ■ To enable CHAP, select the check box. ■ To disable CHAP, clear the check box.
CHAP Local Identity (available if CHAP is enabled)		
Use System Host Name	Specifies that the serial interface use the Services Router's system hostname in CHAP challenge and response packets.	<ul style="list-style-type: none"> ■ To enable, select the check box (the default). ■ To disable, clear the check box.
Local Name	If Use System Host Name is disabled, specifies the local name for CHAP to use.	Type a local name for this serial interface.
CHAP Peer Identity	Identifies the client or peer with which the Services Router communicates on this serial interface.	Type the CHAP client name.
CHAP Secret	Specifies the secret password for CHAP authentication, known to both sides of the connection.	Type a password that is known to the other side of the connection. Use a combination of letters and numbers that is difficult for others to guess.
Serial Options		
MTU	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.

Table 23: Serial Quick Configuration Summary (Continued)

Field	Function	Your Action
Clocking Mode	<p>Specifies the clock source to determine the timing on serial interfaces.</p> <p>If you use an externally timed clocking mode—dce or loop—long cables might introduce a phase shift of DTE-transmitted clock and data. At high speeds, this phase shift might cause errors.</p> <p>Inverting the transmit clock corrects the phase shift, thereby reducing error rates. By default, the transmit clock is not inverted. To invert the transmit clock, do either of the following:</p> <ul style="list-style-type: none"> ■ In the J-Web configuration editor, set the Transmit clock value to invert on the Interfaces > Interface-name > Serial options page. ■ In the CLI configuration editor, include the transmit-clock invert statement at the [edit interfaces se-fpc / O / port serial-options] hierarchy level. 	<p>From the drop-down list, select one of the following timing sources:</p> <ul style="list-style-type: none"> ■ dce—Uses a transmit clock generated by the data circuit-terminating equipment (DCE) for the Services Router's DTE. ■ internal—Uses the Services Router's internal clock. ■ loop—Uses the DCE's or DTE's receive clock (the default). <p>For X.21 serial interfaces, you must use the loop clocking mode.</p> <p>When the Services Router is functioning as DTE, you must use the dce clocking mode for all interfaces except X.21 serial interfaces.</p> <p>When the Services Router is functioning as DCE, we recommend using the internal clocking mode for all interfaces.</p>
Clock Rate	<p>Specifies the line speed in kilohertz or megahertz for serial interfaces that use the DTE clocking mode.</p>	<p>From the drop-down 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 Configuration Page, as described in “Configuring Network Interfaces with Quick Configuration” on page 47. You can perform the same configuration tasks using the J-Web or CLI configuration editors. 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 64
- Adding an ATM-for-ADSL Network Interface with a Configuration Editor on page 66
- Deleting a Network Interface with a Configuration Editor on page 69

For information about using the J-Web and CLI configuration editors, see “Using J-series 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 24.
3. When you are finished configuring the interface, click **Commit**.
4. To verify that the network interface is configured correctly, see “Verifying Interface Configuration” on page 70.

Table 24: 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. 	From the top of the configuration hierarchy, enter edit interfaces

Table 24: Adding an Interface (Continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
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 “Interfaces and Interface Naming” on page 45. 3. Click OK. 	<p>Create and name the interface:</p> <pre>set interface-name</pre>
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.</p>

Table 24: 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 <code>edit interface-options</code> 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 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 <code>set unit logical-unit-number</code> 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-for-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-for-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 DSL only, not for direct ATM connections.

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

To configure ATM-for-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 25.
3. When you are finished configuring the interface, click **Commit**.
4. Go on to one of the following procedures:
 - To verify that the ATM-for-ADSL network interface is configured correctly, see “Verifying Interface Configuration” on page 70.
 - To configure Point-to-Point Protocol over Ethernet (PPPoE) for ADSL encapsulation on the interface, see “Configuring PPPoE Encapsulation on an ATM-for-ADSL Interface” on page 83.

Table 25: Adding an ATM-for-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. 	From the top of the configuration hierarchy, create and name the interface: edit interfaces at-2/0/0
Create the new interface—for example, at-2/0/0.	<ol style="list-style-type: none"> 1. Next to Interface, click Add new entry. 2. In the Interface name box, type the name of the new interface. 3. Click OK. 	
Configuring Physical Properties		
Configure ATM options for the interface.	<ol style="list-style-type: none"> 1. Next to Atm options, click Configure. 2. Next to Vpi, click Add new entry. 3. In the Vpi number box, enter a value between 0 and 255 in the box—for example, 25. 4. Click OK. 	Enter set atm-options vpi 25

Table 25: Adding an ATM-for-ADSL Network Interface (Continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Configure the ADSL operating mode.	1. Next to Dsl options, click Configure .	Enter
Annex A and Annex B support the following operating modes:	2. From the Operating Mode drop-down list, select the type of DSL operating mode—for example, 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. 		
Annex A supports an additional operating mode:		
<ul style="list-style-type: none"> ■ 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.	1. From the Encapsulation drop-down list, select the encapsulation type—for example, ether-over-atm .	Enter
<ul style="list-style-type: none"> ■ atm-pvc—ATM permanent virtual circuits. For PPP over ATM-for-ADSL (PPPoA) interfaces, use this type of encapsulation. ■ ether-over-atm—Ethernet over ATM encapsulation. For PPP over Ethernet (PPPoE) over ATM for ADSL interfaces that carry IPv4 traffic, use this type of encapsulation. 	2. Click OK to save the configuration so far.	set encapsulation ether-over-atm
Configuring Logical Properties		
Add the logical interface.	1. Scroll down the page to Unit, and click Add new entry .	Enter
	2. In the Interface unit number box, type a value from 0 to 16385—for example, 3.	set unit 3
	3. Enter values in the fields required by your network.	

Table 25: Adding an ATM-for-ADSL Network Interface (Continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Configure encapsulation for the logical unit.</p> <p>The ATM-for-ADSL interface supports the following encapsulations on the logical interface:</p> <ul style="list-style-type: none"> ■ atm-vc-mux—Use ATM VC multiplex encapsulation. You can only configure the inet family when you use this type of encapsulation. ■ atm-nlpid—Use ATM network layer protocol identifier (NLPID) encapsulation. You can only configure the inet family when you use this type of encapsulation. ■ atm-cisco-nlpid—Use Cisco NLPID encapsulation. You can only configure the inet family when you use this type of encapsulation. ■ atm-snap—Use ATM SNAP encapsulation. ■ atm-ppp-vc-mux—Use PPP over ATM AAL5 multiplex encapsulation. (For PPPoA interfaces, use this type of encapsulation.) ■ atm-ppp-llc—Use PPP over ATM AAL5 LLC encapsulation. (For PPPoA interfaces, use this type of 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. ■ 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. 	<p>From the Encapsulation drop-down list, select the type of encapsulation—for example, atm-nlpid.</p>	<p>Enter</p> <p>set unit 3 encapsulation atm-nlpid</p>
Add the Family protocol type.	Select the protocol type—for example, inet —and then click Configure .	<p>Enter set unit 3 family inet</p> <p>Commands vary depending on the protocol type.</p>
Configure the virtual channel identifier (VCI) type and value.	<ol style="list-style-type: none"> 1. From the Vci type menu, select a value. 2. Enter the ATM VCI value in the Vci field. 3. Click OK. 	Enter set unit 3 vci 10

Deleting a Network Interface with a Configuration Editor

To delete an interface on a Services Router:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 26.



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 26: Deleting an Interface

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

Verifying Interface Configuration

To verify an interface configuration, perform these tasks:

- Verifying the Link State of All Interfaces on page 70
- Verifying Interface Properties on page 71
- Verifying ADSL Interface Properties on page 72

Verifying the Link State of All Interfaces

Purpose By using the ping tool on each peer address in the network, verify that all interfaces on the Services Router are operational.

- Action** For each interface on the Services Router:
1. In the J-Web interface, select **Diagnose > Ping Host**.
 2. In the Remote Host box, type the address of the interface for which you want to verify the link state.
 3. Click **Start**. Output appears on a separate page.

Sample Output

```
PING 10.10.10.10 : 56 data bytes
64 bytes from 10.10.10.10: icmp_seq=0 ttl=255 time=0.382 ms
64 bytes from 10.10.10.10: icmp_seq=1 ttl=255 time=0.266 ms
```

What It Means If the interface is operational, it generates an ICMP response. If this response is received, the round-trip time in milliseconds is listed in the *time* field. For more information about the output, see the *J-series Services Router Administration Guide*.

For more information about using the J-Web interface to ping a host, see the *J-series Services Router Administration Guide*.

For information about the `ping` command, see the *J-series Services Router Administration Guide* or the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying Interface Properties

Purpose Verify that the interface properties are correct.

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

Sample Output

```
user@host> show interfaces detail

Physical interface: fe-1/0/0, Enabled, Physical link is Up
  Interface index: 134, SNMP ifIndex: 27, Generation: 17
  Link-level type: Ethernet, MTU: 1514, Speed: 100mbps, Loopback: Disabled,
  Source filtering: Disabled, Flow control: Enabled
  Device flags   : Present Running
  Interface flags: SNMP-Traps 16384
  Link flags     : None
  CoS queues     : 4 supported
  Hold-times     : Up 0 ms, Down 0 ms
  Current address: 00:90:69:87:44:9d, Hardware address: 00:90:69:87:44:9d
  Last flapped   : 2004-08-25 15:42:30 PDT (4w5d 22:49 ago)
  Statistics last cleared: Never
  Traffic statistics:
    Input  bytes :                  0          0 bps
    Output bytes :                  0          0 bps
    Input  packets:                  0          0 pps
    Output packets:                  0          0 pps
  Queue counters:      Queued packets  Transmitted packets  Dropped packets
    0 best-effort             0                0                0
```

```

1 expedited-fo          0          0          0
2 assured-forw          0          0          0
3 network-cont          0          0          0
Active alarms   : None
Active defects  : None

```

What It Means The output shows a summary of interface information. Verify the following information:

- The physical interface is **Enabled**. If the interface is shown as **Disabled**, do either of the following:
 - In the CLI configuration editor, delete the **disable** statement at the [edit interfaces *interface-name*] level of the configuration hierarchy.
 - In the J-Web configuration editor, clear the **Disable** check box on the **Interfaces > interface-name** page.
- The physical link is **Up**. A link state of **Down** indicates a problem with the interface module, interface port, or physical connection (link-layer errors).
- The **Last Flapped** time is an expected value. The **Last Flapped** time indicates the last time the physical interface became unavailable and then available again. Unexpected flapping indicates likely link-layer errors.
- The traffic statistics reflect expected input and output rates. Verify that the number of inbound and outbound bytes and packets matches expected throughput for the physical interface. To clear the statistics and see only new changes, use the **clear interfaces statistics interface-name** command.

For more information about **show interfaces detail**, see the *JUNOS Network and Services 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-4/0/0 extensive

Physical interface: at-4/0/0, Enabled, Physical link is Up
  Interface index: 141, SNMP ifIndex: 38, Generation: 102
  Link-level type: ATM-PVC, MTU: 4482, Clocking: Internal, ADSL mode, Speed: ADSL
  Device flags   : Present Running
  Link flags     : None
  CoS queues     : 8 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   : 2005-01-25 15:42:30 PDT (4w5d 22:49 ago)
  Statistics last cleared: Never
  Traffic statistics:

```

```

Input bytes : 0
Output bytes : 0
Input packets: 0
Output packets: 0
Input errors:
Errors: 0, Drops: 0, Invalid VCs: 0, Framing Errors: 0, Policed Discards: 0,
L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts:0, Resource errors:0
Carrier transitions: 9, Errors: 0, Drops:0, Aged packets: 0, MTU errors:0,
ADSL alarms: none
ADSL defects: none
ADSL media:
      Seconds Count State
LOF      0      0 OK
LOS      0      0 OK
LOM      0      0 OK
LOP      0      0 OK
LOC DI   0      0 OK
LOC DNI  0      0 OK
ADSL Status:
Modem status:Showtime
DSL Mode: Auto
ADSL Statistics:
      ATU-R ATU-C
Attainable Bit Rate (kbps): 0 0
Attenuation (dB): 8 0
Capacity used (%): 0 0
Noise Margin (dB): 24 33
Output Power )dBm): 24 24
      Interleave Fast Interleave Fast
Bitrate (kbps): 0 8128 0 832
CRC: 0 1 0 0
FEC: 0 0 0 0
HEC: 0 0 0 0
Received Cells: 0 164548
Transmitted Cells: 0 263100
ATM status:
HCS state: Hunt
LOC: OK
ATM Statistics:
Uncorrectable HCS errors:0, Correctable HCS errors:0,
Tx cell FIFO overruns:0,Rx cell overruns:0,
Rx cell FIFO underruns:0, Input cell count:0,
Output cell count:202811720590, Output idle cell count:202811720562
Output VC queue drops:0, Input no buffers:0, Input length errors:0,
Input timeouts:0, Input valid VCs: 0, Input bad CRCs:0,
Input OAM cell no buffers:0
Packet Forwarding Engine configuration:
Destination Slot:0
CoS transmit queue Bandwidth Buffer Priority Limit %bps %bytes
0 best-effort 95 14774400 95 0 low none
3 network-control 5 7776000 5 0 low 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.
 - 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 pointer
 - LOS—Loss of signal

Examine the operational statistics for an ADSL interface. Statistics in the **ATU-R** (ADSL transceiver unit–remote) column are for the far end. Statistics in the **ATU-C** (ADSL transceiver unit–central office) column are for the near 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 Network and Services Interfaces Command Reference*.

Chapter 3

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 77
- PPPoE Overview on page 78
- Before You Begin on page 82
- Configuring PPPoE with a Configuration Editor on page 82
- Verifying a PPPoE Configuration on page 88

PPPoE Terms

Before configuring PPPoE on a Services Router, become familiar with the terms defined in Table 27.

Table 27: PPPoE Terms

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

PPPoE Overview

On the Services Router, PPPoE establishes a point-to-point connection between the client (Services Router) and the server, also called an access concentrator. Multiple hosts can be connected to the Services Router, and their data can be authenticated, encrypted, and compressed before the traffic is sent to the PPPoE session on the Services Router's Fast Ethernet 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 79

- PPPoE Stages on page 80
- Optional CHAP Authentication on page 81

PPPoE Interfaces

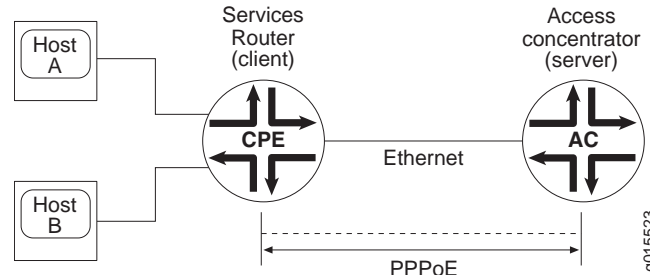
The PPPoE interface to the access concentrator can be either a Fast Ethernet interface on any Services Router or an ATM-for-ADSL interface on a J4300 or J6300 Services Router. The PPPoE configuration is the same for both interfaces. The only difference is the encapsulation for the underlying interface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

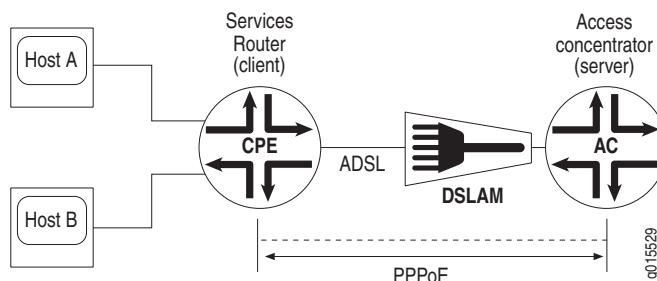
The Services Router encapsulates each PPP frame in an Ethernet frame and transports the frames over an Ethernet loop. Figure 17 shows a typical PPPoE session between a Services Router and an access concentrator on the Ethernet loop.

Figure 17: PPPoE Session on the Ethernet Loop



ATM-for-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 18 shows a typical PPPoE over ATM session between a Services Router and an access concentrator on an ADSL loop.

Figure 18: PPPoE Session on an ADSL Loop

PPPoE Stages

PPPoE has two stages, the discovery stage and the PPPoE session stage. In the discovery stage, the client discovers the access concentrator by identifying the Ethernet media access control (MAC) address of the access concentrator and establishing a PPPoE session ID. In the PPPoE session stage, the client and the access concentrator build a point-to-point connection over Ethernet, based on the information collected in the discovery stage.

PPPoE Discovery Stage

A Services Router initiates the PPPoE discovery stage by broadcasting a PPPoE Active Discovery Initiation (PADI) packet. To provide a point-to-point connection over Ethernet, each PPPoE session must learn the Ethernet MAC address of the access concentrator and establish a session with a unique session ID. Because the network might have more than one access concentrator, the discovery stage allows the client to communicate with all of them and select one.



NOTE: A Services Router cannot receive PPPoE packets from two different access concentrators on the same physical interface.

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 Session Stage

The PPPoE session stage starts after the PPPoE discovery stage is over. The access concentrator can start the PPPoE session after it sends the PADS packet to the client, or the client can start the PPPoE session after it receives a PADS packet from the access concentrator. A Services Router supports multiple PPPoE sessions on each interface, but no more than 256 PPPoE sessions per Services Router.

Each PPPoE session is uniquely identified by the Ethernet address of the peer and the session ID. After the PPPoE session is established, data is sent as in any other PPP encapsulation. The PPPoE information is encapsulated within an Ethernet frame and is sent to a unicast address. In this stage, both the client and the server must allocate resources for the PPPoE logical interface.

After a session is established, the client or the access concentrator can send a PPPoE Active Discovery Termination (PADT) packet anytime to terminate the session. The PADT packet contains the destination address of the peer and the session ID of the session to be terminated. After this packet is sent, the session is closed to PPPoE traffic during that session.

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.

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 41.

For a PPPoE over ATM interface, see “Adding an ATM-for-ADSL Network Interface with a Configuration Editor” on page 66.

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 82
- Configuring a PPPoE Interface (Required) on page 85
- Configuring CHAP (Optional) on page 87

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-for-ADSL interface, you must configure encapsulation on both the physical and logical interfaces. To configure encapsulation on an ATM-for-ADSL physical interface, use Ethernet over ATM encapsulation. To configure encapsulation on an ATM-for-ADSL logical interface, use the PPPoE over AAL5 Link Layer 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 83
- Configuring PPPoE Encapsulation on an ATM-for-ADSL Interface on page 83

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 28.
3. If you are finished configuring the network, 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 85.
 - To enable authentication on the interface, see “Configuring CHAP (Optional)” on page 87.
 - To check the configuration, see “Verifying a PPPoE Configuration” on page 88.

Table 28: Configuring PPPoE Encapsulation on an Ethernet Interface

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

Configuring PPPoE Encapsulation on an ATM-for-ADSL Interface

To configure PPPoE encapsulation on an ATM-for-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 29.

3. If you are finished configuring the network, 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 85.
 - To enable authentication on the interface, see “Configuring CHAP (Optional)” on page 87.
 - To check the configuration, see “Verifying a PPPoE Configuration” on page 88.

Table 29: Configuring PPPoE Encapsulation on an ATM-for-ADSL Interface

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

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 30.
3. If you are finished configuring the network, 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 87.
 - To check the configuration, see “Verifying a PPPoE Configuration” on page 88.

Table 30: Configuring a PPPoE Interface

Task	J-Web Configuration Editor	CLI Configuration Editor
Navigate to the Interfaces level in the configuration hierarchy.	In the configuration editor hierarchy select Interfaces .	From the top of the configuration hierarchy, enter edit interfaces
Create a PPPoE interface with a logical interface unit 0.	<ol style="list-style-type: none"> 1. Next to Interface, click Add new entry. 2. In the Interface name box, type pp0 and click OK. 3. Under Interface name, click pp0. 4. Next to Unit, click Add new entry. 5. In the Interface unit number box, type 0. 	Enter edit pp0 unit 0
Specify the logical Ethernet interface or the logical ATM interface as the underlying interface for the PPPoE session—for example, fe-0/0/1.0 or at-2/0/0.0 .	<ol style="list-style-type: none"> 1. Next to Pppoe options, click Edit. 2. In the Underlying Interface box, type the following: For the logical Ethernet interface, type fe-0/0/1.0. For the logical ATM interface type, at-2/0/0.0. 	Enter one of the following: ■ For the logical Ethernet interface, type set pppoe-options underlying-interface fe-0/0/1.0 . ■ For the logical ATM interface, type set pppoe-options underlying-interface at-2/0/0.0 .
Identify the access concentrator by a unique name—for example, ispl.com .	In the Access concentrator box type ispl.com .	Enter set pppoe-options access-concentrator ispl.com

Table 30: Configuring a PPPoE Interface (Continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
Specify the time in seconds to reconnect after a PPPoE session is terminated—for example, 100 seconds .	In the Auto reconnect box, type 100 .	Enter 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, video@ispl.com .	<ol style="list-style-type: none"> 1. In the Service name box, type video@ispl.com. 2. Click OK to apply your entries to the configuration. 	Enter set pppoe-options service-name video@ispl.com
Configure the maximum transmission unit (MTU) of the protocol—for example, 1492 .	<ol style="list-style-type: none"> 1. In the Inet box, select Yes and click Configure. 2. In the Mtu box, type 1492. 	Enter up set pp0 mtu 1492

Table 30: Configuring a PPPoE Interface (Continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
<p>Configure the PPPoE interface address in one of the following ways:</p> <ul style="list-style-type: none"> Assign source and destination addresses—for example, 192.168.1.1/32 and 192.168.1.2. 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. Obtain an IP address by negotiation with the remote end. This method might require the access concentrator to use a RADIUS authentication server. 	<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"> Next to Address, click Add new entry. In the Source box, type 192.168.1.1/32. In the Destination box, type 192.168.1.2. Click OK until you return to the Unit page. <p>To derive the source address and assign the destination address:</p> <ol style="list-style-type: none"> Next to Unnumbered address, select the Yes check box and click Configure. In the Destination box, type 192.168.1.2. In the Source box, type lo0.0. Click OK until you return to the Unit page. <p>To obtain an IP address from the remote end:</p> <ol style="list-style-type: none"> Next to Negotiate address, select the Yes check box. Click OK until you return to the Unit page. 	<p>Enter up, then do one of the following:</p> <ul style="list-style-type: none"> To assign the source and destination address, type set pp0.0 family inet address 192.168.1.1/32 destination 192.168.1.2. To derive the source address and assign the destination address, type set pp0.0 family inet unnumbered-address lo0.0 destination 192.168.1.2. To obtain an IP address from the remote end, type set pp0.0 family inet negotiate-address.
<p>Disable the sending of keepalives on a logical interface—for example, no-keepalives.</p>	<ol style="list-style-type: none"> From the Keepalive choices drop-down list, select no keepalives. Click OK to apply your entries to the configuration. 	<p>Enter</p> <p>set pp0.0 no-keepalives</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.

2. Perform the configuration tasks described in Table 31.
3. If you are finished configuring the network, commit the configuration.
4. To check the configuration, see “Verifying a PPPoE Configuration” on page 88.

Table 31: Configuring CHAP

Task	J-Web Configuration Editor	CLI Configuration Editor
Define a CHAP access profile—for example, A-ppp-client.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Access. 2. Next to Profile, click Add new entry. 3. In the Profile name box, type A-ppp-client. 4. Next to Client, type client1. 5. In the Chap secret box, type my-secret. 6. Click OK three times. 	<p>Enter</p> <pre>set access profile A-ppp-client client client1 chap-secret my-secret</pre>
Navigate to the pp0 unit 0 interface level in the configuration hierarchy.	<ol style="list-style-type: none"> 1. In the configuration editor hierarchy, select Interfaces. 2. In the Interface name box, click pp0. 3. In the Interface unit number box, click 0. 	<p>From the top of the configuration hierarchy, enter</p> <pre>edit interfaces pp0 unit 0</pre>
Configure CHAP on the PPPoE interface and specify a unique profile name containing a client list and access parameters—for example, A-ppp-client.	<ol style="list-style-type: none"> 1. Next to Ppp options, click Configure. 2. Next to Chap, click Configure. 3. In the Access-profile box, type A-ppp-client. 	<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-fe-0/0/1.0 or A-at-2/0/0.0.	<p>In the Local name box, type one of the following:</p> <ul style="list-style-type: none"> ■ For the Ethernet interface, type A-fe-0/0/1.0. ■ For the ATM interface, type A-at-2/0/0.0. 	<p>Enter one of the following:</p> <ul style="list-style-type: none"> ■ For the Ethernet interface, type set ppp-options chap local-name A-fe-0/0/1.0. ■ For the ATM interface, type set ppp-options chap local-name A-at-2/0/0.0.
Set the passive option to handle incoming CHAP packets only.	<ol style="list-style-type: none"> 1. In the Passive box, click Yes. 2. Click OK to apply your entries to the configuration. 	<p>Enter</p> <pre>set ppp-options chap passive</pre>

Verifying a PPPoE Configuration

To verify PPPoE configuration perform the following tasks:

- Displaying a PPPoE Configuration for an ATM-for-ADSL Interface on page 89
- Verifying PPPoE Interfaces on page 90
- Verifying PPPoE Sessions on page 91
- Verifying the PPPoE Version on page 92
- Verifying PPPoE Statistics on page 92

Displaying a PPPoE Configuration for an ATM-for-ADSL Interface

Purpose Verify the PPPoE configuration for an ATM-for-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-for-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 Network and Services Interfaces Command Reference*.

Verifying PPPoE Sessions

Purpose Verify that a PPPoE session is running properly on the logical interface.

Action From the CLI, enter the `show pppoe interfaces` command.

Sample Output

```
user@host> show pppoe interfaces

pp0.0 Index 67
  State: Session up, Session ID: 31,
  Service name: video@ispl.com, Configured AC name: ispl.com,
  Session AC name: belur, AC MAC address: 00:90:1a:40:f6:4e,
  Auto-reconnect timeout: 1 seconds,
  Underlying interface: fe-0/0/1.0 Index 69
```

- What It Means** The output shows information about the PPPoE sessions. Verify the following information:
- The PPPoE session is running on the correct logical interface.
 - Under **State**, the session is active (**up**).
 - Under **Underlying interface**, the physical interface on which the PPPoE session is running is correct:
 - For an Ethernet connection, the underlying interface is Fast Ethernet—for example, **fe-0/0/1.0**.
 - For an ATM-for-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 Network and Services 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 Network and Services 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 Network and Services Interfaces Command Reference*.

Part 3

Configuring Routing Protocols

- Routing Overview on page 97
- Configuring Static Routes on page 127
- Configuring a RIP Network on page 139
- Configuring an OSPF Network on page 155
- Configuring BGP Sessions on page 177

Chapter 4

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 97
- Routing Overview on page 101
- RIP Overview on page 107
- OSPF Overview on page 111
- BGP Overview on page 116

Routing Terms

To understand routing, become familiar with the terms defined in Table 32 .

Table 32: 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 32: 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 32: 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 32: 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 102
- Autonomous Systems on page 102
- Interior and Exterior Gateway Protocols on page 102
- Routing Tables on page 103
- Forwarding Tables on page 103

- Dynamic and Static Routing on page 104
- Route Advertisements on page 105
- Route Aggregation on page 105

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 19 shows a simple network of routers.

Figure 19: 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 19 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 19, 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 20 shows a network that uses static routes.

Figure 20: Static Routing Example




In Figure 20, 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 21.

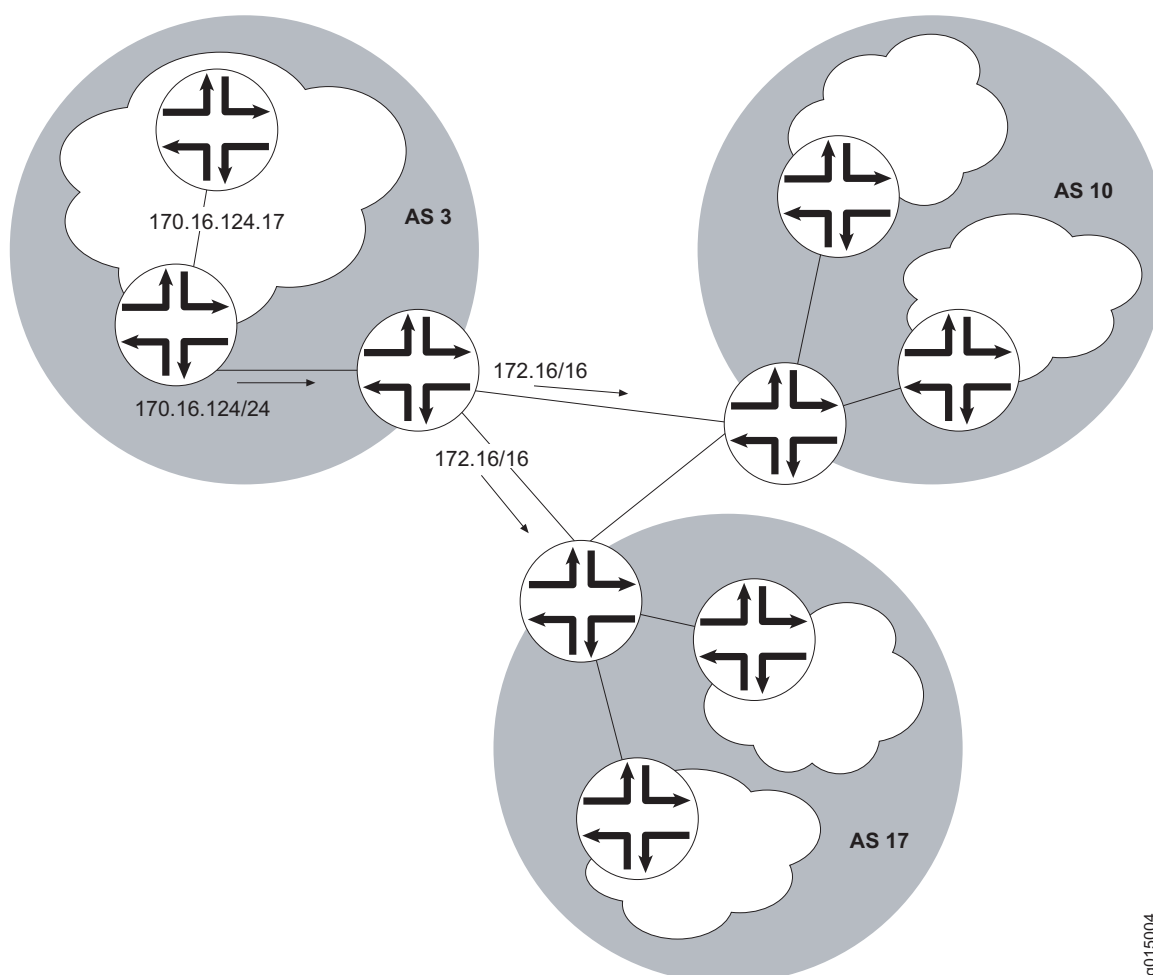
Figure 21: Route Advertisement



In Figure 21, 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 22.

Figure 22: Route Aggregation

g015004

Figure 22 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:

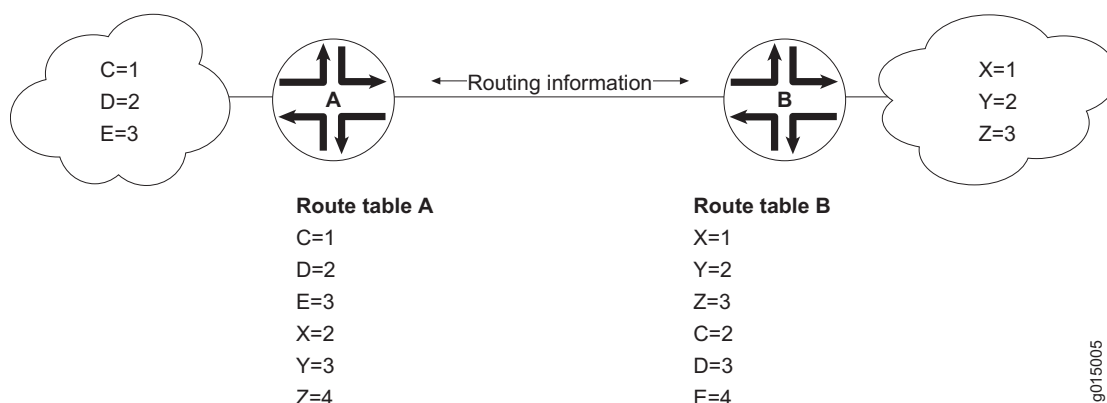
- Distance-Vector Routing Protocols on page 107
- Maximizing Hop Count on page 108
- RIP Packets on page 109
- Split Horizon and Poison Reverse Efficiency Techniques on page 109
- Limitations of Unidirectional Connectivity on page 110



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 23 shows how distance-vector routing works.

Figure 23: Distance-Vector Protocol

In Figure 23, 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 24 shows an example of the split horizon technique.


Figure 24: Split Horizon Example



In Figure 24, 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 25 shows an example of the poison reverse technique.


Figure 25: Poison Reverse Example



In Figure 25, 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 26 shows, RIP networks are limited by their unidirectional connectivity.

Figure 26: Limitations of Unidirectional Connectivity

In Figure 26, 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 112.

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 112
- Role of the Designated Router on page 112
- Path Cost Metrics on page 113
- Areas and Area Border Routers on page 113
- Role of the Backbone Area on page 114
- Stub Areas and Not-So-Stubby Areas on page 115




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 27.

Figure 27: OSPF Three-Way Handshake



In Figure 27, 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 28 shows an OSPF topology of three areas connected by two area border routers.


Figure 28: 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 28, 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 29 shows a virtual link between a noncontiguous area and the backbone area through an area connected to both.

Figure 29: OSPF Topology with a Virtual Link

In the topology shown in Figure 29, 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 30 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 30: OSPF AS Network with Stub Areas and NSSAs

9015012

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 30 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 30, 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 117
- BGP Messages for Session Establishment on page 118
- BGP Messages for Session Maintenance on page 118
- IBGP and EBGP on page 118
- Route Selection on page 119
- Local Preference on page 120
- AS Path on page 121
- Origin on page 121
- Multiple Exit Discriminator on page 122
- Scaling BGP for Large Networks on page 122

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 31 shows an example of a BGP peering session.

Figure 31: BGP Peering Session



In Figure 31, 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 122. For information about routing confederations, see “Scaling BGP for Large Networks” on page 122.

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 120.)
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 121.)
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 121.)
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 122.)


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 32 illustrates how to use local preference to determine BGP route selection.

Figure 32: Local Preference



The network in Figure 32 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 33 shows how BGP creates an AS path.

Figure 33: BGP AS Path



In the network shown in Figure 33, 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 34 illustrates how to use an MED metric to determine route selection.

Figure 34: MED Example




Figure 34 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 123

- Confederations—for Subdivision on page 125

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 35.



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 35: Simple Route Reflector Topology (One Cluster)




Figure 35 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 36).


Figure 36: Basic Route Reflection (Multiple Clusters)

Figure 36 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 37).

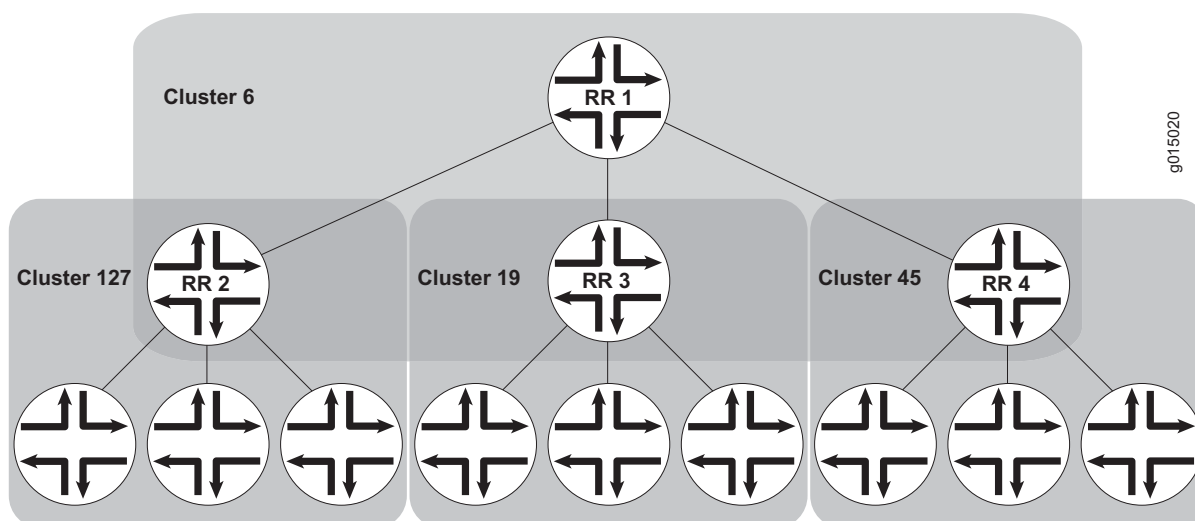
Figure 37: Hierarchical Route Reflection (Clusters of Clusters)

Figure 37 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 38 shows an AS divided into four confederations.


Figure 38: BGP Confederations

Figure 38 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 5

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 127
- Before You Begin on page 130
- Configuring Static Routes with Quick Configuration on page 130
- Configuring Static Routes with a Configuration Editor on page 132
- Verifying the Static Route Configuration on page 137

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 128
- Qualified Next Hops on page 128
- Control of Static Routes on page 128
- Default Properties on page 129

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 129.
- **no-readvertise**—Prevents the route from being readvertised to other routing protocols. For more information, see “Readvertisement Prevention” on page 129.
- **passive**—Rejects traffic destined for the route. For more information, see “Forced Rejection of Passive Route Traffic” on page 129.

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 41.

Configuring Static Routes with Quick Configuration

J-Web Quick Configuration allows you to configure static routes. Figure 39 shows the Quick Configuration Routing page for static routing.

Figure 39: 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"/>	<u>10.74.10.0/24</u>	
<input type="checkbox"/>	<u>172.16.0.0/12</u>	192.168.124.254
<input type="checkbox"/>	<u>192.168.0.0/18</u>	192.168.124.254
<input type="checkbox"/>	<u>192.168.64.0/18</u>	192.168.124.254
<input type="checkbox"/>	<u>207.17.136.192/32</u>	192.168.124.254
<input type="checkbox"/>	<u>192.168.40.0/22</u>	192.168.124.254

To configure static routes with Quick Configuration:

1. In the J-Web user interface, select **Configuration > Routing > Static Routing**.
2. Enter information into the Static Routing Quick Configuration page, as described in Table 33.
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 137.

Table 33: 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 132
- Controlling Static Route Selection (Optional) on page 133
- Controlling Static Routes in the Routing and Forwarding Tables (Optional) on page 135
- Defining Default Behavior for All Static Routes (Optional) on page 136

For information about using the J-Web and CLI configuration editors, see “Using J-series 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 40 shows a sample network.

Figure 40: Customer Routes Connected to a Stub Network



To configure customer routes as static routes, like the ones in Figure 40, 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 34.


3. If you are finished configuring static routes, 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 133.
 - 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 135.
 - To define default properties for static routes, see “Defining Default Behavior for All Static Routes (Optional)” on page 136.
 - To check the configuration, see “Verifying the Static Route Configuration” on page 137.

Table 34: 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 41), you can specify how traffic is to be routed to the destination.

Figure 41: 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 35.
3. If you are finished configuring static routes, 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 135.
 - To define default properties for static routes, see “Defining Default Behavior for All Static Routes (Optional)” on page 136.
 - To check the configuration, see “Verifying the Static Route Configuration” on page 137.

Table 35: 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 36.
3. If you are finished configuring static routes, 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 136.
- To check the configuration, see “Verifying the Static Route Configuration” on page 137.

Table 36: 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.	<ol style="list-style-type: none"> Next to Retain, select the Yes check box. Click OK. 	Set the retain attribute: set retain
Specify that the static route is not to be readvertised. By default, static routes are eligible to be readvertised.	<ol style="list-style-type: none"> Next to Readvertise, select the No check box. Click OK. 	Set the no-readvertise attribute: set no-readvertise
Specify that the static route is to be included in the routing table whether the route is active or not. By default, passive routes are not included in the routing table.	<ol style="list-style-type: none"> From the Passive flag list, select Passive. Click OK. 	Set the passive attribute: set passive

Defining Default Behavior for All Static Routes (Optional)

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

- Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
- Perform the configuration tasks described in Table 37.
- If you are finished configuring static routes, commit the configuration.
- To check the configuration, see “Verifying the Static Route Configuration” on page 137.

Table 37: 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 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

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 6

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 139
- Before You Begin on page 140
- Configuring a RIP Network with Quick Configuration on page 140
- Configuring a RIP Network with a Configuration Editor on page 143
- Verifying the RIP Configuration on page 151

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 41.

Configuring a RIP Network with Quick Configuration

J-Web Quick Configuration allows you to create RIP networks. Figure 42 shows the Quick Configuration Routing page for RIP.

Figure 42: Quick Configuration Routing Page for RIP

Juniper NETWORKS

ROUTER - J4300

Logged in as: **regress**

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

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

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

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 38.
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 151.

Table 38: 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 143
- Controlling Traffic in a RIP Network (Optional) on page 146
- Enabling Authentication for RIP Exchanges (Optional) on page 149

For information about using the J-Web and CLI configuration editors, see “Using J-series 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 43.

Figure 43: 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 43, 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 39.
3. If you are finished configuring the network, 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 146.
 - To authenticate RIP exchanges, see “Enabling Authentication for RIP Exchanges (Optional)” on page 149.
 - To check the configuration, see “Verifying the RIP Configuration” on page 151.

Table 39: 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 39: 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 drop-down 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 drop-down 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 drop-down menu, select rip. 8. Click OK until you return to the Policy statement page. 9. Next to Then, click Configure. 10. From the Accept reject drop-down menu, 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 146
- Controlling Traffic with the Outgoing Metric on page 147

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

Figure 44: 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 40.
3. If you are finished configuring the network, 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 149.
 - To check the configuration, see “Verifying the RIP Configuration” on page 151.

Table 40: 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 45 shows a network with alternate routes between routers A and D.

Figure 45: 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 41.
3. If you are finished configuring the network, 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 149.
 - To check the configuration, see “Verifying the RIP Configuration” on page 151.

Table 41: 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 149
- Enabling Authentication with MD5 Authentication on page 150

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 42.
3. If you are finished configuring the network, commit the configuration.
4. To check the configuration, see “Verifying the RIP Configuration” on page 151.

Table 42: 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 42: 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 43.
3. If you are finished configuring the network, commit the configuration.
4. To check the configuration, see “Verifying the RIP Configuration” on page 151.

Table 43: 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 151
- Verifying the Exchange of RIP Messages on page 151
- Verifying Reachability of All Hosts in the RIP Network on page 152

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	Address	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

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

Chapter 7

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 155
- Before You Begin on page 156
- Configuring an OSPF Network with Quick Configuration on page 156
- Configuring an OSPF Network with a Configuration Editor on page 160
- Tuning an OSPF Network for Efficient Operation on page 167
- Verifying an OSPF Configuration on page 171

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.

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 41.

Configuring an OSPF Network with Quick Configuration

J-Web Quick Configuration allows you to create single-area OSPF networks. Figure 46 shows the Quick Configuration Routing page for OSPF.

Figure 46: 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**

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

Quick Configuration

Routing

Router Identification

Router Identifier ?

OSPF

Enable OSPF ☒

OSPF Area ID

Area Type ?

Enable OSPF on All Interfaces ☒

OSPF Interfaces

fe-0/0/0.0
lo0.0

OSPF-Enabled Interfaces

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 44.
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 171.

Table 44: OSPF Routing Quick Configuration Summary

Field	Function	Your Action
Router Identification		
Router Identifier (required)	Uniquely identifies the router.	Type the Services Router's 32-bit IP address, in dotted decimal notation.
OSPF		
Enable OSPF	Enables or disables OSPF.	<ul style="list-style-type: none"> ■ To enable OSPF, select the check box. ■ To disable OSPF, clear the check box.
OSPF Area ID	Uniquely identifies the area within its AS.	<p>Type a 32-bit numeric identifier for the area, or an integer.</p> <p>If you enter an integer, the value is converted to a 32-bit equivalent. For example, if you enter 3, the value assigned to the area is 0.0.0.3.</p>

Table 44: OSPF Routing Quick Configuration Summary (Continued)

Field	Function	Your Action
Area Type	Designates the type of OSPF area.	<p>From the drop-down 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 160
- Configuring a Single-Area OSPF Network (Required) on page 161
- Configuring a Multiarea OSPF Network (Optional) on page 162
- Configuring Stub and Not-So-Stubby Areas (Optional) on page 165

For information about using the J-Web and CLI configuration editors, see “Using J-series 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 45.
3. Go on to “Configuring a Single-Area OSPF Network (Required)” on page 161.


Table 45: 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 47.

Figure 47: Typical Single-Area OSPF Network Topology



To configure a single-area OSPF network with a backbone area, like the one in Figure 47, 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 46.
3. If you are finished configuring the network, 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 162.
 - To control external route advertisement in the AS, see “Configuring Stub and Not-So-Stubby Areas (Optional)” on page 165.
 - To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 167.
 - To check the configuration, see “Verifying an OSPF Configuration” on page 171.

Table 46: 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 48.

Figure 48: Typical Multiarea OSPF Network Topology

To configure a multiarea OSPF network shown in Figure 48, 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 163
- Creating Additional OSPF Areas on page 163
- Configuring Area Border Routers on page 164

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 161.

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 47.
3. If you are finished configuring the network, commit the configuration.

Table 47: 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 48 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 48.
3. If you are finished configuring the network, 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 165.
 - To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 167.
 - To check the configuration, see “Verifying an OSPF Configuration” on page 171.

Table 48: 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.	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. For example, Services Router B in Figure 48 has the following interfaces enabled for OSPF in the backbone area: <ul style="list-style-type: none"> ■ Interface fe-0/0/0.0 ■ Interface fe-0/0/1.0 To enable an interface on the backbone area, see “Configuring a Single-Area OSPF Network (Required)” on page 161.	View the configuration using the show command: <code>show</code> For example, Services Router B in Figure 48 has the following interfaces enabled for OSPF in the backbone area: <code>area 0.0.0.0 { interface fe-0/0/0.0; interface fe-0/0/1.0; }</code> To enable an interface on the backbone area, see “Configuring a Single-Area OSPF Network (Required)” on page 161.

Table 48: Configuring Area Border Routers (Continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
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: <pre>set area 0.0.0.2 interface fe-0/0/0</pre>
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 49, 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 49: OSPF Network Topology with Stub Areas and NSSAs

To configure stub areas and NSSAs in an OSPF network like the one shown in Figure 49:

1. Create the area and enable OSPF on the interfaces within that area.
For instructions, see “Creating Additional OSPF Areas” on page 163.
2. Configure an area border router to bridge the areas.
For instructions, see “Configuring Area Border Routers” on page 164.
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 49.
5. If you are finished configuring the network, commit the configuration.
6. Go on to one of the following procedures:
 - To improve network operation, see “Tuning an OSPF Network for Efficient Operation” on page 167.
 - To check the configuration, see “Verifying an OSPF Configuration” on page 171.

Table 49: 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.

Table 49: Configuring Stub Area and Not-So-Stubby Area Routers (Continued)

Task	J-Web Configuration Editor	CLI Configuration Editor
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 167
- Controlling the Cost of Individual Network Segments on page 168
- Enabling Authentication for OSPF Exchanges on page 169
- Controlling Designated Router Election on page 170

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 50.

Table 50: 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 51.

Table 51: 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.	<ol style="list-style-type: none"> 1. In the Metric box, type an interface metric value—for example, 5. 2. Click OK. 	<ol style="list-style-type: none"> 1. Set the interface metric. For example: set metric 5 2. Set the external preference. For example: set external-preference 130 Changes in the CLI are applied automatically when you execute the set command.

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 52.

Table 52: 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 edit protocols ospf area 0.0.0.0
Set the authentication type.	<ol style="list-style-type: none"> From the Authentication type list, select the type of authentication to enable on the stub area: simple md5 Click OK. 	<p>Set the authentication type to either simple or md5. For example:</p> <p>set authentication-type md5</p> <p>Changes in the CLI are applied automatically when you execute the set command.</p>
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:

- 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 53.

Table 53: 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.	<ul 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.	Set the priority to a value between 0 and 255. The default value is 128 . For example: set priority 200 Changes in the CLI are applied automatically when you execute the set command.

Verifying an OSPF Configuration

To verify an OSPF configuration, perform these tasks:

- Verifying OSPF-Enabled Interfaces on page 171
- Verifying OSPF Neighbors on page 172
- Verifying the Number of OSPF Routes on page 173
- Verifying Reachability of All Hosts in an OSPF Network on page 174

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	<pre>user@host> show ospf interface</pre> <table><tr><th>Intf</th><th>State</th><th>Area</th><th>DR ID</th><th>BDR ID</th><th>Nbrs</th></tr></table>	Intf	State	Area	DR ID	BDR ID	Nbrs
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 Protocols, Class of Service, and System Basics 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 Protocols, Class of Service, and System Basics 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 50 shows a sample network with an OSPF topology.

Figure 50: 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 50, 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 Protocols, Class of Service, and System Basics 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 173.

For information about the `traceroute` command and its output, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Chapter 8

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 177
- Before You Begin on page 179
- Configuring BGP Sessions with Quick Configuration on page 179
- Configuring BGP Sessions with a Configuration Editor on page 181
- Verifying a BGP Configuration on page 190

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 123

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 125

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 41.

Configuring BGP Sessions with Quick Configuration

J-Web Quick Configuration allows you to create BGP peering sessions. Figure 51 shows the Quick Configuration Routing page for BGP.

Figure 51: 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 54.
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 190.

Table 54: 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 181
- Configuring BGP Within a Network (Required) on page 184
- Configuring a Route Reflector (Optional) on page 185
- Configuring BGP Confederations (Optional) on page 188

For information about using the J-Web and CLI configuration editors, see “Using J-series 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 52 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 52: Typical Network with BGP Peering Sessions

To configure the BGP peering sessions shown in Figure 52:

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 55.
3. If you are finished configuring the network, 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 184.
 - To configure route reflector clusters, see “Configuring a Route Reflector (Optional)” on page 185.
 - To subdivide autonomous systems (ASs), see “Configuring BGP Confederations (Optional)” on page 188.
 - To check the configuration, see “Verifying a BGP Configuration” on page 190.

Table 55: 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 drop-down menu, select external. 2. Click OK. 	From the [edit protocols bgp group external-peers] hierarchy level: set type external

Configuring BGP Within a Network (Required)

To configure BGP sessions between peering networks, you must configure point-to-point sessions between the external peers of the networks. Additionally, you must configure BGP internally to provide a means by which BGP route advertisements can be forwarded throughout the network. Because of the full mesh requirement of IBGP, you must configure individual peering sessions between all internal nodes of the network—unless you use route reflectors or confederations.

Figure 53 shows a typical network with external and internal peer sessions. In the sample network, the Services Router in AS 17 is fully meshed with its internal peers in the group internal-peers, which have IP addresses starting at 192.168.6.4.

Figure 53: Typical Network with EBGP External Sessions and IBGP Internal Sessions



To configure IBGP in the network shown in Figure 53:

1. Configure all external peering sessions as described in “Configuring a Point-to-Point Peering Session (Required)” on page 181.
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 56.
4. If you are finished configuring the network, commit the configuration.
5. Go on to one of the following procedures:
 - To configure route reflector clusters, see “Configuring a Route Reflector (Optional)” on page 185.
 - To subdivide autonomous systems (ASs), see “Configuring BGP Confederations (Optional)” on page 188.

- To check the configuration, see “Verifying a BGP Configuration” on page 190.

Table 56: 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 drop-down menu, 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 322.	

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 54 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 54: 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 181.
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 57.
4. If you are finished configuring the network, commit the configuration.
5. Go on to one of the following procedures:
 - To subdivide autonomous systems (ASs), see “Configuring BGP Confederations (Optional)” on page 188.

- To check the configuration, see “Verifying a BGP Configuration” on page 190.

Table 57: 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 drop-down menu, 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 57: 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 322.	

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 55 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 55: Typical Network Using BGP Confederations



To configure the BGP confederations shown in Figure 55:

- 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 network, commit the configuration.
- 4. To check the configuration, see “Verifying a BGP Configuration” on page 190.

Table 58: 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 58: 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).	See “Configuring a Point-to-Point Peering Session (Required)” on page 181.	
When setting the peer AS number for these sessions, use the sub-AS number rather than the AS number.		
Using IBGP, configure internal sessions within a sub-AS. You can configure an IBGP full mesh, or you can configure a route reflector.	<ul style="list-style-type: none"> ■ To configure an IBGP full mesh, see “Configuring BGP Within a Network (Required)” on page 184. ■ To configure a route reflector, see “Configuring a Route Reflector (Optional)” on page 185. 	

Verifying a BGP Configuration

To verify a BGP configuration, perform these tasks:

- Verifying BGP Neighbors on page 190
- Verifying BGP Groups on page 191
- Verifying BGP Summary Information on page 192
- Verifying Reachability of All Peers in a BGP Network on page 193

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/bgpggr 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 Protocols, Class of Service, and System Basics 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 Protocols, Class of Service, and System Basics 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/R
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

The output shows a summary of BGP session information. Verify the following information:

- For Groups, the total number of configured groups is shown.
- For Peers, the total number of BGP peers is shown.
- For Down Peers, the total number of unestablished peers is 0. If this value is not zero, one or more peering sessions are not yet established.
- Under Peer, the IP address for each configured peer is shown.
- Under AS, the peer AS for each configured peer is correct.
- Under Up/Dwn State, the BGP state reflects the number of paths received from the neighbor, the number of these paths that have been accepted, and the number of routes being damped (such as 0/0/0). If the field is **Active**, it indicates a problem in the establishment of the BGP session.

For more information about `show bgp summary`, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

Verifying Reachability of All Peers in a BGP Network

Purpose

By using the ping tool on each peer address in the network, verify that all peers in the network are reachable from each Services Router.

Action

For each Services Router in the BGP network:

1. In the J-Web interface, select **Diagnose > Ping Host**.
2. In the Remote Host box, type the name of a host for which you want to verify reachability from the Services Router.
3. Click **Start**. Output appears on a separate page.

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 a host is active, it generates an ICMP response. If this response is received, the round-trip time is listed in the `time` field. For more information about the ping 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*.

Part 4

Configuring Private Communications over Public Networks with MPLS

- Multiprotocol Label Switching Overview on page 197
- Configuring Signaling Protocols for Traffic Engineering on page 213
- Configuring Virtual Private Networks on page 227
- Configuring IPSec for Secure Packet Exchange on page 251

Chapter 9

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 197
- MPLS Overview on page 199
- Signaling Protocols Overview on page 204
- VPN Overview on page 208

MPLS and VPN Terms

To understand MPLS and VPNs, become familiar with the terms defined in Table 59 .

Table 59: 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 59: 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 59: 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 200
- Label-Switched Paths on page 200
- Label-Switching Routers on page 201
- Labels on page 202

- Label Operations on page 202
- Penultimate Hop Popping on page 203
- LSP Establishment on page 203

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 56 shows a typical LSP topology.

Figure 56: Typical LSP Topology

In the topology shown in Figure 56, 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 204
- Resource Reservation Protocol on page 204

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 205
- Bandwidth Reservation Requirement on page 205
- Explicit Route Objects on page 205
- Constrained Shortest Path First on page 207
- Link Coloring on page 207

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 57 shows a typical RSVP-signaled LSP that uses EROs.

Figure 57: Typical RSVP-Signaled LSP with EROs



In the topology shown in Figure 57, 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

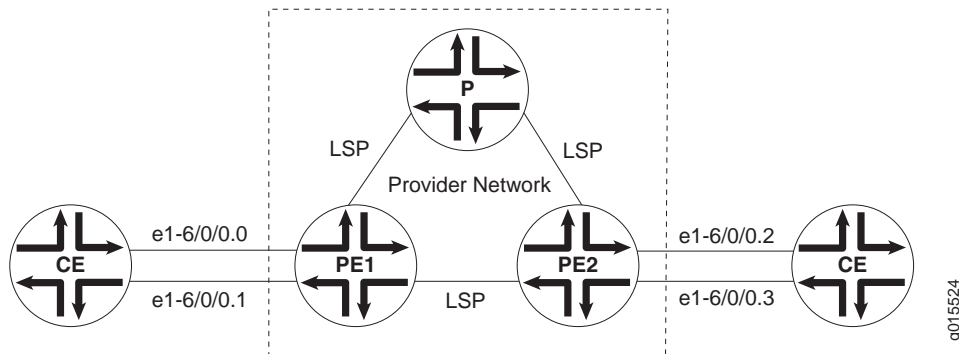
Virtual private networks (VPNs) are private networks that use a public network to connect two or more remote sites. In place of dedicated connections between networks, VPNs use virtual connections routed (tunneled) through public networks that are typically service provider networks. The type of the 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 208
- VPN Routing Requirements on page 209
- VPN Routing Information on page 210
- Types of VPNs on page 211

VPN Components

All types of VPNs share certain components. Figure 58 shows a typical VPN topology.

Figure 58: 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 10

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 213
- Before You Begin on page 214
- Configuring LDP and RSVP with a Configuration Editor on page 215
- Verifying an MPLS Configuration on page 220

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 41.
- Configure an interior gateway protocol (IGP) across your network. See “Configuring an OSPF Network” on page 155 or “Configuring a RIP Network” on page 139. 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 215
- Configuring RSVP-Signaled LSPs on page 217

For information about using the J-Web and CLI configuration editors, see “Using J-series 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 59.

Figure 59: 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 59, perform these steps:

1. Navigate to the top of the configuration hierarchy in either the J-Web or CLI configuration editor.
2. Perform the configuration tasks described in Table 60.

3. If you are finished configuring the network, commit the configuration.
4. Go on to “Verifying an LDP-Signaled LSP” on page 220.

Table 60: 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 <code>edit interfaces</code>
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: <code>set fe-0/0/0 unit 0 family mpls</code> 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 <code>all</code>. 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 <code>edit protocols mpls</code> 2. Enter <code>set interface all</code> 3. Repeat Steps 1 and 2 for each transit interface on the routers in the MPLS network.


Table 60: 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 <code>edit protocols ldp</code> 2. Enable LDP on a transit interface. For example: <code>set interface fe-0/0/0</code> 3. Repeat Steps 1 and 2 for each transit interface on the routers in the MPLS network.
Set the keepalive interval to 5 seconds.	<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. 	On each router in the MPLS network, enter <code>set keepalive-interval 5</code>
The keepalive interval specifies the number of seconds between the transmission of keepalive messages along the LDP link.		

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 60.

Figure 60: 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 60, 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 61.
3. If you are finished configuring the network, commit the configuration.
4. Go on to “Verifying an RSVP-Signaled LSP” on page 223.

Table 61: 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 61: 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 61: 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 220
- Verifying an RSVP-Signaled LSP on page 223

Verifying an LDP-Signaled LSP

Suppose that LDP is configured to establish an LSP as shown in Figure 59.

To verify the LDP configuration, perform these verification tasks:

- Verifying LDP Neighbors on page 220
- Verifying LDP Sessions on page 221
- Verifying the Presence of LDP-Signaled LSPs on page 222
- Verifying Traffic Forwarding over the LDP-Signaled LSP on page 222

Verifying LDP Neighbors

Purpose	Verify that each Services Router shows the appropriate LDP neighbors—for example, that router R5 has both router R6 and router R7 as LDP neighbors.
Action	From the CLI, enter the show ldp neighbor command.

Sample Output

```
user@r5> show ldp neighbor

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

What It Means

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

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

Verifying LDP Sessions

Purpose

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

Action

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

Sample Output

```
user@r5> show ldp session detail

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

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

Verifying the Presence of LDP-Signaled LSPs

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

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

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

## Verifying Traffic Forwarding over the LDP-Signaled LSP

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

**Sample Output**

```

user@c1> traceroute 220.220.0.0 source 200.200.0.1 bypass-routing gateway
172.16.0.1

traceroute to 220.220.0.1 (172.16.0.1) from 200.200.0.1, 30 hops max, 40 byte packets
 1 172.16.0.1 (172.16.0.1) 0.661 ms 0.538 ms 0.449 ms
 2 10.0.8.9 (10.0.8.9) 0.511 ms 0.479 ms 0.468 ms
 MPLS Label=100004 CoS=0 TTL=1 S=1
 3 10.0.8.5 (10.0.8.5) 0.476 ms 0.512 ms 0.441 ms
 4 220.220.0.1 (220.220.0.1) 0.436 ms 0.420 ms 0.416 ms

```

**What It Means**

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.

## Verifying an RSVP-Signaled LSP

Suppose that RSVP is configured to establish an LSP as shown in Figure 60.

To verify the RSVP configuration, perform these verification tasks:

- Verifying RSVP Neighbors on page 223
- Verifying RSVP Sessions on page 224
- Verifying the Presence of RSVP-Signaled LSPs on page 224

### Verifying RSVP Neighbors

**Purpose** Verify that each Services Router shows the appropriate RSVP neighbors—for example, that router R1 lists both router R3 and router R2 as RSVP neighbors.

**Action** From the CLI, enter the `show rsvp neighbor` command.

**Sample Output**

```

user@r1> show rsvp neighbor

RSVP neighbor: 2 learned
Address Idle Up/Dn LastChange HelloInt HelloTx/Rx
10.0.6.2 0 3/2 13:01 3 366/349
10.0.3.3 0 1/0 22:49 3 448/448

```

**What It Means**

The output shows the IP addresses of the neighboring routers. Verify that each neighboring RSVP router loopback address is listed.

## Verifying RSVP Sessions

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

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

**Sample Output**

```
user@r1> show rsvp session detail

Ingress RSVP: 1 sessions

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

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

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

## Verifying the Presence of RSVP-Signaled LSPs

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

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

**Sample Output**

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

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

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

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





## Chapter 11

# 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 227
- Before You Begin on page 230
- Configuring VPNs with a Configuration Editor on page 230
- Verifying a VPN Configuration on page 248

## 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:

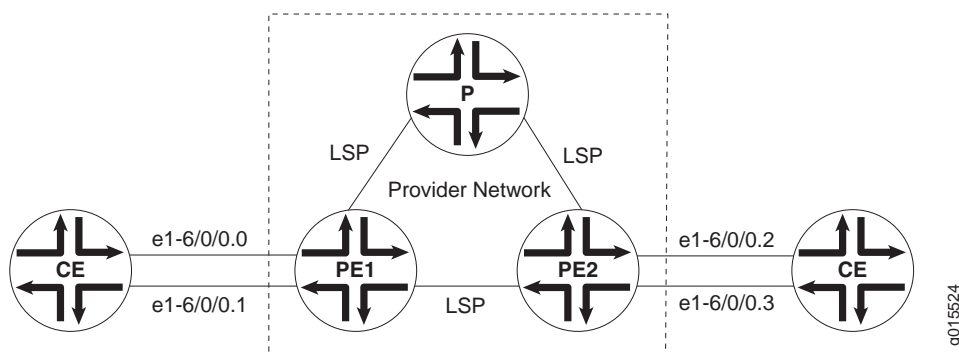
- Sample VPN Topology on page 228
- Basic Layer 2 VPN Configuration on page 228
- Basic Layer 2 Circuit Configuration on page 229

- Basic Layer 3 VPN Configuration on page 229

## Sample VPN Topology

Figure 61 shows the overview of a basic VPN topology for the sample configurations in this chapter.

**Figure 61: 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 41.
- Determine the protocols to use in the VPN configuration. These protocols include
  - MPLS—See “Multiprotocol Label Switching Overview” on page 197 and the *JUNOS Routing Protocols Configuration Guide*.
  - BGP, EBGP, and internal BGP (IBGP)—See “Configuring BGP Sessions” on page 177 and the *JUNOS Routing Protocols Configuration Guide*.
  - LDP and Resource Reservation Protocol (RSVP)—See “Configuring Signaling Protocols for Traffic Engineering” on page 213 and the *JUNOS MPLS Applications Configuration Guide*.
  - OSPF—See “Configuring an OSPF Network” on page 155 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 62 to help you select the tasks for your VPN type. For information about using the J-Web and CLI configuration editors, see “Using J-series Configuration Tools” on page 3.

- Configuring Interfaces Participating in a VPN on page 231
- Configuring Protocols Used by a VPN on page 233
- Configuring a VPN Routing Instance on page 241
- Configuring a VPN Routing Policy on page 243

**Table 62: 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 231 | All Services Routers                                                             | All Services Routers                                    | All Services Routers   |
| “Configuring a VPN Routing Instance” on page 241            | PE Services Routers                                                              | PE Services Routers                                     | N/A                    |
| “Configuring a VPN Routing Policy” on page 243              | 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 41.

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 63 for each interface involved in the VPN, except Layer 3 loopback interfaces, which do not require other configuration.
3. When you have finished configuring the network, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 248.
5. Go on to “Configuring Protocols Used by a VPN” on page 233.

**Table 63: Configuring an Interface for a VPN**

| Task                                                                                                                                                                                                                                                                                                       | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | CLI Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure IPv4.<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 drop-down list. For Fast Ethernet interfaces, you also must select <b>Vlan tagging</b> from the Vlan tag mode drop-down 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> | <p>■ For all interfaces except loopback, and a Layer 2 VPN interface facing a CE router:</p> <p>From the top of the configuration hierarchy, enter</p> <pre>edit interfaces interface-name unit logical_interface family inet address ipv4_address</pre> <p>■ For a loopback address on a Layer 2 configuration:</p> <p>From the top of the configuration hierarchy, enter</p> <pre>edit interfaces lo0 unit logical_interface family inet address ipv4_address primary</pre> <p>■ For a Layer 2 VPN interface facing a CE router:</p> <p>From the top of the configuration hierarchy, enter</p> <pre>set interfaces interface-name vlan-tagging encapsulation vlan-ccc unit logical_interface encapsulation vlan-ccc vlan-id id-number</pre> |
| 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                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|                                                                                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | set unit <i>logical_interface</i> family mpls                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| 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 drop-down list.</li> <li>2. Click <b>OK</b>.</li> <li>3. On the Interface page, select an encapsulation type from the Encapsulation drop-down 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</li> </ol> <pre>set encapsulation encapsulation_type</pre> <ol style="list-style-type: none"> <li>2. Enter</li> </ol> <pre>set unit logical_interface encapsulation encapsulation_type</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## 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 64 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*.

- Configuring MPLS for VPNs on page 233
- Configuring a BGP Session on page 235
- Configuring Routing Options for VPNs on page 236
- Configuring an IGP and a Signaling Protocol on page 237
- Configuring LDP for Signaling on page 237
- Configuring RSVP for Signaling on page 239
- Configuring a Layer 2 Circuit on page 240

**Table 64: VPN Protocol Configuration Task Summary**

| Section                                                                                                                                                                                                                                               | Layer 3 VPN                      | Layer 2 VPN                      | Layer 2 Circuit      |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------|----------------------------------|----------------------|
| “Configuring MPLS for VPNs” on page 233                                                                                                                                                                                                               | N/A unless you are using RSVP    | PE and provider Services Routers | PE Services Routers  |
| “Configuring a BGP Session” on page 235                                                                                                                                                                                                               | PE Services Routers              | PE Services Routers              | PE Services Routers  |
| “Configuring Routing Options for VPNs” on page 236                                                                                                                                                                                                    | All Services Routers             | All Services Routers             | All Services Routers |
| “Configuring an IGP and a Signaling Protocol” on page 237— <i>one</i> of the following tasks: <ul style="list-style-type: none"> <li>■ “Configuring LDP for Signaling” on page 237</li> <li>■ “Configuring RSVP for Signaling” on page 239</li> </ul> | PE and provider Services Routers | PE Services Routers              | PE Services Routers  |
| “Configuring a Layer 2 Circuit” on page 240                                                                                                                                                                                                           | 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 197 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 65 on each PE Services Router and provider Services Router interface that communicates with another PE Services Router.
3. When you have finished configuring the network, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 248.
5. Go on to “Configuring a BGP Session” on page 235.

**Table 65: 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>                                                                                                                                                                                                                          | <p>From the top of the configuration hierarchy, enter the following command for each interface you want to enable:</p> <pre>edit protocols mpls interface <i>interface-name</i></pre>                                                                                                                                                   |
| For RSVP only, configure an MPLS label-switched path (LSP) to the destination point on the PE router for LSP. During configuration, you specify the IP address of the LSP destination point, which is an address on the remote PE router.<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/><pre>edit protocols mpls label-switched-path <i>path-name</i></pre></li> <li>2. Enter<br/><pre>set to <i>ip-address</i></pre></li> <li>3. Enter <i>up</i>.</li> <li>4. Enter<br/><pre>interface <i>interface-name</i></pre></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 177 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 66 on each PE router.
3. When you have finished configuring the network, commit the configuration.
4. To verify the configuration, “Verifying a VPN Configuration” on page 248.
5. Go on to “Configuring an IGP and a Signaling Protocol” on page 237.

**Table 66: 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                                                   |
|                                                                                                                | 2. Next to Group, click <b>Add New Entry</b> .                                                      | <code>edit protocols bgp group group-name</code>                                                        |
|                                                                                                                | 3. Type a name in the Group name box.                                                               | 2. Enter                                                                                                |
|                                                                                                                | 4. From the Type drop-down list, select <b>Internal</b> .                                           | <code>set type internal</code>                                                                          |
|                                                                                                                | 5. In the Local address box, type the local loopback IP address.                                    | 3. Enter                                                                                                |
|                                                                                                                | 6. In the Family group, select <b>L2vpn</b> for a Layer 2 VPN or <b>Inet vpn</b> for a Layer 3 VPN. | <code>set local-address loopback-interface-ip-address</code>                                            |
|                                                                                                                | 7. Select <b>Unicast</b> .                                                                          | 4. Enter                                                                                                |
|                                                                                                                | 8. Click <b>OK</b> .                                                                                | <code>set family family-type unicast</code>                                                             |
|                                                                                                                | 9. In the Neighbor group, click <b>Add new entry</b> .                                              | Replace <i>family-type</i> with <i>l2vpn</i> for a Layer 2 VPN or <i>inet-vpn</i> for a Layer 3 VPN.    |
|                                                                                                                | 10. In the Address box, type the loopback IP address of the neighboring PE router.                  | 5. Enter <code>up</code> .                                                                              |
|                                                                                                                | 11. Click <b>OK</b> until you return to the BGP page.                                               | 6. Enter the loopback address of the neighboring PE router:<br><br><code>set neighbor ip-address</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 67.
3. When you have finished configuring the network, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 248.
5. Go on to “Configuring an IGP and a Signaling Protocol” on page 237.

**Table 67: 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 204.

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 155, “Configuring Static Routes” on page 127, and the *JUNOS Routing Protocols Configuration Guide*.

Configure the appropriate signaling protocol for your VPN:

- “Configuring LDP for Signaling” on page 237
- “Configuring RSVP for Signaling” on page 239

## 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 155.

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 68 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 231.

3. When you have finished configuring the network, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 248.
5. Go on to “Configuring a VPN Routing Instance” on page 241.

**Table 68: 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 <b>(0.0.0.0)</b>.</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 <b>[edit protocols ospf]</b> 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 155.

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 69 on each PE router and provider router, as specified.
3. When you have finished configuring the network, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 248.
5. Go on to “Configuring a VPN Routing Instance” on page 241.

**Table 69: 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 70 on each PE router and provider router, as specified.
3. When you have finished configuring the network, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 248.

**Table 70: 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 71 on each PE router.
3. When you have finished configuring the network, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 248.
5. Go on to “Configuring a VPN Routing Policy” on page 243.

**Table 71: 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 " <i>text</i> "</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 drop-down 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 71: 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 244.<br><br>(PE Services Router) | <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>                                                                                                                                                        |
| 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 247.<br><br>(PE Services Router)    | <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> |

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

- Configuring a Routing Policy for Layer 2 VPNs on page 244
- Configuring a Routing Policy for Layer 3 VPNs on page 247

## 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 72 and Table 73 on each PE router.
3. When you have finished configuring the network, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 248.

**Table 72: 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 72: 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 drop-down 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 drop-down 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 73: 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 drop-down 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 drop-down 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 <i>as-number: number</i> or <i>ip-address: number</i>.</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 74 on each CE Services Router.
3. When you have finished configuring the network, commit the configuration.
4. To verify the configuration, see “Verifying a VPN Configuration” on page 248.

**Table 74: 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 74: 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 drop-down 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 drop-down 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 drop-down 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 249
- Pinging a Layer 3 VPN on page 249
- Pinging a Layer 2 Circuit on page 249

### **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 12

# 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 251
- Before You Begin on page 252
- Configuring an IPSec Tunnel with Quick Configuration on page 252
- Configuring an IPSec Tunnel with a Configuration Editor on page 254
- Verifying the IPSec Tunnel Configuration on page 264

## 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.

## Securing Incoming Traffic

Incoming (ingress) traffic across the tunnel must be secured to ensure that the IPsec tunnel is protected. Typically, you secure incoming traffic by configuring a stateful firewall filter that acts on the incoming flow through the tunnel. By filtering all traffic that does not match the remote gateway address, you ensure that only traffic sent by the tunnel endpoint reaches destinations through the IPsec tunnel.

## 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 41.
- Configure one or more routing protocols. See “Configuring Static Routes” on page 127, “Configuring a RIP Network” on page 139, “Configuring an OSPF Network” on page 155, or “Configuring BGP Sessions” on page 177.

## Configuring an IPsec Tunnel with Quick Configuration

---

J-Web Quick Configuration allows you to create IPsec tunnels. Figure 62 shows the Quick Configuration page for IPsec tunnels.

**Figure 62: 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**

| <input type="text"/> | / | <input type="text"/> |
|----------------------|---|----------------------|

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

To configure an IPSec tunnel with Quick Configuration:

1. In the J-Web user interface, select **Configuration > IPSec Tunnels**.
2. Enter information into the Quick Configuration page for IPSec Tunnels, as described in Table 75.
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 264.

**Table 75: 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 255
- Configuring IPSec Service Sets on page 256
- Configuring an IPSec Stateful Firewall Filter on page 260
- Configuring a NAT Pool on page 262

## 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 76.
3. If you are finished configuring the network, commit the configuration.
4. Go on to “Configuring IPsec Service Sets” on page 256.

**Table 76: 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 76: 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 drop-down menu.  |                                                                      |
|                                                                                                                                                                                                                                                                  | 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 drop-down menu. |                                                                      |
|                                                                                                                                                                                                                                                                  | 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 77.
3. If you are finished configuring the network, commit the configuration.
4. Go on to “Configuring an IPSec Stateful Firewall Filter” on page 260.

**Table 77: 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 drop-down menu.</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, <b>sp-0/0/0.1001</b>.</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, <b>sp-0/0/0.2002</b>.</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/><b>edit services</b></li> <li>2. Set the inside-service interface:<br/><br/><b>set service-set service-set-name next-hop-service inside-service-interface sp-0/0/0.1001</b></li> <li>3. Set the outside-service interface:<br/><br/><b>set service-set service-set-name next-hop-service outside-service-interface sp-0/0/0.2001</b></li> </ol> |

**Table 77: 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 77: 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 drop-down menu.</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/><br/><code>edit services ipsec-vpn ike</code></li> <li>Configure the IKE pre-shared key in ASCII text format:<br/><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/><br/><code>edit services ipsec-vpn rule-name term term-name then</code></li> <li>Configure a dynamic security association that applies the IKE policy:<br/><br/><code>set dynamic ike-policy policy-name</code></li> </ol> |

**Table 77: 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 drop-down menu.</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 drop-down menu.</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

Configure stateful firewall filter rules to ensure that only desired traffic is permitted. This firewall is applied to all inbound traffic from the WAN. For this IPSec tunnel, desired traffic must be from the remote tunnel endpoint, destined for the local tunnel endpoint, and using either IPSec or IKE as an application protocol.

For more information about firewall filters, see “Configuring Firewall Filters and NAT” on page 331.

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 78.
- If you are finished configuring the network, commit the configuration.
- Go on to “Configuring a NAT Pool” on page 262.

**Table 78: Configuring an IPSec Stateful Firewall Filter**

| <b>Task</b>                                                        | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | <b>CLI Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Create the stateful firewall rule and apply it to inbound traffic. | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, click <b>Services &gt; Stateful firewall</b>.</li> <li>2. In the rule field, click <b>Add new entry</b>.</li> <li>3. In the Rule name box, type the name of the rule. It can be any unique string.</li> <li>4. In the Match direction field, select <b>Input</b> from the drop-down menu.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | <ol style="list-style-type: none"> <li>1. From the top of the configuration hierarchy, enter<br/><code>edit services stateful-firewall</code></li> <li>2. 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>1. In the Term field, click <b>Add new entry</b>.</li> <li>2. In the Term name box, type the name of the term. It can be any unique string.</li> <li>3. Click <b>From</b>.</li> <li>4. In the Destination address field, click <b>Add new entry</b>.</li> <li>5. In the address field, select <b>Enter specific value</b> from the drop-down menu.</li> <li>6. In the Address box, type the IP address of the local tunnel endpoint, in dotted decimal notation, and click <b>OK</b>.</li> <li>7. In the Source address field, click <b>Add new entry</b>.</li> <li>8. In the address field, select <b>Enter specific value</b> from the drop-down menu.</li> <li>9. In the Address box, type the IP address of the remote tunnel endpoint, in dotted decimal notation, and click <b>OK</b>.</li> <li>10. In the Applications field, click <b>Add new entry</b>.</li> <li>11. In the Application name field, type <code>junos-ipsec-esp</code>, and click <b>OK</b>.</li> <li>12. In the Applications field, click <b>Add new entry</b>.</li> <li>13. In the Application name field, type <code>junos-ike</code>, and click <b>OK</b>.</li> </ol> | <ol style="list-style-type: none"> <li>1. 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>2. 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>3. 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>4. 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 78: Configuring an IPSec Stateful Firewall Filter (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 drop-down menu, 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 drop-down menu.</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 79.
3. If you are finished configuring the network, commit the configuration.
4. To check the configuration, see “Verifying the IPSec Tunnel Configuration” on page 264.

**Table 79: 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 drop-down menu.</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 79: 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 drop-down menu.</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 drop-down menu.</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 drop-down menu.</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 Network and Services Interfaces Command Reference*.





## **Part 5**

# **Managing Multicast Transmissions**

- Multicast Overview on page 269
- Configuring a Multicast Network on page 279



## Chapter 13

# 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 279.

- Multicast Terms on page 269
- Multicast Architecture on page 272
- Dense and Sparse Routing Modes on page 274
- Strategies for Preventing Routing Loops on page 274
- Multicast Protocol Building Blocks on page 275

## Multicast Terms

---

To understand multicast routing, you must be familiar with the terms defined in Table 80. See Figure 63 for a general view of some of the elements commonly used in an IP multicast network architecture.

**Table 80: Multicast Terms**

| <b>Term</b>                                               | <b>Definition</b>                                                                                                                                                                                                                                                                                                                                                                                                                 |
|-----------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>administrative scoping</b>                             | Multicast routing strategy that limits the routers and interfaces used to forward a multicast packet by reserving a range of multicast addresses.                                                                                                                                                                                                                                                                                 |
| <b>Auto-RP</b>                                            | Cisco multicast routing protocol that allows sparse-mode routing protocols to find rendezvous points (RPs) within a routing domain.                                                                                                                                                                                                                                                                                               |
| <b>bootstrap router (BSR)</b>                             | Multicast mechanism that allows routers running PIM sparse mode to find rendezvous points (RPs) within a routing domain.                                                                                                                                                                                                                                                                                                          |
| <b>branch</b>                                             | 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.                                                                                                                                                                                  |
| <b>broadcast routing protocol</b>                         | Protocol that distributes traffic from a particular source to all destinations.                                                                                                                                                                                                                                                                                                                                                   |
| <b>dense mode</b>                                         | Multicast routing mode appropriate for LANs with many interested receivers.                                                                                                                                                                                                                                                                                                                                                       |
| <b>Distance Vector Multicast Routing Protocol (DVMRP)</b> | Distributed multicast routing protocol that dynamically generates IP multicast distribution trees using reverse-path multicasting (RPM) to forward multicast traffic to downstream interfaces.                                                                                                                                                                                                                                    |
| <b>distribution tree</b>                                  | 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. |
| <b>downstream interface</b>                               | Interface on a multicast router that is leading toward the receivers. You can configure all the logical interfaces except one as downstream interfaces.                                                                                                                                                                                                                                                                           |
| <b>group address</b>                                      | 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.                                                                                                                                                                  |
| <b>Internet Group Management Protocol (IGMP)</b>          | Multicast routing protocol that runs between receiver hosts and routers to determine whether group members are present. Services Routers support IGMPv1, IGMPv2, and IGMPv3.                                                                                                                                                                                                                                                      |
| <b>leaf</b>                                               | 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.                                                                                                                            |
| <b>listener</b>                                           | Another name for a receiver in a multicast network.                                                                                                                                                                                                                                                                                                                                                                               |

**Table 80: 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***

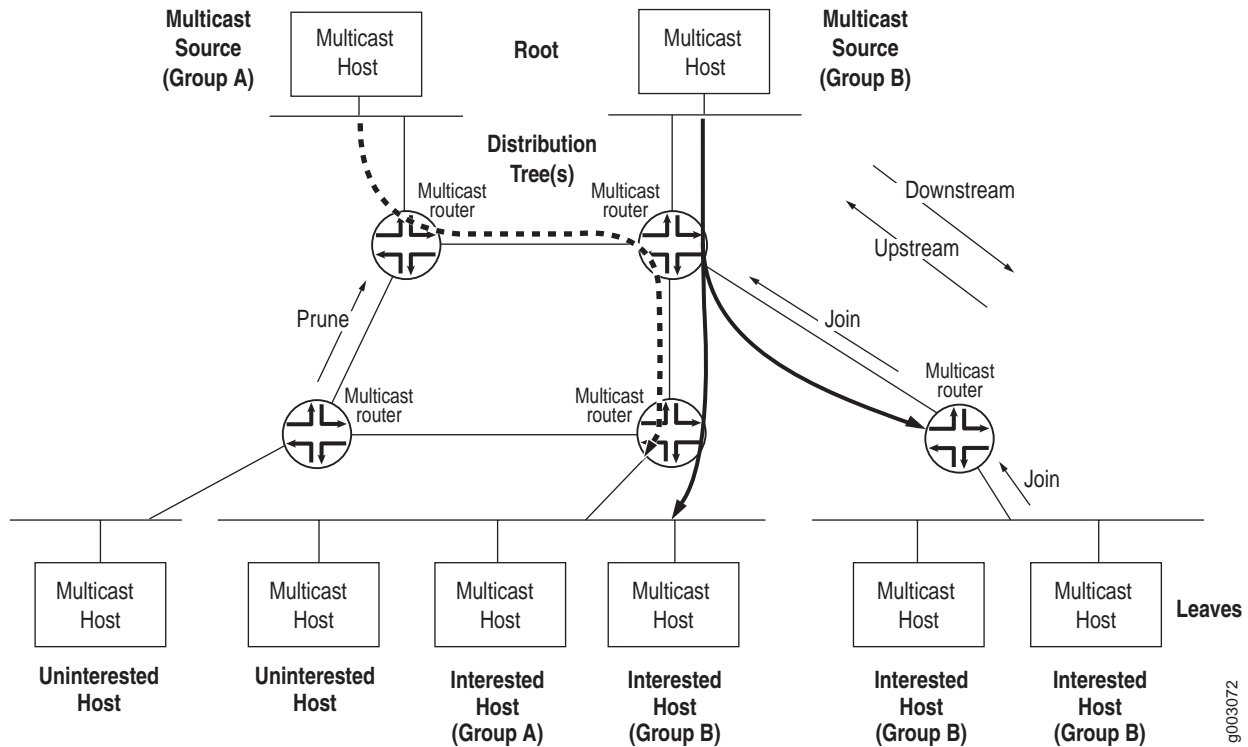
A single upstream interface on the router leads toward the source to receive multicast packets. The downstream interfaces on the router lead toward the receivers to transmit packets. A router can have as many downstream interfaces as it has logical interfaces, minus 1. To prevent looping, the router's 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 63). 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 63: 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 81.



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

**Table 81: 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 82 lists and summarizes these protocols.

**Table 82: 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 82: 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 82: 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 14

# 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 280
- Configuring a Multicast Network with a Configuration Editor on page 280
- Verifying a Multicast Configuration on page 285

## 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 269.
- Determine whether the Services Router is directly attached to any multicast sources. Receivers must be able to locate these sources.
- Determine whether the Services Router is directly attached to any multicast group receivers. If receivers are present, IGMP is needed.
- Determine whether to use the sparse, dense, or sparse-dense mode of multicast operation. Each mode has different configuration considerations.
- Determine the address of the rendezvous point (RP) if sparse or sparse-dense mode is used.
- Determine whether to locate the RP with the static configuration, bootstrap router (BSR), or Auto-RP method.
- Determine whether to configure multicast to use its own reverse-path forwarding (RPF) routing table when configuring PIM in sparse, dense, or sparse-dense modes.

## Configuring a Multicast Network with a Configuration Editor

---

To configure the Services Router as a node in a multicast network, you must perform the following tasks marked *(Required)*. For information about using the J-Web and CLI configuration editors, see “Using J-series Configuration Tools” on page 3.

- Configuring SAP and SDP (Optional) on page 280
- Configuring IGMP (Required) on page 281
- Configuring the PIM Static RP (Optional) on page 282
- Configuring a PIM RPF Routing Table (Optional) on page 284

### 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 83.
3. Go on to “Configuring IGMP (Required)” on page 281.

**Table 83: 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 84.
3. If you are finished configuring the network, 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 282.
  - To check the configuration, see “Verifying a Multicast Configuration” on page 285.

**Table 84: 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> <p><b>edit protocols igmp</b></p>                                                                                                                                                                                                                              |
| Set the IGMP version. By default, the Services Router uses IGMPv2, but this version can be changed through negotiation with hosts unless explicitly configured. | <ol style="list-style-type: none"> <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:<br/><br/><b>set igmp interface all</b></li> <li>2. Set the <b>version</b> value to <b>1</b>, <b>2</b>, or <b>3</b>. For example:<br/><br/><b>set igmp interface all version 2</b></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 85.
3. Go on to “Configuring a PIM RPF Routing Table (Optional)” on page 284.

**Table 85: 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 85: 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 86.
3. To check the configuration, see “Verifying a Multicast Configuration” on page 285.

**Table 86: 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                                                   |

**Table 86: Configuring a PIM RPF Routing Table (Continued)**

| <b>Task</b>                                                                                                                         | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                         | <b>CLI Configuration Editor</b>                                                                                                                                                                                   |
|-------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Configure a name for the RPF routing table group, and use <code>inet.2</code> 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, <code>multicast-rfp-rib</code>.</li> <li>2. In the Export rib box, type <code>inet.2</code>.</li> </ol>                                                                                      | <p>Type the name for the RPF routing table and set the export routing table to <code>inet.2</code>. For example:</p> <pre>set multicast-rpf-rib export-rib inet.2</pre>                                           |
| 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 <code>inet.2</code>.</li> <li>3. Click <b>OK</b> three times.</li> </ol>                                                                                                                      | <p>Set the import routing table to <code>inet.2</code>. For example:</p> <pre>set multicast-rpf-rib import-rib inet.2</pre>                                                                                       |
| 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> .                                                                                                                                                                                                                                  | <p>From the top of the configuration hierarchy, enter</p> <pre>edit protocols pim</pre>                                                                                                                           |
| 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, <code>multicast-rpf-rib</code>.</li> <li>2. Click <b>OK</b> three times.</li> </ol>                                                                                                                  | <p>Enter</p> <pre>set rib-group multicast-rpf-rib</pre>                                                                                                                                                           |
| 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>                                                                                                                             | <p>From the top of the configuration hierarchy, enter</p> <pre>edit routing-options rib-groups.</pre>                                                                                                             |
| Configure a name for the RPF routing table group, and use <code>inet.2</code> and <code>inet.0</code> 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, <code>if-rib</code>.</li> <li>2. Click <b>Add new entry</b> next to Import rib.</li> <li>3. In the Value box, type <code>inet.2 inet.0</code>.</li> <li>4. Click <b>OK</b> twice.</li> </ol> | <p>Type the name for the RPF routing table and set the export routing table to <code>inet.2</code> and <code>inet.0</code>. For example:</p> <pre>set if-rib import-rib inet.2 set if-rib import-rib inet.0</pre> |
| 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, <code>if-rib</code>.</li> <li>3. Click <b>OK</b>.</li> </ol>                                              | <p>From the top of the configuration hierarchy, enter</p> <pre>edit routing-options interface-routes set rib-group inet if-rib</pre>                                                                              |

## Verifying a Multicast Configuration

To verify a multicast configuration, perform these tasks:

- Verifying SAP and SDP Addresses and Ports on page 286
- Verifying the IGMP Version on page 286
- Verifying the PIM Mode and Interface Configuration on page 287
- Verifying the PIM RP Configuration on page 287
- Verifying the RPF Routing Table Configuration on page 288

## Verifying SAP and SDP Addresses and Ports

**Purpose** Verify that SAP and SDP are configured to listen on the correct group addresses and ports.

**Action** From the CLI, enter the `show sap listen` command.

### Sample Output

```
user@host> show sap listen
```

```
Group Address Port
224.2.127.254 9875
```

**What It Means** The output shows a list of the group addresses and ports that SAP and SDP listen on. Verify the following information:

- Each group address configured, especially the default 224.2.127.254, is listed.
- Each port configured, especially the default 9875, is listed.

For more information about `show sap listen`, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

## Verifying the IGMP Version

**Purpose** Verify that IGMP version 2 is configured on all applicable interfaces.

**Action** From the CLI, enter the `show igmp interface` command.

### Sample Output

```
user@host> show igmp interface
```

```
Interface: fe-0/0/0.0
 Querier: 192.168.4.36
 State: Up Timeout: 197 Version: 2 Groups: 0
```

```
Configured Parameters:
IGMP Query Interval: 125.0
IGMP Query Response Interval: 10.0
IGMP Last Member Query Interval: 1.0
IGMP Robustness Count: 2
```

```

Derived Parameters:
IGMP Membership Timeout: 260.0
IGMP Other Querier Present Timeout: 255.0

```

**What It Means** The output shows a list of the Services Router interfaces that are configured for IGMP. Verify the following information:

- Each interface on which IGMP is enabled is listed.
- Next to Version, the number 2 appears.

For more information about `show igmp interface`, see the *JUNOS Protocols, Class of Service, and System Basics 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
pim0.32769 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 Protocols, Class of Service, and System Basics Command Reference*.

## Verifying the PIM RP Configuration

**Purpose** Verify that the PIM RP is statically configured with the correct IP address.

**Action** From the CLI, enter the `show pim rps` command.

### Sample Output

```

user@host> show pim rps

Instance: PIM.master
Address family INET
RP address Type Holdtime Timeout Active groups Group prefixes

```

```
192.168.14.27 static 0 None 2 224.0.0.0/4
```

**What It Means** The output shows a list of the RP addresses that are configured for PIM. At least one RP must be configured. Verify the following information:

- The configured RP is listed with the proper IP address.
- Under **Type**, the word **static** appears.

## Verifying the RPF Routing Table Configuration

**Purpose** Verify that the PIM RPF routing table is configured correctly.

**Action** From the CLI, enter the `show multicast rpf` command.

**Sample Output**

```
user@host> show multicast rpf

Multicast RPF table: inet.0 , 2 entries...
```

**What It Means** The output shows the multicast RPF table that is configured for PIM. If no multicast RPF routing table is configured, RPF checks use `inet.0`. Verify the following information:

- The configured multicast RPF routing table is `inet.0`.
- The `inet.0` table contains entries.

For more information about `show multicast rpf`, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

## **Part 6**

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

- Policy, Firewall Filter, and Class-of-Service Overview on page 291
- Configuring Routing Policies on page 317
- Configuring Firewall Filters and NAT on page 331
- Configuring Class of Service with DiffServ on page 371





## Chapter 15

# 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 291
- Routing Policy Overview on page 293
- Firewall Filter Overview on page 298
- Class-of-Service Overview on page 307

## 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 87.

**Table 87: 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 87: 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 293
- Applying Routing Policies on page 298

### 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.

- Routing Policy Terms on page 294
- Routing Policy Match Conditions on page 294
- Routing Policy Actions on page 296

- Default and Final Actions on page 298

## 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 88 summarizes the routing policy match conditions.

**Table 88: Summary of Routing Policy Match Conditions**

| Match Condition                    | Description                                                                                                                                                                                                                                                                  |
|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>aggregate-contributor</code> | Matches routes that are contributing to a configured aggregate. This match condition can be used to suppress a contributor in an aggregate route.                                                                                                                            |
| <code>area area-id</code>          | Matches a route learned from the specified OSPF area during the exporting of OSPF routes into other protocols.                                                                                                                                                               |
| <code>as-path name</code>          | Name of an AS path regular expression. BGP routes whose AS path matches the regular expression are processed.                                                                                                                                                                |
| <code>color preference</code>      | 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. |

**Table 88: Summary of Routing Policy Match Conditions (Continued)**

| Match Condition                                 | Description                                                                                                                                                                                                                                                                                                                                                                         |
|-------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 metric-type]                     | Matches external OSPF routes, including routes exported from one level to another. In this construct <b>type</b> is an optional keyword. The <b>metric-type</b> value can be either <b>1</b> or <b>2</b> . When you do not specify <b>type</b> , this condition matches all external routes.                                                                                        |
| interface interface-name                        | 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 level                                     | Matches the IS-IS level. Routes that are from the specified level or are being advertised to the specified level are processed.                                                                                                                                                                                                                                                     |
| local-preference value                          | BGP local preference attribute. The preference value can be from <b>0</b> through <b>4,294,967,295</b> ( $2^{32} - 1$ ).                                                                                                                                                                                                                                                            |
| metric metric<br>metric2 metric                 | Metric value. The <b>metric</b> value corresponds to the multiple exit discriminator (MED), and <b>metric2</b> corresponds to the interior gateway protocol (IGP) metric if the BGP next hop runs back through another route.                                                                                                                                                       |
| neighbor address                                | 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 address                                | 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 value                                    | 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 [ policy-names ]                         | Name of one or more policies to evaluate as a subroutine.                                                                                                                                                                                                                                                                                                                           |
| preference preference<br>preference2 preference | Preference value. You can specify a primary preference value ( <b>preference</b> ) and a secondary preference value ( <b>preference2</b> ). The preference value can be a number from <b>0</b> through <b>4,294,967,295</b> ( $2^{32} - 1$ ). A lower number indicates a more preferred route.                                                                                      |
| prefix-list name                                | 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 protocol                               | 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>dvmp</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> .                                                     |

**Table 88: Summary of Routing Policy Match Conditions (Continued)**

| Match Condition                                                                  | Description                                                                                                                                                                                                                                                                                                                                                        |
|----------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>route-filter destination-prefix match-type &lt;actions&gt;</code>          | <p>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.</p> <p>Route filters can be used on import policies only.</p>            |
| <code>route-type value</code>                                                    | Type of route. The value can be either <b>external</b> or <b>internal</b> .                                                                                                                                                                                                                                                                                        |
| <code>source-address-filter destination-prefix match-type &lt;actions&gt;</code> | <p>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.</p> <p>Source-address filters can be used on import policies only.</p> |

## 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 89 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 89: 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>                                                                                                                                                                                                                                                                                                                                                                                    |
| local-preference <i>value</i>                       | Sets the BGP local preference attribute. The preference can be a number from 0 through 4,294,967,295 ( $2^{32} - 1$ ).                                                                                                                                                                                                                                                                                                                                                                                                                                                   |

**Table 89: Summary of Key Routing Policy Actions (Continued)**

| Action                  | Description                                                                                                                                                                                                                                                                                                                                                    |
|-------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 299
- Process for Configuring a Stateful Firewall Filter and NAT on page 300
- Summary of Stateful Firewall Filter and NAT Match Conditions and Actions on page 300
- Planning a Stateless Firewall Filter on page 302
- Stateless Firewall Filter Match Conditions, Actions, and Action Modifiers on page 303

## 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 251 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 300.

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

Table 90 lists the match conditions you can specify in stateful firewall filter and NAT terms. Table 91 and Table 92 list actions you can specify in stateful firewall filter and NAT terms.

**Table 90: 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 91: 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 92: 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 92: NAT Actions (Continued)**

| Actions                                           | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| translation-type (destination type   source type) | <p>Translate the destination and source port using the specified type:</p> <ul style="list-style-type: none"> <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 <b>[edit nat pool]</b> 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 <b>[edit nat pool]</b> 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 <b>[edit nat pool]</b> 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 303. For more information about stateless firewall filters, see the *JUNOS Policy Framework Configuration Guide*.

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

Table 93 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 drop-down 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 94 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 93: 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 93: 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 93: 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 94: Stateless Firewall Filter Bit-Field Logical Operators**

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

Table 95 lists the actions and action modifiers you can specify in stateless firewall filter terms.

**Table 95: 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>routing-instance | Routes the packet using the specified routing instance.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>Action Modifiers</b>              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |



**Table 95: 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 371.

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 307
- DSCPs and Forwarding Service Classes on page 308
- JUNOS CoS Functions on page 309
- How Forwarding Classes and Schedulers Work on page 311

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

**Table 96: 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> |

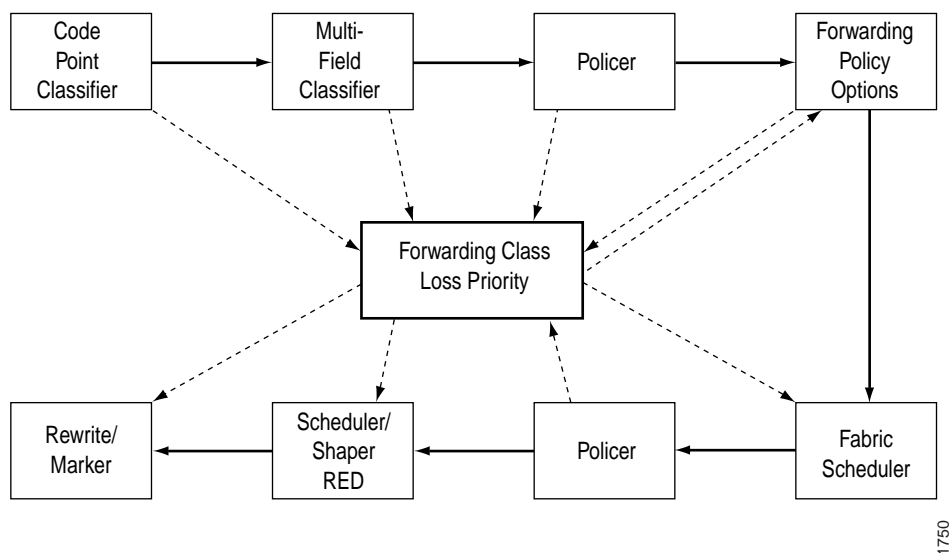
**Table 96: Default Forwarding Service Class-to-DSCP Mapping (Continued)**

| <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 64 shows the components of the JUNOS CoS features, illustrating the sequence in which they interact. Table 97 defines the components and explains their use.

**Figure 64: Packet Flow Through JUNOS CoS-Configurable Components****Table 97: 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 97: 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 311
- Default Scheduler Settings on page 312
- Default Behavior Aggregate (BA) Classifiers on page 313
- DSCP Rewrites on page 314
- Sample BA Classification on page 314

### **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 98 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 98: 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 99 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 371.

**Table 99: 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 99: 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 379.

## Sample BA Classification

Table 100 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 371.

**Table 100: 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 100: 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 16

# 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 318
- Configuring a Routing Policy with a Configuration Editor on page 318

## 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 293.
- 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 41.
- Configure an Interior Gateway Protocol (IGP) and Border Gateway Protocol (BGP), if necessary. See “Configuring BGP Sessions” on page 177.
- Configure the router interface to reject or accept routes, if necessary. See “Configuring Firewall Filters and NAT” on page 331.
- Configure static routes, if necessary. See “Configuring Static Routes” on page 127.

## Configuring a Routing Policy with a Configuration Editor

---

A routing policy has a major impact on the flow of routing information or packets within and through the Services Router. The match conditions and actions allow you to configure a customized policy to fit your needs.

To configure a routing policy, you must perform the following tasks marked *(Required)*. Perform additional tasks as needed for your router. For information about using the J-Web and CLI configuration editors, see “Using J-series Configuration Tools” on page 3.

- Configuring the Policy Name (Required) on page 319
- Configuring a Policy Term (Required) on page 319
- Rejecting Known Invalid Routes (Optional) on page 320
- Injecting OSPF Routes into the BGP Routing Table (Optional) on page 322
- Grouping Source and Destination Prefixes in a Forwarding Class (Optional) on page 324
- Configuring a Policy to Prepend the AS Path (Optional) on page 325
- Configuring Damping Parameters (Optional) on page 327

## 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 101.
3. Go on to “Configuring a Policy Term (Required)” on page 319.

**Table 101: 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 102.
3. If you are finished configuring the policy, commit the configuration.
4. Go on to one of the following procedures:

- To remove useless routes, see “Rejecting Known Invalid Routes (Optional)” on page 320.
- To advertise additional routes, see “Injecting OSPF Routes into the BGP Routing Table (Optional)” on page 322.
- To create a forwarding class, see “Grouping Source and Destination Prefixes in a Forwarding Class (Optional)” on page 324.
- To make a route less preferable to BGP, see “Configuring a Policy to Prepend the AS Path (Optional)” on page 325.
- To suppress route information, see “Configuring Damping Parameters (Optional)” on page 327.

**Table 102: 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 103 lists route list match types.

**Table 103: 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 103: 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 104.
3. If you are finished configuring the policy, commit the configuration.
4. Go on to one of the following procedures:
  - To advertise additional routes, see “Injecting OSPF Routes into the BGP Routing Table (Optional)” on page 322.
  - To create a forwarding class, see “Grouping Source and Destination Prefixes in a Forwarding Class (Optional)” on page 324.
  - To make a route less preferable to BGP, see “Configuring a Policy to Prepend the AS Path (Optional)” on page 325.
  - To suppress route information, see “Configuring Damping Parameters (Optional)” on page 327.

**Table 104: Creating a Policy to Reject Known Invalid 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 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 105.
3. If you are finished configuring the policy `injectpolicy1`, commit the configuration.
4. 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 324.
  - To make a route less preferable to BGP, see “Configuring a Policy to Prepend the AS Path (Optional)” on page 325.
  - To suppress route information, see “Configuring Damping Parameters (Optional)” on page 327.

**Table 105: 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><code>edit policy-options policy-statement injectpolicy1 term injectterm1</code> |
| 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><code>set from ospf</code>                                                                            |
| 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><code>set from area 1</code>                                                                       |
| 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><code>set then accept</code>                                                                              |

**Table 105: 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 106.
3. If you are finished configuring the policy, commit the configuration.
4. 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 325.
  - To suppress route information, see “Configuring Damping Parameters (Optional)” on page 327.

**Table 106: 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 107.
3. If you are finished configuring the policy, commit the configuration.
4. Go on to “Configuring Damping Parameters (Optional)” on page 327.

**Table 107: 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 107: 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 108.
3. If you are finished configuring the policy, commit the configuration.

**Table 108: 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 108: Creating a Policy to Accept and Apply Damping on Routes (Continued)**

| Task                                                                                                                                                                                                   | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | CLI Configuration Editor                                                                                                                                                                                                          |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 17

# 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 332
- Configuring a Stateful Firewall Filter with Quick Configuration on page 332
- Configuring a Stateful Firewall Filter with a Configuration Editor on page 336
- Configuring a Stateless Firewall Filter with a Configuration Editor on page 342
- Verifying Firewall Filter Configuration on page 359

## 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 298.
- 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 41.

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 65 and Figure 66 show the Firewall/NAT Quick Configuration main and application pages.

**Figure 65: 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 66: Firewall/NAT Quick Configuration Application Page**

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](#) > [Firewall/NAT](#)

### Quick Configuration

## Firewall/NAT

### Allow an Application Through the Firewall

---

#### Application

\* **Application**

---

#### Source Address

**Any Unicast WAN Address** ☒

#### Source Addresses and Prefixes

| Source Address |                                                                          |
|----------------|--------------------------------------------------------------------------|
|                | <input type="button" value="Add"/> <input type="button" value="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 109.
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 359.
- To verify a stateful firewall filter, see “Verifying Firewall Filter Configuration” on page 359.

**Table 109: 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 109: Firewall/NAT Quick Configuration Pages Summary (Continued)**

| Field                              | Function                                                                                               | Your Action                                                                                                                                                                                                                                        |
|------------------------------------|--------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Application</b>                 |                                                                                                        |                                                                                                                                                                                                                                                    |
| Application (required)             | Designate which applications are allowed to operate from the untrusted network to the trusted network. | From the drop-down 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 110.

**Table 110: 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 300.

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 111.
3. To apply the stateful firewall filter and NAT to the interface, perform the configuration tasks described in Table 112.
4. If you are finished configuring the network, commit the configuration.
5. Go on to one of the following procedures:
  - To display the configuration, see “Displaying Firewall Filter Configurations” on page 359.
  - To verify the stateful firewall filter, see “Verifying a Stateful Firewall Filter” on page 364.

**Table 111: 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 <b>edit services stateful-firewall</b> .                                                                                           |
| 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 drop-down list, select <b>output</b>.</li> </ol>                                          | Set the rule name, match direction, term name, and match condition:<br><br><b>set rule to-wan-rule match-direction output term app-term from application-sets junos-algs-outbound</b> |
| 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 drop-down list, select <b>Accept</b>.</li> <li>3. Click <b>OK</b> twice.</li> </ol>                                          | Set the action:<br><br><b>set rule to-wan-rule term app-term then accept</b>                                                                                                          |



**Table 111: Configuring a Stateful Firewall Filter and NAT (Continued)**

| Task                                                                                                                          | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                       | CLI Configuration Editor                                                                                                                                                                      |
|-------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Define <b>accept-all-term</b> for <b>to-wan-rule</b> .                                                                        | <ol style="list-style-type: none"> <li>On the Rule <b>to-wan-rule</b> page, next to Term, click <b>Add new entry</b>.</li> <li>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>Next to Then, click <b>Configure</b>.</li> <li>From the Designation drop-down list, select <b>Accept</b>.</li> <li>Next to Accept, select the check box.</li> <li>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>On the Rule page, next to Rule, click <b>Add new entry</b>.</li> <li>In the Rule name box, type <b>from-wan-rule</b>.</li> <li>From the Match direction drop-down 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>Next to Term, click <b>Add new entry</b>.</li> <li>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>Next to From, click <b>Configure</b>.</li> <li>Next to Source address, click <b>Add new entry</b>.</li> <li>From the Address drop-down list, select <b>Enter Specific Value—&gt;</b>.</li> <li>In the Prefix box, type <b>192.168.33.0/24</b>.</li> <li>Click <b>OK</b> twice.</li> </ol> |                                                                                                                                                                                               |
| Define an action for <b>wan-src-addr-term</b> .                                                                               | <ol style="list-style-type: none"> <li>On the Term <b>wan-src-addr-term</b> page, next to Then, click <b>Configure</b>.</li> <li>In the Designation drop-down list, select <b>Accept</b>.</li> <li>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>On the Rule <b>from-wan-rule</b> page, next to Term, click <b>Add new entry</b>.</li> <li>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>Next to Then, click <b>Configure</b>.</li> <li>From the Designation drop-down list, select <b>Discard</b>.</li> <li>Click <b>OK</b> three times.</li> </ol>                                                                                                                               |                                                                                                                                                                                               |

**Table 111: 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 drop-down 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>                                                                                                                                                                                                                          | Set the address pool name and the range:<br><br><code>set pool public-pool address-range low 10.148.2.1 high 10.148.2.32</code>                                                                               |
| 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 drop-down list, select <b>Automatic</b>.</li> <li>3. Click <b>OK</b> twice.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                        | Configure the source port translation to be automatic:<br><br><code>set pool public-pool port automatic</code>                                                                                                |
| 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 drop-down 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> | Set the rule name, match direction, term name, and the term's pool name:<br><br><code>set rule nat-to-wan-rule match-direction output term private-public-term then translated source-pool public-pool</code> |
| 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 drop-down list, select <b>dynamic</b>.</li> <li>4. Click <b>OK</b> five times.</li> </ol>                                                                                                                                                                                                                                                                                                                           | Set the NAT translation type:<br><br><code>set rule nat-to-wan-rule match-direction output term private-public-term then translated translation-type source dynamic</code>                                    |

**Table 112: 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 drop-down 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 drop-down 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 drop-down 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 112: 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 set interfaces t1-0/0/0 unit 0 family inet 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 303.

- Stateless Firewall Filter Strategies on page 343
- Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources on page 344

- Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods on page 347
- Configuring a Routing Engine Firewall Filter to Handle Fragments on page 353
- Applying a Stateless Firewall Filter to an Interface on page 358

## 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 344 and “Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods” on page 347.

### 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 353.

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

**Table 113: Sample Stateless Firewall Filter `protect-RE` Terms to Allow Packets from Trusted Sources**

| Term                           | Purpose                                                                                                                                                                                                                                                                                                                                       |
|--------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <code>ssh-term</code>          | Accepts TCP packets with a source address of <code>192.168.122.0/24</code> and a destination port that specifies SSH.                                                                                                                                                                                                                         |
| <code>bgp-term</code>          | Accepts TCP packets with a source address of <code>10.2.1.0/24</code> and a destination port that specifies the BGP protocol.                                                                                                                                                                                                                 |
| <code>discard-rest-term</code> | For all packets that are not accepted by <code>ssh-term</code> or <code>bgp-term</code> , creates a firewall filter log and system logging records, then discards all packets. To view the log, enter the <code>show firewall log</code> operational mode command. (For more information, see “Displaying Firewall Filter Logs” on page 365.) |

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 114.
3. If you are finished configuring the network, commit the configuration.
4. Go on to one of the following procedures:
  - To display the configuration, see “Displaying Firewall Filter Configurations” on page 359.
  - To apply the firewall filter to the Routing Engine, see “Applying a Stateless Firewall Filter to an Interface” on page 358.
  - To verify the firewall filter, see “Verifying a Services, Protocols, and Trusted Sources Firewall Filter” on page 367.

**Table 114: Configuring a Protocols and Services Firewall Filter for the Routing Engine**

| 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>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 drop-down list, select <b>Protocol</b>.</li> <li>Next to Protocol, click <b>Add new entry</b>.</li> <li>In the Value keyword drop-down list, select <b>tcp</b>.</li> <li>Click <b>OK</b>.</li> <li>In the Destination port choice drop-down list, select <b>Destination port</b>.</li> <li>Next to Destination port, click <b>Add new entry</b>.</li> <li>In the Value keyword drop-down 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 drop-down 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 114: 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>On the Filter <b>protect-RE</b> page, next to Term, click <b>Add New Entry</b>.</li> <li>In the Rule name box, type <b>bgp-term</b>.</li> <li>Next to From, click <b>Configure</b>.</li> <li>In the Protocol choice drop-down list, select <b>Protocol</b>.</li> <li>Next to Protocol, click <b>Add new entry</b>.</li> <li>In the Value keyword drop-down list, select <b>tcp</b>.</li> <li>Click <b>OK</b>.</li> <li>In the Destination port choice drop-down list, select <b>Destination port</b>.</li> <li>Next to Destination port, click <b>Add new entry</b>.</li> <li>In the Value keyword drop-down 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 <b>10.2.1.0/24</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 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>On the Term <b>bgp-term</b> page, next to Then, click <b>Configure</b>.</li> <li>In the Designation drop-down list, select <b>Accept</b>.</li> <li>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>On the Filter <b>protect-RE</b> page, next to Term, click <b>Add New Entry</b>.</li> <li>In the Rule name box, type <b>discard-rest-term</b>.</li> <li>Next to Then, click <b>Configure</b>.</li> <li>Next to Log, select the check box.</li> <li>Next to Syslog, select the check box.</li> <li>In the Designation drop-down list, select <b>Discard</b>.</li> <li>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 344), 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 115 lists the terms that are configured in this sample filter.

**Table 115: 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 116.
3. To configure the prefix lists and the firewall filter, perform the configuration tasks described in Table 117.
4. If you are finished configuring the network, commit the configuration.
5. Go on to one of the following procedures:
  - To display the configuration, see “Displaying Firewall Filter Configurations” on page 359.
  - To apply the firewall filter to the Routing Engine, see “Applying a Stateless Firewall Filter to an Interface” on page 358.
  - To verify the firewall filter, see “Verifying a TCP and ICMP Flood Firewall Filter” on page 368.

**Table 116: Configuring Policers for TCP and ICMP**

| <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>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 <b>Policer</b>, click <b>Add new entry</b>.</li> <li>In the <b>Policer name</b> box, type <b>tcp-connection-policer</b>.</li> <li>Next to <b>Filter specific</b>, select the check box.</li> <li>Next to <b>If Exceeding</b>, select the check box and click <b>Configure</b>.</li> <li>In the <b>Burst size limit</b> box, type <b>15k</b>. The burst size limit can be from 1,500 through 100,000,000 bytes.</li> <li>In the <b>Bandwidth drop-down list</b>, select <b>Bandwidth limit</b>.</li> <li>In the <b>Bandwidth limit</b> 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 <b>Policer tcp-connection-policer</b> page, next to <b>Then</b>, click <b>Configure</b>.</li> <li>Next to <b>Discard</b>, 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 116: 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>■ k (1000)<br>■ m (1,000,000)<br>■ g (1,000,000,000) | <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 drop-down 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 117: 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 117: 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 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</pre> <pre>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 drop-down list, select <b>Protocol</b>.</li> <li>Next to Protocol, click <b>Add new entry</b>.</li> <li>In the Value keyword drop-down 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 117: 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>1. On the Term tcp-connection-term page, next to Then, click <b>Configure</b>.</li> <li>2. In the Policer box, type tcp-connection-policer.</li> <li>3. In the Designation drop-down list, select <b>Accept</b>.</li> <li>4. 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>1. On the Filter protect-RE page, next to Term, click <b>Add New Entry</b>.</li> <li>2. In the Rule name box, type icmp-term.</li> <li>3. Next to From, click <b>Configure</b>.</li> <li>4. In the Protocol choice drop-down list, select <b>Protocol</b>.</li> <li>5. Next to Protocol, click <b>Add new entry</b>.</li> <li>6. In the Value keyword drop-down list, select <b>icmp</b>.</li> <li>7. 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 117: Configuring a TCP and ICMP Flood Firewall Filter for the Routing Engine (Continued)**

| Task                                            | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | CLI Configuration Editor                                                                                                                                                  |
|-------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Define the ICMP type match conditions.          | <ol style="list-style-type: none"> <li>1. In the <code>Icmp</code> type choice drop-down list, select <b>Icmp type</b>.</li> <li>2. Next to <code>Icmp</code> type, click <b>Add new entry</b>.</li> <li>3. In the Value keyword drop-down list, select <b>echo-request</b>.</li> <li>4. Click <b>OK</b>.</li> <li>5. Next to <code>Icmp</code> type, click <b>Add new entry</b>.</li> <li>6. In the Value keyword drop-down list, select <b>echo-reply</b>.</li> <li>7. Click <b>OK</b>.</li> <li>8. Next to <code>Icmp</code> type, click <b>Add new entry</b>.</li> <li>9. In the Value keyword drop-down list, select <b>unreachable</b>.</li> <li>10. Click <b>OK</b>.</li> <li>11. Next to <code>Icmp</code> type, click <b>Add new entry</b>.</li> <li>12. In the Value keyword drop-down 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 <code>icmp-term</code> . | <ol style="list-style-type: none"> <li>1. On the <code>icmp-term</code> page, next to Then, click <b>Configure</b>.</li> <li>2. In the Count box, type <code>icmp-counter</code>.</li> <li>3. In the Policer box, type <code>icmp-policer</code>.</li> <li>4. In the Designation drop-down 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 118 lists the terms that are configured in this sample filter.

**Table 118: 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 119.
3. If you are finished configuring the network, commit the configuration.
4. Go on to one of the following procedures:
  - To display the configuration, see “Displaying Firewall Filter Configurations” on page 359.
  - To apply the firewall filter to the Routing Engine, see “Applying a Stateless Firewall Filter to an Interface” on page 358.
  - To verify the firewall filter, see “Verifying a Firewall Filter That Handles Fragments” on page 369.



**Table 119: 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 drop-down 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 drop-down 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 119: 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>On the Filter <b>fragment-RE</b> page, next to Term, click <b>Add New Entry</b>.</li> <li>In the Term name box, type <b>not-fragmented-term</b>.</li> <li>Next to From, click <b>Configure</b>.</li> <li>In the Fragment flags box, type <b>0x0</b>.</li> <li>In the Fragment offset choice drop-down 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>0</b>.</li> <li>Click <b>OK</b>.</li> <li>In the Protocol choice drop-down list, select <b>Protocol</b>.</li> <li>Next to Protocol, click <b>Add new entry</b>.</li> <li>In the Value keyword drop-down list, select <b>tcp</b>.</li> <li>Click <b>OK</b>.</li> <li>In the Destination port choice drop-down list, select <b>Destination port</b>.</li> <li>Next to Destination port, click <b>Add new entry</b>.</li> <li>In the Value keyword drop-down 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 <b>10.2.1.0/24</b>.</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 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>On the Term <b>not-fragmented-term</b> page, next to Then, click <b>Configure</b>.</li> <li>In the Designation drop-down 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 not-fragmented-term then accept</pre>                                                                                                                           |

**Table 119: 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 drop-down list, select <b>Protocol</b>.</li> <li>Next to Protocol, click <b>Add new entry</b>.</li> <li>In the Value keyword drop-down list, select <b>tcp</b>.</li> <li>Click <b>OK</b>.</li> <li>In the Destination port choice drop-down list, select <b>Destination port</b>.</li> <li>Next to Destination port, click <b>Add new entry</b>.</li> <li>In the Value keyword drop-down 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 drop-down 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 119: 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 drop-down 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> | <p>Set the term name and define match conditions:</p> <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 drop-down list, select <b>Accept</b>.</li> <li>3. Click <b>OK</b> four times.</li> </ol>                                                                                                                                                                                                                                                                   | <p>Set the action:</p> <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 120.
2. If you are finished configuring the network, commit the configuration.

**Table 120: Applying a Firewall Filter to the Routing Engine Interface**

| Task                                                                    | J-Web Configuration Editor                                                                                                                                                        | CLI Configuration Editor                                                        |
|-------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------|
| 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 359
- Verifying a Stateful Firewall Filter on page 364
- Displaying Firewall Filter Logs on page 365
- Displaying Firewall Filter Statistics on page 366
- Verifying a Services, Protocols, and Trusted Sources Firewall Filter on page 367
- Verifying a TCP and ICMP Flood Firewall Filter on page 368
- Verifying a Firewall Filter That Handles Fragments on page 369

## Displaying Firewall Filter Configurations

**Purpose** Verify the configuration of the firewall filter. You can analyze the flow of the firewall filter terms by displaying the entire configuration.

**Action** From the J-Web interface, select **Configuration > View and Edit > View Configuration Text**. Alternatively, from configuration mode in the CLI, enter the `show services` or `show firewall` command for stateful and stateless firewall filters.

The sample output in this section displays the following firewall filters (in order):

- Stateful firewall filter and NAT configured in “Configuring a Stateful Firewall Filter with a Configuration Editor” on page 336
- Stateless protect-RE filter configured in “Configuring a Routing Engine Firewall Filter for Services and Protocols from Trusted Sources” on page 344
- Stateless protect-RE filter configured in “Configuring a Routing Engine Firewall Filter to Protect Against TCP and ICMP Floods” on page 347
- Stateless fragment-RE filter configured in “Configuring a Routing Engine Firewall Filter to Handle Fragments” on page 353

**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 & lack) | 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 336.

**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 Protocols, Class of Service, and System Basics 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 344 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 Protocols, Class of Service, and System Basics 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 347 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 Protocols, Class of Service, and System Basics 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 344.

**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 Protocols, Class of Service, and System Basics 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 347.

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

## Verifying a Firewall Filter That Handles Fragments

- |                |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Purpose</b> | Verify the firewall filter configured in “Configuring a Routing Engine Firewall Filter to Handle Fragments” on page 353.                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| <b>Action</b>  | <p>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 <i>not</i> taken for packets that do not match.</p> <ul style="list-style-type: none"> <li>■ Verify that packets with small fragment offsets are recorded in the router’s system logging facility.</li> <li>■ Use the <code>show route summary</code> command to verify that the routing table does not contain any entries with a protocol other than Direct, Local, BGP, or Static.</li> </ul> |

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



## Chapter 18

# 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 121.

**Table 121: 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 372
- Configuring CoS with DiffServ with a Configuration Editor on page 372
- Verifying a DiffServ Configuration on page 402

## 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 291.
- Determine whether the Services Router needs to support different traffic streams, such as voice or video. If so, CoS with DiffServ helps to make sure this traffic receives more than basic best-effort packet delivery service.
- Determine whether the Services Router is directly attached to any applications that send DiffServ packets. If no sources are enabled for DiffServ, you must configure and apply rewrite rules on the interfaces to the sources.
- Determine whether the Services Router must support DiffServ assured forwarding (AF) classes. Assured forwarding usually requires random early detection (RED) drop profiles to be configured and applied.
- Determine whether the Services Router must support DiffServ expedited forwarding (EF) classes with a policer. Policers require you to apply a burst size and bandwidth limit to the traffic flow, and set a consequence for packets that exceed these limits—usually a high loss priority, so that packets exceeding the policer limits are discarded first.

## Configuring CoS with DiffServ with a Configuration Editor

---

To configure the Services Router as a node in a network supporting CoS with DiffServ, you must perform the following tasks marked *(Required)*. For information about using the J-Web and CLI configuration editors, see “Using J-series Configuration Tools” on page 3.

- Configuring a Policer for a Firewall Filter (Required) on page 373
- Configuring and Applying a Firewall Filter for a Multifield Classifier (Required) on page 374
- Assigning Forwarding Classes to Output Queues (Required) on page 378
- Configuring and Applying Rewrite Rules (Required) on page 379
- Configuring and Applying Behavior Aggregate Classifiers (Required) on page 384
- Configuring RED Drop Profiles for Assured Forwarding Congestion Control (Required) on page 388
- Configuring Schedulers (Optional) on page 390
- Configuring and Applying Scheduler Maps (Optional) on page 394

- Configuring and Applying Virtual Channels (Optional) on page 397
- Configuring and Applying Adaptive Shaping (Optional) on page 401

**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 331 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 122.
3. Go on to “Configuring and Applying a Firewall Filter for a Multifield Classifier (Required)” on page 374.

**Table 122: 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.                 | <ol style="list-style-type: none"><li>1. Click <b>Add new entry</b> next to Policer.</li><li>2. In the Policer name box, type a name for the EF policer—for example, <code>ef-policer</code>.</li></ol> | Enter<br><br>edit policer ef-policer                                    |

**Table 122: 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 123.

**Table 123: 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 373. | 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 331 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 124.
3. Go on to “Assigning Forwarding Classes to Output Queues (Required)” on page 378.

**Table 124: 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 124: 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 124: 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 373.)</li> <li>4. Click <b>OK</b> twice.</li> </ol> | <p>From the top of the configuration hierarchy, enter</p> <pre>edit firewall filter mf-classifier term expedited-forwarding</pre> <pre>set then forwarding-class ef-class</pre> <pre>set then policer ed-policer</pre> |
| 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>                                                                                                                                                                                                                                                                                                   | <pre>Enter</pre> <pre>edit term network-control</pre>                                                                                                                                                                  |
| 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>                                                                                                                                                                              | <pre>Enter</pre> <pre>set from traffic-class net-control</pre>                                                                                                                                                         |
| 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> <pre>edit firewall filter mf-classifier term network-control</pre> <pre>set then forwarding-class nc-class</pre>                                             |
| 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>                                                                                                                                                                                                                                                                                                      | <pre>Enter</pre> <pre>edit term best-effort-data</pre>                                                                                                                                                                 |
| 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> <pre>set then forwarding-class be-class</pre>                                                                                                                |

**Table 124: 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 125.

**Table 125: 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 374.

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 126.
- Go on to “Configuring and Applying Rewrite Rules (Required)” on page 379.



**Table 126: 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>1. Click <b>Configure</b> next to Forwarding classes.</li> <li>2. Click <b>Add new entry</b> next to Queue.</li> <li>3. In the Queue num box, type <b>0</b>.</li> <li>4. In the Class name box, type the previously configured name of the best-effort class—<b>be-class</b>.</li> <li>5. 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>1. Click <b>Add new entry</b> next to Queue.</li> <li>2. In the Queue num box, type <b>1</b>.</li> <li>3. In the Class name box, type the previously configured name of the expedited forwarding class—<b>ef-class</b>.</li> <li>4. 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>1. Click <b>Add new entry</b> next to Queue.</li> <li>2. In the Queue num box, type <b>2</b>.</li> <li>3. In the Class name box, type the previously configured name of the assured forwarding class—<b>af-class</b>.</li> <li>4. 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>1. Click <b>Add new entry</b> next to Queue.</li> <li>2. In the Queue num box, type <b>3</b>.</li> <li>3. In the Class name box, type the previously configured name of the expedited forwarding class—<b>nc-class</b>.</li> <li>4. 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 127.

**Table 127: 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<br>High-priority code point: 000001 |
| ef-class                       | Expedited forwarding traffic | Low-priority code point: 101110<br>High-priority code point: 101111 |
| af-class                       | Assured forwarding traffic   | Low-priority code point: 001010<br>High-priority code point: 001100 |
| nc-class                       | Network control traffic      | Low-priority code point: 110000<br>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 128.
3. If you are finished configuring the network, commit the configuration.
4. Go on to “Configuring and Applying Behavior Aggregate Classifiers (Required)” on page 384.

**Table 128: 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>edit class-of-service |

**Table 128: Configuring and Applying Rewrite Rules (Continued)**

| <b>Task</b>                                           | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | <b>CLI Configuration Editor</b>                                                                                                                                                                     |
|-------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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, <b>rewrite-dscps</b>.</li> </ol>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   | <p>Enter</p> <p><code>edit rewrite-rules dscp rewrite-dscps</code></p>                                                                                                                              |
| 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, <b>000000</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 best-effort traffic—for example, <b>000001</b>.</li> <li>10. Click <b>OK</b> twice.</li> </ol> | <p>Enter</p> <p><code>set forwarding-class be-class<br/>loss-priority low code points 000000</code></p> <p><code>set forwarding-class be-class<br/>loss-priority high code points 000001</code></p> |

**Table 128: 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 128: 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 128: 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 129, to incoming packets in the four forwarding classes.

**Table 129: 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 130.
3. If you are finished configuring the network, commit the configuration.
4. Go on to “Configuring RED Drop Profiles for Assured Forwarding Congestion Control (Required)” on page 388.

**Table 130: 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 130: 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, 00001.</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, 101111.</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 130: 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 131.

**Table 131: 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 132.
3. Go on to one of the following tasks:
  - “Configuring Schedulers (Optional)” on page 390
  - “Verifying a DiffServ Configuration” on page 402

**Table 132: Configuring RED Drop Profiles for Assured Forwarding Congestion Control**

| 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 132: 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 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> | <p>Enter</p> <p>edit drop-profiles af-normal interpolate</p> <p>set drop-probability 0</p> <p>set drop-probability 100</p> |
| 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>                                                                                                                                                                                                                                                                                             | <p>Enter</p> <p>set fill-level 95</p> <p>set fill-level 100</p>                                                            |

**Table 132: 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 133.

**Table 133: 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 134.
3. Go on to “Configuring and Applying Scheduler Maps (Optional)” on page 394.

**Table 134: 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 134: 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 134: 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 134: 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 374 to the schedulers configured in “Configuring Schedulers (Optional)” on page 390, as shown in Table 135.



**Table 135: 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 136.
3. If you are finished configuring the network, commit the configuration.
4. To check the configuration, see “Verifying a DiffServ Configuration” on page 402.

**Table 136: 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                               |
| 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> | Enter<br><br>set forwarding-class be-class scheduler be-scheduler               |

**Table 136: Configuring Scheduler Maps (Continued)**

| Task                                                   | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                            | CLI Configuration Editor                                                       |
|--------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------|
| Configure an expedited forwarding class and scheduler. | <ol style="list-style-type: none"> <li>Click <b>Add new entry</b> next to Forwarding class.</li> <li>In the Class name box, type the name of the previously configured expedited forwarding class—<b>ef-class</b>.</li> <li>In the Scheduler box, type the name of the previously configured expedited forwarding scheduler—<b>ef-scheduler</b>.</li> <li>Click <b>OK</b>.</li> </ol> | <p>Enter</p> <pre>set forwarding-class ef-class scheduler ef-scheduler</pre>   |
| Configure an assured forwarding class and scheduler.   | <ol style="list-style-type: none"> <li>Click <b>Add new entry</b> next to Forwarding class.</li> <li>In the Class name box, type the name of the previously configured assured forwarding class—<b>af-class</b>.</li> <li>In the Scheduler box, type the name of the previously configured assured forwarding scheduler—<b>af-scheduler</b>.</li> <li>Click <b>OK</b>.</li> </ol>     | <p>Enter</p> <pre>set forwarding-class af-class scheduler af-scheduler</pre>   |
| Configure a network control class and scheduler.       | <ol style="list-style-type: none"> <li>Click <b>Add new entry</b> next to Forwarding class.</li> <li>In the Class name box, type the name of the previously configured network control class—<b>nc-class</b>.</li> <li>In the Scheduler box, type the name of the previously configured network control scheduler—<b>nc-scheduler</b>.</li> <li>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>Click <b>Add new entry</b> next to Interfaces.</li> <li>In the Interface name box, type the name of the interface—for example, <b>fe-0/0/0</b>.</li> <li>In the Scheduler map box, type the name of the previously configured scheduler map—<b>diffserv-cos-map</b>.</li> <li>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 137.
- 3. If you are finished configuring the network, commit the configuration.

Table 137: 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 <code>branch1-vc</code> , <code>branch2-vc</code> , <code>branch3-vc</code> , and the default virtual channel. You must specify a default virtual channel. | <div>1. Click <b>Add new entry</b> next to Virtual channels.</div> <div>2. In the Channel name box, type the name of the virtual channel—for example, <code>branch1-vc</code>.</div> <div>3. Click <b>OK</b>.</div> <div>4. Create additional virtual channels for <code>branch2-vc</code>, <code>branch3-vc</code>, and <code>default-vc</code>.</div> | <div>Enter</div> <div><b>set virtual-channels branch1-vc</b></div> <div>Repeat this statement for <code>branch2-vc</code>, <code>branch3-vc</code>, and <code>default-vc</code>.</div> |

**Table 137: Configuring and Applying Virtual Channels (Continued)**

| <b>Task</b>                                                                                                                                                         | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | <b>CLI Configuration Editor</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/> <pre>set virtual-channel-groups wan-vc-group branch1-vc scheduler-map bestscheduler</pre> </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/> <pre>set virtual-channel-groups wan-vc-group default-vc default</pre> </li> </ol> |
| 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>                                                                                                                                  |

**Table 137: Configuring and Applying Virtual Channels (Continued)**

| <b>Task</b>                                                          | <b>J-Web Configuration Editor</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                | <b>CLI Configuration Editor</b>                                                           |
|----------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| 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—t3-1/0/0.</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—0.</li> <li>5. In the Virtual channel group box, type the name of the previously configured virtual channel group—wan-vc-group.</li> <li>6. Click <b>OK</b>.</li> </ol> | <p>Enter</p> <pre>set interfaces t3-1/0/0 unit 0 virtual-channel-group wan-vc-group</pre> |

**Table 137: 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 138.
- 3. If you are finished configuring the network, commit the configuration.

Table 138: 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 138: Configuring and Applying an Adaptive Shaper (Continued)**

| Task                                                                   | J-Web Configuration Editor                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | CLI Configuration Editor                                                             |
|------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|
| Define the adaptive shaper name and maximum transmit rate.             | <ol style="list-style-type: none"> <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 drop-down 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

**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**

```
user@host> show sap listen
```



```
Group Address Port
224.2.127.254 9875
```

**What It Means** The output shows a list of the group addresses and ports that SAP and SDP listen on. Verify the following information:

- Each group address configured, especially the default 224.2.127.254, is listed.
- Each port configured, especially the default 9875, is listed.

For more information about `show sap listen`, see the *JUNOS Protocols, Class of Service, and System Basics Command Reference*.

**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**

```
user@host> show class-of-service adaptive-shaper

Adaptive shaper: fr-shaper, Index: 35320
 Trigger type Shaping rate
 BECN 64000 bps

user@host> show class-of-service interface t1-0/0/2

Physical interface: t1-0/0/2, Index: 137
Queues supported: 8, Queues in use: 4
Scheduler map: <default>, Index: 2

Logical interface: t1-0/0/2.0, Index: 69
 Object Name Type Index
 Adaptive-shaper fr-shaper 35320
 Classifier ipprec-compatibility ip 11
```

**What It Means** Verify the following information:

- The trigger type and shaping rate are consistent with the configured adaptive shaper.
- The adaptive shaper applied to the logical interface is displayed under Name.



## **Part 7**

# **Index**



# Index

## Symbols

|                                               |    |
|-----------------------------------------------|----|
| [ ], in configuration statements .....        | xx |
| { }, in configuration statements .....        | xx |
| ( ), in syntax descriptions .....             | xx |
| < >, in syntax descriptions .....             | xx |
| (pipe), in syntax descriptions .....          | xx |
| #, comments in configuration statements ..... | xx |

## A

|                                                      |     |
|------------------------------------------------------|-----|
| ABRs <i>See</i> area border routers                  |     |
| access concentrator                                  |     |
| as a PPPoE server .....                              | 78  |
| naming for PPPoE .....                               | 85  |
| action modifiers .....                               | 306 |
| actions                                              |     |
| default, routing policy .....                        | 298 |
| final, routing policy .....                          | 298 |
| NAT .....                                            | 301 |
| route list match types .....                         | 320 |
| routing policy .....                                 | 296 |
| routing policy, summary of .....                     | 297 |
| stateful firewall filters .....                      | 301 |
| stateless firewall filters .....                     | 306 |
| active routes, versus passive routes .....           | 129 |
| adaptive shaping, applying CoS rules to logical      |     |
| interfaces .....                                     | 401 |
| Add button .....                                     | 8   |
| Add new entry link .....                             | 10  |
| address match conditions .....                       | 305 |
| address translation <i>See</i> NAT                   |     |
| addresses                                            |     |
| BGP external peer address (configuration             |     |
| editor) .....                                        | 183 |
| BGP internal peer address (configuration             |     |
| editor) .....                                        | 185 |
| BGP local address (Quick Configuration) .....        | 180 |
| BGP peer address (Quick Configuration) .....         | 180 |
| multicast ranges .....                               | 273 |
| translating <i>See</i> NAT                           |     |
| administrative groups, for MPLS path selection ..... | 207 |
| administrative scoping .....                         | 275 |
| ADSL ports <i>See</i> ATM-for-ADSL interfaces        |     |
| advertisements <i>See</i> LSAs; route advertisements |     |
| AF <i>See</i> DiffServ, assured forwarding           |     |

|                                                             |     |
|-------------------------------------------------------------|-----|
| aggregation, route .....                                    | 105 |
| alternate mark inversion <i>See</i> AMI                     |     |
| AMI (alternate mark inversion)                              |     |
| E1 .....                                                    | 51  |
| T1 .....                                                    | 56  |
| Annex A PIMs                                                |     |
| ATM-for-ADSL interfaces .....                               | 66  |
| operating modes .....                                       | 68  |
| <i>See also</i> ATM-for-ADSL interfaces                     |     |
| Annex B PIMs                                                |     |
| ATM-for-ADSL interfaces .....                               | 66  |
| operating modes .....                                       | 68  |
| <i>See also</i> ATM-for-ADSL interfaces                     |     |
| ANSI DMT operating mode .....                               | 68  |
| Apply button .....                                          | 8   |
| area border routers                                         |     |
| adding interfaces .....                                     | 165 |
| area ID (configuration editor) .....                        | 165 |
| backbone area <i>See</i> backbone area                      |     |
| backbone area interface .....                               | 164 |
| description .....                                           | 113 |
| areas <i>See</i> area border routers; backbone area; NSSAs; |     |
| stub areas                                                  |     |
| AS path                                                     |     |
| description .....                                           | 121 |
| forcing by MED .....                                        | 122 |
| prepending .....                                            | 325 |
| role in route selection .....                               | 119 |
| ASs (autonomous systems)                                    |     |
| area border routers .....                                   | 113 |
| AS number (configuration editor) .....                      | 183 |
| AS number (Quick Configuration) .....                       | 180 |
| AS number, in VPNs .....                                    | 236 |
| breaking into confederations .....                          | 125 |
| description .....                                           | 102 |
| group AS number (configuration editor) .....                | 183 |
| individual AS number (configuration editor) .....           | 183 |
| LSPs through .....                                          | 200 |
| sample BGP confederation .....                              | 189 |
| stub areas <i>See</i> stub areas                            |     |
| sub-AS number .....                                         | 189 |
| assured forwarding .....                                    | 388 |
| asymmetrical digital subscriber line (ADSL)                 |     |
| <i>See</i> ATM-for-ADSL interfaces                          |     |

|                                                |                                    |
|------------------------------------------------|------------------------------------|
| Asynchronous Transfer Mode (ATM) interface     |                                    |
| <i>See</i> ATM-for-ADSL interfaces             |                                    |
| at-0/0/0                                       | <i>See</i> ATM-for-ADSL interfaces |
| ATM interface                                  | <i>See</i> ATM-for-ADSL interfaces |
| ATM NLPID encapsulation                        | 69                                 |
| ATM PPP over AAL5 LLC encapsulation            | 69                                 |
| ATM PVC encapsulation                          | 68                                 |
| ATM SNAP encapsulation                         | 69                                 |
| ATM VC multiplex encapsulation                 | 69                                 |
| ATM-for-ADSL interfaces                        | 66                                 |
| adding                                         | 66                                 |
| CHAP for PPPoE                                 | 87                                 |
| encapsulation types, logical                   | 69                                 |
| encapsulation types, physical                  | 68                                 |
| logical properties                             | 68                                 |
| operating modes                                | 68                                 |
| physical properties                            | 67                                 |
| PPPoE configuration                            | 85                                 |
| PPPoE encapsulation                            | 83                                 |
| PPPoE session on                               | 79                                 |
| statistics                                     | 75                                 |
| VCI                                            | 69                                 |
| verifying                                      | 72                                 |
| verifying a PPPoE configuration                | 89                                 |
| <i>See also</i> PPPoE; PPPoE over ATM-for-ADSL |                                    |
| authentication                                 |                                    |
| OSPF, MD5                                      | 170                                |
| OSPF, plain-text passwords                     | 170                                |
| RIPv2, MD5                                     | 150                                |
| RIPv2, plain-text passwords                    | 149                                |
| auto operating mode                            | 68                                 |
| Auto-RP                                        | 276                                |

## B

|                                                     |                               |
|-----------------------------------------------------|-------------------------------|
| BA classifiers                                      | <i>See</i> classifiers        |
| backbone area                                       |                               |
| area ID (configuration editor)                      | 162                           |
| area ID (Quick Configuration)                       | 158                           |
| area type (Quick Configuration)                     | 159                           |
| configuring                                         | 161                           |
| description                                         | 114                           |
| interface                                           | 164                           |
| bandwidth, for RSVP-signaled LSPs                   | 220                           |
| behavior aggregate classifiers                      | <i>See</i> classifiers        |
| best-effort service                                 | 307                           |
| BGP (Border Gateway Protocol)                       |                               |
| AS number (Quick Configuration)                     | 180                           |
| <i>See also</i> ASs (autonomous systems), AS number |                               |
| AS path                                             | 121                           |
| <i>See also</i> AS path                             |                               |
| confederations                                      | <i>See</i> BGP confederations |
| enabling (Quick Configuration)                      | 180                           |
| external                                            | 118                           |
| <i>See also</i> EBGp                                |                               |

|                                                             |                                    |
|-------------------------------------------------------------|------------------------------------|
| external group type (configuration editor)                  | 183                                |
| external neighbor (peer) address (configuration editor)     | 183                                |
| full mesh requirement                                       | 119, 178                           |
| injecting OSPF routes into BGP                              | 322                                |
| internal                                                    | 118                                |
| <i>See also</i> IBGP                                        |                                    |
| internal group type (configuration editor)                  | 185                                |
| internal neighbor (peer) address (configuration editor)     | 185                                |
| local address (Quick Configuration)                         | 180                                |
| local preference                                            | 120                                |
| MED metric                                                  | 122                                |
| origin value                                                | 121                                |
| overview                                                    | 116, 177                           |
| peer address (Quick Configuration)                          | 180                                |
| peer AS number (Quick Configuration)                        | 180                                |
| peering sessions                                            | <i>See</i> BGP peers; BGP sessions |
| point-to-point internal peer session (configuration editor) | 184                                |
| point-to-point peer session (configuration editor)          | 181                                |
| policy to make routes less preferable                       | 325                                |
| Quick Configuration                                         | 179                                |
| requirements                                                | 179                                |
| route reflectors                                            | <i>See</i> BGP route reflectors    |
| route selection process                                     | 119                                |
| <i>See also</i> route selection                             |                                    |
| route-flap damping                                          | 327                                |
| router ID (Quick Configuration)                             | 180                                |
| routing policy (configuration editor)                       | 185                                |
| <i>See also</i> routing policies                            |                                    |
| sample BGP peer network                                     | 182                                |
| sample confederation                                        | 189                                |
| sample full mesh                                            | 184                                |
| sample route reflector                                      | 186                                |
| scaling techniques                                          | 122                                |
| session establishment                                       | 118                                |
| session maintenance                                         | 118                                |
| verifying BGP configuration                                 | 192                                |
| verifying BGP groups                                        | 191                                |
| verifying BGP peers (neighbors)                             | 190                                |
| verifying peer reachability                                 | 193                                |
| VPNs                                                        | 235                                |
| BGP confederations                                          |                                    |
| confederation members                                       | 190                                |
| confederation number                                        | 189                                |
| creating (configuration editor)                             | 188                                |
| description                                                 | 125, 178                           |
| route-flap damping                                          | 327                                |
| sample network                                              | 189                                |
| sub-AS number                                               | 189                                |
| BGP groups                                                  |                                    |
| cluster identifier (configuration editor)                   | 187                                |
| confederations (configuration editor)                       | 188                                |

external group type (configuration editor)..... 183  
 external, creating (configuration editor)..... 183  
 group AS number (configuration editor)..... 183  
 internal group type (configuration editor)..... 185  
 internal, creating (configuration editor)..... 185  
 internal, creating for a route reflector  
     (configuration editor)..... 187  
 verifying..... 191  
 BGP messages  
     to establish sessions..... 118  
     update, to maintain sessions..... 118  
 BGP page..... 179  
 BGP peers  
     directing traffic by local preference..... 120  
     external (configuration editor)..... 181  
     internal (configuration editor)..... 184  
     internal, sample full mesh..... 184  
     internal, sample route reflector..... 186  
     peer address (Quick Configuration)..... 180  
     peer AS number (Quick Configuration)..... 180  
     point-to-point connections..... 117  
     routing policy (configuration editor)..... 185  
         *See also* routing policies  
     sample peer network..... 182  
     sessions between peers..... 177  
     verifying..... 190, 192  
     verifying reachability..... 193  
 BGP route reflectors  
     cluster (configuration editor)..... 187  
     cluster identifier (configuration editor)..... 187  
     cluster of clusters..... 124  
     creating (configuration editor)..... 185  
     description..... 123, 178  
     group type (configuration editor)..... 187  
     multiple clusters..... 123  
     sample IBGP network..... 186  
 BGP sessions  
     configured at both ends..... 177  
     establishment..... 118  
     maintenance..... 118  
     point-to-point external (configuration editor)..... 181  
     point-to-point internal (configuration editor)..... 184  
     sample peering session..... 117  
     types..... 178  
 bit-field logical operators, stateless firewall filters..... 306  
 bit-field match conditions..... 305  
 bit-field synonym match conditions..... 305  
 bootstrap router..... 276  
 Border Gateway Protocol *See* BGP  
 braces, in configuration statements.....xx  
 brackets  
     angle, in syntax descriptions.....xx  
     square, in configuration statements.....xx  
 branches..... 272  
     *See also* multicast

BSR (bootstrap router)..... 276  
 buttons..... 11  
     Add (Quick Configuration)..... 8  
     Apply (Quick Configuration)..... 8  
     Cancel (J-Web configuration editor)..... 11  
     Cancel (Quick Configuration)..... 8  
     Commit (J-Web configuration editor)..... 11  
     CONFIG button..... 21, 33  
     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  
     *See also* radio buttons

## C

C-bit parity..... 60  
 cables  
     T1 cable length..... 57  
     T3 cable length..... 60  
 Cancel button  
     J-Web configuration editor..... 11  
     Quick Configuration..... 8  
 canceling a commit..... 31–32  
 CE (customer edge) routers..... 228  
     description..... 209  
     VPN task overview..... 230  
     VPN topology..... 228  
     *See also* VPNs  
 Challenge Handshake Authentication Protocol  
     *See* CHAP  
 channel number, in interface name..... 47  
 CHAP (Challenge Handshake Authentication Protocol)  
     E1 local identity..... 51  
     enabling for PPPoE..... 87  
     enabling on E1..... 50  
     enabling on serial interfaces..... 62  
     enabling on T1..... 55  
     enabling on T3..... 59  
     serial interface local identity..... 62  
     T1 local identity..... 55  
     T3 local identity..... 59  
 CHAP secret *See* CHAP, local identity  
 checksum  
     E1 frame..... 51  
     T1 frame..... 57  
     T3 frame..... 60  
 circuit *See* Layer 2 circuits  
 Cisco NLPID encapsulation..... 69  
 class of service *See* CoS  
 classifiers  
     applying BA classifiers..... 384–385  
     default BA classifiers..... 313  
     description..... 310  
     sample BA classification..... 314

|                                                       |     |                                                          |        |
|-------------------------------------------------------|-----|----------------------------------------------------------|--------|
| sample BA classifier assignments.....                 | 385 | rescue configuration (J-Web) .....                       | 21     |
| sample, for firewall filter .....                     | 375 | scheduling (CLI configuration editor) .....              | 31     |
| clear system commit command .....                     | 32  | storage location .....                                   | 5      |
| CLI configuration editor .....                        |     | summaries .....                                          | 17     |
| activating a configuration .....                      | 31  | verifying (CLI configuration editor) .....               | 30     |
| BGP .....                                             | 181 | viewing previous (CLI configuration editor) .....        | 33     |
| command summary .....                                 | 5   | confederations <i>See</i> BGP confederations             |        |
| committing files .....                                | 30  | CONFIG button .....                                      |        |
| confirming a configuration .....                      | 31  | 15-second caution .....                                  | 21, 33 |
| exiting .....                                         | 22  | configuration .....                                      |        |
| IPSec tunnels .....                                   | 254 | activating (CLI configuration editor) .....              | 31     |
| managing files .....                                  | 34  | adding a statement (CLI configuration editor) .....      | 26     |
| modifying a configuration .....                       | 25  | basic .....                                              | 7      |
| MPLS traffic engineering .....                        | 215 | changing part of a file (CLI configuration editor) ..... | 35     |
| network interfaces .....                              | 64  | CLI commands .....                                       | 5      |
| OSPF .....                                            | 160 | CLI configuration mode .....                             | 22     |
| PPPoE .....                                           | 82  | committed .....                                          | 4      |
| PPPoE over ATM-for-ADSL .....                         | 82  | committing (CLI configuration editor) .....              | 30     |
| RIP .....                                             | 143 | committing (J-Web) .....                                 | 12     |
| saving files .....                                    | 37  | committing as a text file, with caution (J-Web) .....    | 13     |
| starting .....                                        | 22  | confirming (CLI configuration editor) .....              | 31     |
| static routes .....                                   | 132 | copying a statement .....                                | 27     |
| using show commands with .....                        | 33  | deactivating a statement .....                           | 29     |
| verifying a configuration .....                       | 30  | deleting a statement .....                               | 26     |
| VPNs .....                                            | 230 | discarding changes (J-Web) .....                         | 11     |
| clickable configuration .....                         | 8   | downloading (J-Web) .....                                | 20     |
| committing .....                                      | 12  | editing (J-Web) .....                                    | 8      |
| discarding changes .....                              | 11  | editing as a text file, with caution (J-Web) .....       | 13     |
| viewing and editing .....                             | 8   | history .....                                            | 16     |
| <i>See also</i> J-Web configuration editor            |     | <i>See also</i> configuration history                    |        |
| clock rate, serial interface .....                    | 63  | inserting an identifier .....                            | 28     |
| clocking .....                                        |     | J-Web options .....                                      | 5      |
| E1 .....                                              | 51  | loading new (CLI configuration editor) .....             | 34     |
| serial interface .....                                | 63  | loading previous (CLI configuration editor) .....        | 32     |
| serial interface, inverting the transmit clock .....  | 63  | loading previous (J-Web) .....                           | 21     |
| T1 .....                                              | 56  | locked, with the configure exclusive command .....       | 23     |
| T3 .....                                              | 60  | managing files (CLI configuration editor) .....          | 34     |
| clusters <i>See</i> BGP route reflectors              |     | managing files (J-Web) .....                             | 15     |
| coloring, link, for MPLS path selection .....         | 207 | merging (CLI configuration editor) .....                 | 35     |
| comments, in configuration statements .....           | xx  | modifying (CLI configuration editor) .....               | 25     |
| commit and-quit command .....                         | 31  | modifying a statement (CLI configuration editor) .....   | 26     |
| commit at command .....                               | 31  | overriding (CLI configuration editor) .....              | 35     |
| Commit button .....                                   | 11  | renaming an identifier .....                             | 27     |
| commit check command .....                            | 30  | replacing configuration statements (CLI                  |        |
| commit command .....                                  | 30  | configuration editor) .....                              | 35     |
| commit confirmed command .....                        | 31  | requirements .....                                       | 7      |
| committed configuration .....                         |     | rescuing (CLI configuration editor) .....                | 33     |
| activating (CLI configuration editor) .....           | 31  | rescuing (J-Web) .....                                   | 21     |
| canceling a commit (CLI configuration editor) .....   | 32  | rollback (CLI configuration editor) .....                | 32     |
| comparing two configurations .....                    | 18  | rollback (J-Web) .....                                   | 21     |
| confirming (CLI configuration editor) .....           | 31  | saving (CLI configuration editor) .....                  | 37     |
| description .....                                     | 4   | uploading (J-Web) .....                                  | 14     |
| methods .....                                         | 17  | users-editors, viewing .....                             | 18     |
| replacing (CLI configuration editor) .....            | 32  | verifying (CLI configuration editor) .....               | 30     |
| rescue configuration (CLI configuration editor) ..... | 33  | 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 ..... 33

configuration text

- editing and committing, with caution ..... 13
- viewing ..... 12

configuration tools ..... 3

- See also* CLI configuration editor; configuration; configuration history; J-Web configuration editor; Quick Configuration

configure command ..... 23

configure exclusive command ..... 23

Configure link ..... 10

configure private command ..... 23

confirming a configuration ..... 31

congestion control, with DiffServ assured forwarding ..... 388

connectivity

- bidirectional (BGP) ..... 116
- bidirectional (OSPF) ..... 111
- unidirectional (RIP) ..... 110

Constrained Shortest Path First *See* CSPF

conventions

- for interface names ..... 45
- how to use this guide ..... xviii
- notice icons ..... xix
- text and syntax ..... xix

copy command ..... 27

CoS (class of service)

- adaptive shaping for rules ..... 401
- assigning forwarding classes to output queues ..... 378
- BA classifiers ..... 384
- configuration tasks ..... 372
- default BA classifiers ..... 313
- default forwarding class queue assignments ..... 311
- default scheduler settings ..... 312
- DiffServ benefits ..... 307
- See also* DiffServ
- DSCP rewrites ..... 314
- DSCPs ..... 308
- See also* DSCPs
- firewall filter for a multifeild classifier ..... 374
- JUNOS components ..... 310
- JUNOS implementation ..... 309
- policer for firewall filter ..... 373

- preparation ..... 372
- RED drop profiles ..... 388
- rewrite rules ..... 379
- sample BA classification ..... 314
- scheduler maps ..... 394
- schedulers ..... 390
- uses ..... 371
- verifying adaptive shaper configuration ..... 403
- verifying multicast session announcements ..... 402
- virtual channels for rules ..... 397

cost, of a network path *See* path cost metrics

CPE device, Services Router as, with PPPoE ..... 77

- See also* PPPoE

CSPF (Constrained Shortest Path First)

- constraints ..... 207
- disabling ..... 220
- link coloring ..... 207
- rules ..... 207

CSPF algorithm *See* CSPF

curly braces, in configuration statements ..... xx

customer edge routers *See* CE routers

customer premises equipment (CPE) device, Services Router as, with PPPoE ..... 77

- See also* PPPoE

customer support ..... xxii

- contacting JTAC ..... xxii

## D

data inversion

- E1 ..... 51
- T1 ..... 56

Database Information page ..... 16

deactivate command ..... 29

deactivating configuration statements or identifiers ..... 29

default gateway, static routing ..... 131

defaults

- BA classifiers ..... 313
- CoS forwarding class assignments ..... 312
- junos-algs-outbound group, stateful firewall filters ..... 300
- routing policy actions ..... 298
- setting for static routes ..... 136

Delete button ..... 8

delete command ..... 26

Delete Configuration Below This Point radio button ..... 11

deleting

- current rescue configuration (CLI configuration editor) ..... 33
- current rescue configuration (J-Web) ..... 22
- network interfaces ..... 69

denial-of-service attacks, preventing ..... 347

dense routing mode, caution for use ..... 274

- See also* multicast routing modes

designated router, OSPF

- controlling election ..... 170

|                                                      |          |                                                    |      |
|------------------------------------------------------|----------|----------------------------------------------------|------|
| description .....                                    | 112      | default BA classifiers.....                        | 313  |
| Deutsche Telekom UR-2 operating mode.....            | 68       | default forwarding class queue assignments .....   | 311  |
| diagnosis                                            |          | default scheduler settings .....                   | 312  |
| displaying firewall filter configurations .....      | 359      | DSCP rewrites.....                                 | 314  |
| displaying firewall filter statistics .....          | 366      | firewall filter for a multifield classifier .....  | 374  |
| displaying static routes in the routing table .....  | 137      | forwarding service classes .....                   | 308  |
| LDP neighbors .....                                  | 220      | interoperability .....                             | 308  |
| LDP sessions .....                                   | 221      | JUNOS implementation .....                         | 309  |
| LDP-signaled LSP .....                               | 222      | policer for firewall filter .....                  | 373  |
| RSVP neighbors.....                                  | 223      | preparation.....                                   | 372  |
| RSVP sessions .....                                  | 224      | RED drop profiles.....                             | 388  |
| RSVP-signaled LSP.....                               | 224      | rewrite rules.....                                 | 379  |
| traffic forwarding over LDP-signaled LSPs.....       | 222      | sample BA classification.....                      | 314  |
| verifying adaptive shaper configuration.....         | 403      | scheduler maps.....                                | 394  |
| verifying BGP configuration .....                    | 192      | schedulers.....                                    | 390  |
| verifying BGP groups .....                           | 191      | uses.....                                          | 371  |
| verifying BGP peer reachability .....                | 193      | virtual channels for rules .....                   | 397  |
| verifying BGP peers (neighbors) .....                | 190, 402 | Discard All Changes radio button .....             | 11   |
| verifying firewall filter actions .....              | 367      | Discard button .....                               | 11   |
| verifying firewall filter DoS protection.....        | 368      | Discard Changes Below This Point radio button..... | 11   |
| verifying firewall filter flood protection.....      | 368      | discard rule                                       |      |
| verifying firewall filter handles fragments.....     | 369      | firewall filters .....                             | 299  |
| verifying firewall filters with packet logs .....    | 365      | stateful firewall filters.....                     | 299  |
| verifying IPsec tunnel operation .....               | 265      | stateless firewall filters .....                   | 302  |
| verifying MPLS traffic engineering.....              | 220      | discarding configuration changes.....              | 11   |
| verifying multicast IGMP versions .....              | 286      | discovery packets, PPPoE .....                     | 80   |
| verifying multicast SAP and SDP configuration ..     | 286      | Distance Vector Multicast Routing Protocol.....    | 275  |
| verifying OSPF host reachability .....               | 174      | distance-vector routing protocols .....            | 107  |
| verifying OSPF neighbors .....                       | 172      | <i>See also</i> RIP                                |      |
| verifying OSPF routes .....                          | 173      | documentation set                                  |      |
| verifying OSPF-enabled interfaces.....               | 171      | comments on .....                                  | xxii |
| verifying PIM mode and interface configuration ..    | 287      | DoS (denial-of-service) attacks, preventing .....  | 347  |
| verifying PIM RPF routing table.....                 | 288      | downloading, configuration files (J-Web) .....     | 20   |
| verifying PIM RPs.....                               | 287      | downstream interfaces .....                        | 272  |
| verifying PPPoE interfaces .....                     | 90       | <i>See also</i> multicast                          |      |
| verifying PPPoE over ATM-for-ADSL                    |          | DS1 ports <i>See</i> T1 ports                      |      |
| configuration .....                                  | 89       | DS3 ports <i>See</i> T3 ports                      |      |
| verifying PPPoE sessions.....                        | 91       | DSCPs (DiffServ code points)                       |      |
| verifying PPPoE statistics.....                      | 92       | corresponding forwarding service classes .....     | 308  |
| verifying PPPoE version information .....            | 92       | default forwarding class queue assignments .....   | 311  |
| verifying RIP host reachability .....                | 152      | description .....                                  | 308  |
| verifying RIP message exchange.....                  | 151      | replacing with rewrite rules .....                 | 380  |
| verifying RIP-enabled interfaces .....               | 151      | rewrites .....                                     | 314  |
| verifying stateful firewall filters .....            | 364      | sample BA classification.....                      | 314  |
| verifying VPN connectivity .....                     | 248      | DSL access multiplexer (DSLAM) connection          |      |
| Differentiated Services <i>See</i> DiffServ          |          | <i>See</i> DSLAM connection                        |      |
| DiffServ (Differentiated Services)                   |          | DSLAM connection                                   |      |
| assigning forwarding classes to output queues... 378 |          | ATM-for-ADSL interface for.....                    | 66   |
| assured forwarding .....                             | 388      | PPPoE over ATM-for-ADSL topology .....             | 79   |
| BA classifiers .....                                 | 384      | DVMRP (Distance Vector Multicast Routing           |      |
| benefits for CoS.....                                | 307      | Protocol) .....                                    | 275  |
| code points.....                                     | 308      | dynamic LSPs.....                                  | 203  |
| <i>See also</i> DSCPs                                |          | dynamic routing .....                              | 104  |
| configuration tasks .....                            | 372      |                                                    |      |

**E**

## E1 ports

|                                  |    |
|----------------------------------|----|
| CHAP .....                       | 50 |
| clocking .....                   | 51 |
| configuring .....                | 49 |
| data inversion .....             | 51 |
| encapsulation type .....         | 50 |
| fractional, channel number ..... | 47 |
| frame checksum .....             | 51 |
| framing .....                    | 51 |
| logical interfaces .....         | 50 |
| MTU .....                        | 51 |
| time slots .....                 | 51 |

## EBGP (external BGP)

|                          |     |
|--------------------------|-----|
| description .....        | 118 |
| route-flap damping ..... | 327 |
| sample network .....     | 184 |

## edit command

|       |    |
|-------|----|
| ..... | 24 |
|-------|----|

## Edit Configuration page

|       |   |
|-------|---|
| ..... | 9 |
|-------|---|

## Edit Configuration Text page

|       |    |
|-------|----|
| ..... | 14 |
|-------|----|

## Edit link

|       |    |
|-------|----|
| ..... | 10 |
|-------|----|

## EGPs (exterior gateway protocols)

|       |     |
|-------|-----|
| ..... | 102 |
|-------|-----|

egress router *See* LSPs; outbound router

## encapsulation type

|                                        |    |
|----------------------------------------|----|
| .....                                  | 50 |
| ATM-for-ADSL logical interfaces .....  | 69 |
| ATM-for-ADSL physical interfaces ..... | 68 |
| E1 .....                               | 50 |
| PPPoE .....                            | 77 |
| PPPoE for Ethernet .....               | 83 |
| PPPoE, over ATM for ADSL .....         | 83 |
| serial interfaces .....                | 62 |
| T1 .....                               | 55 |
| T3 .....                               | 59 |

*See also* packet encapsulation

## EROs (Explicit Route Objects)

|                   |     |
|-------------------|-----|
| loose hops .....  | 206 |
| strict hops ..... | 206 |

## Ethernet over ATM encapsulation

|       |    |
|-------|----|
| ..... | 68 |
|-------|----|

## Ethernet over LLC encapsulation

|       |    |
|-------|----|
| ..... | 69 |
|-------|----|

Ethernet ports *See* Fast Ethernet ports

## ETSI operating mode

|       |    |
|-------|----|
| ..... | 68 |
|-------|----|

## exact route list match type

|       |     |
|-------|-----|
| ..... | 320 |
|-------|-----|

## exit command

|                                               |    |
|-----------------------------------------------|----|
| to leave configuration mode .....             | 23 |
| to navigate the configuration hierarchy ..... | 24 |

## exit configuration-mode command

|       |    |
|-------|----|
| ..... | 23 |
|-------|----|

Explicit Route Objects *See* EROs

## export routing policy, for Layer 2 VPNs

|       |     |
|-------|-----|
| ..... | 245 |
|-------|-----|

## export statement, for routing policies

|       |     |
|-------|-----|
| ..... | 298 |
|-------|-----|

## exterior gateway protocols

|       |     |
|-------|-----|
| ..... | 102 |
|-------|-----|

external BGP *See* EBGp**F**

## Fast Ethernet ports

|                      |    |
|----------------------|----|
| CHAP for PPPoE ..... | 87 |
|----------------------|----|

## configuring

|       |    |
|-------|----|
| ..... | 52 |
|-------|----|

## logical interfaces

|       |    |
|-------|----|
| ..... | 53 |
|-------|----|

## PPPoE configuration

|       |    |
|-------|----|
| ..... | 85 |
|-------|----|

## PPPoE encapsulation

|       |    |
|-------|----|
| ..... | 83 |
|-------|----|

## PPPoE session on

|       |    |
|-------|----|
| ..... | 79 |
|-------|----|

## fe-0/0/0, disabling PIM on

|       |     |
|-------|-----|
| ..... | 283 |
|-------|-----|

## file management

|                                                      |    |
|------------------------------------------------------|----|
| configuration files (CLI configuration editor) ..... | 34 |
|------------------------------------------------------|----|

|                                   |    |
|-----------------------------------|----|
| configuration files (J-Web) ..... | 15 |
|-----------------------------------|----|

## firewall filters

|                                                |     |
|------------------------------------------------|-----|
| applying CoS rules to logical interfaces ..... | 397 |
|------------------------------------------------|-----|

|                                 |     |
|---------------------------------|-----|
| displaying configurations ..... | 359 |
|---------------------------------|-----|

|                             |     |
|-----------------------------|-----|
| displaying statistics ..... | 366 |
|-----------------------------|-----|

|                                          |     |
|------------------------------------------|-----|
| multifield classifier filter terms ..... | 375 |
|------------------------------------------|-----|

|                |     |
|----------------|-----|
| overview ..... | 298 |
|----------------|-----|

|                   |     |
|-------------------|-----|
| policer for ..... | 373 |
|-------------------|-----|

|                               |     |
|-------------------------------|-----|
| sample classifier terms ..... | 375 |
|-------------------------------|-----|

|                                 |     |
|---------------------------------|-----|
| stateful firewall filters ..... | 299 |
|---------------------------------|-----|

*See also* stateful firewall filters

|                                  |     |
|----------------------------------|-----|
| stateless firewall filters ..... | 299 |
|----------------------------------|-----|

*See also* stateless firewall filters

|                           |     |
|---------------------------|-----|
| term number caution ..... | 299 |
|---------------------------|-----|

|                               |     |
|-------------------------------|-----|
| verifying configuration ..... | 359 |
|-------------------------------|-----|

|                                  |     |
|----------------------------------|-----|
| verifying flood protection ..... | 368 |
|----------------------------------|-----|

|                                   |     |
|-----------------------------------|-----|
| verifying fragment handling ..... | 369 |
|-----------------------------------|-----|

|                                |     |
|--------------------------------|-----|
| verifying packet logging ..... | 365 |
|--------------------------------|-----|

## Firewall/NAT application page

|       |     |
|-------|-----|
| ..... | 334 |
|-------|-----|

## Firewall/NAT page

|       |     |
|-------|-----|
| ..... | 333 |
|-------|-----|

|                     |     |
|---------------------|-----|
| field summary ..... | 335 |
|---------------------|-----|

## flap damping

|       |     |
|-------|-----|
| ..... | 327 |
|-------|-----|

## Flexible PIM Concentrator, number in interface

|            |    |
|------------|----|
| name ..... | 46 |
|------------|----|

## flooding, preventing

|       |     |
|-------|-----|
| ..... | 347 |
|-------|-----|

## flow control actions, routing policies

|       |     |
|-------|-----|
| ..... | 297 |
|-------|-----|

## font conventions

|       |     |
|-------|-----|
| ..... | xix |
|-------|-----|

## forwarding classes

|                                  |     |
|----------------------------------|-----|
| assigning to output queues ..... | 379 |
|----------------------------------|-----|

|                                 |     |
|---------------------------------|-----|
| default queue assignments ..... | 311 |
|---------------------------------|-----|

|                   |     |
|-------------------|-----|
| description ..... | 310 |
|-------------------|-----|

|                             |     |
|-----------------------------|-----|
| mapping to schedulers ..... | 395 |
|-----------------------------|-----|

|                                                    |     |
|----------------------------------------------------|-----|
| policy to group source and destination prefixes .. | 324 |
|----------------------------------------------------|-----|

|                                |     |
|--------------------------------|-----|
| sample BA classification ..... | 314 |
|--------------------------------|-----|

|                       |     |
|-----------------------|-----|
| sample mappings ..... | 395 |
|-----------------------|-----|

## forwarding policy options

|       |     |
|-------|-----|
| ..... | 310 |
|-------|-----|

## forwarding states, multicast notation

|       |     |
|-------|-----|
| ..... | 273 |
|-------|-----|

## forwarding table

|                                  |     |
|----------------------------------|-----|
| controlling OSPF routes in ..... | 167 |
|----------------------------------|-----|

|                                    |          |
|------------------------------------|----------|
| controlling static routes in ..... | 128, 135 |
|------------------------------------|----------|

|                   |     |
|-------------------|-----|
| description ..... | 103 |
|-------------------|-----|

|                                  |     |
|----------------------------------|-----|
| MED to determine routes in ..... | 122 |
|----------------------------------|-----|

## FPC (Flexible PIM Concentrator), number in interface

|            |    |
|------------|----|
| name ..... | 46 |
|------------|----|

## framing

|          |    |
|----------|----|
| E1 ..... | 51 |
|----------|----|

|          |    |
|----------|----|
| T1 ..... | 56 |
|----------|----|

|                                                      |     |
|------------------------------------------------------|-----|
| T3.....                                              | 60  |
| from statement, routing policy match conditions..... | 294 |
| full mesh requirement                                |     |
| description .....                                    | 119 |
| fulfilling with confederations.....                  | 125 |
| fulfilling with route reflectors .....               | 123 |
| sample network.....                                  | 184 |

## G

|                                                         |     |
|---------------------------------------------------------|-----|
| *,G notation, for multicast forwarding states.....      | 273 |
| gateway, local and remote, for IPSec service sets ..... | 257 |
| glossary                                                |     |
| configuration.....                                      | 3   |
| CoS .....                                               | 291 |
| firewall filters .....                                  | 291 |
| MPLS .....                                              | 197 |
| multicast .....                                         | 269 |
| network interfaces.....                                 | 41  |
| PPPoE .....                                             | 77  |
| routing .....                                           | 97  |
| routing policies .....                                  | 291 |
| VPNs .....                                              | 197 |
| groups                                                  |     |
| BGP <i>See</i> BGP groups                               |     |
| default junos-algs-outbound group, for stateful         |     |
| firewall filters.....                                   | 300 |
| OSPF areas.....                                         | 162 |
| RIP routers .....                                       | 143 |

## H

|                                          |       |
|------------------------------------------|-------|
| handling packet fragments .....          | 355   |
| hierarchy, configuration .....           | 24    |
| history <i>See</i> configuration history |       |
| hold time, to maintain a session .....   | 118   |
| hop count, maximizing.....               | 108   |
| <i>See also</i> RIP                      |       |
| host reachability                        |       |
| verifying for a RIP network .....        | 152   |
| verifying for an OSPF network.....       | 174   |
| hostname, for PPPoE CHAP .....           | 88    |
| how to use this guide .....              | xviii |

## I

|                                                        |     |
|--------------------------------------------------------|-----|
| IBGP (internal BGP)                                    |     |
| description .....                                      | 118 |
| full mesh (configuration editor) .....                 | 184 |
| full mesh requirement.....                             | 178 |
| sample network.....                                    | 184 |
| sample route reflector .....                           | 186 |
| ICMP (Internet Control Message Protocol), policers.... | 349 |
| identifier link.....                                   | 10  |
| identifiers, configuration                             |     |
| adding or modifying.....                               | 26  |
| deactivating .....                                     | 29  |
| deleting.....                                          | 26  |

|                                                         |     |
|---------------------------------------------------------|-----|
| inserting .....                                         | 28  |
| renaming.....                                           | 27  |
| IGMP (Internet Group Management Protocol)               |     |
| IGMPv1 .....                                            | 276 |
| IGMPv2 .....                                            | 276 |
| IGMPv3.....                                             | 276 |
| setting the version .....                               | 281 |
| verifying the version.....                              | 286 |
| IGPs (interior gateway protocols) .....                 | 237 |
| overview .....                                          | 102 |
| VPNs .....                                              | 237 |
| <i>See also</i> OSPF; RIP                               |     |
| IKE (Internet Key Exchange)                             |     |
| description .....                                       | 252 |
| preshared key (configuration editor).....               | 259 |
| preshared key (Quick Configuration) .....               | 254 |
| import routing policy, for Layer 2 VPNs .....           | 244 |
| import statement, for routing policies .....            | 298 |
| inbound router, in an LSP .....                         | 201 |
| incoming metric (RIP)                                   |     |
| description .....                                       | 140 |
| modifying .....                                         | 147 |
| inet routing table.....                                 | 284 |
| ingress router <i>See</i> inbound router; LSPs          |     |
| injecting routes .....                                  | 323 |
| insert command .....                                    | 28  |
| inserting configuration identifiers.....                | 28  |
| interface naming conventions .....                      | 45  |
| interfaces <i>See</i> ATM-for-ADSL interfaces; loopback |     |
| interfaces; management interfaces; network              |     |
| interfaces; services interfaces; ports                  |     |
| Interfaces page.....                                    | 48  |
| for E1 .....                                            | 49  |
| for Fast Ethernet.....                                  | 52  |
| for serial interfaces .....                             | 61  |
| for T1 .....                                            | 54  |
| for T3 (DS3).....                                       | 58  |
| interior gateway protocols <i>See</i> IGPs              |     |
| internal BGP <i>See</i> IBGP                            |     |
| Internet Control Message Protocol policers.....         | 349 |
| Internet Group Management Protocol <i>See</i> IGMP      |     |
| Internet Key Exchange <i>See</i> IKE                    |     |
| Internet routing, with BGP .....                        | 177 |
| invalid configuration, replacing                        |     |
| with J-Web .....                                        | 21  |
| with the CLI .....                                      | 33  |
| invalid routes, rejecting.....                          | 322 |
| inverting the transmit clock .....                      | 63  |
| IP Security <i>See</i> IPSec                            |     |
| IPSec (IP Security)                                     |     |
| IKE <i>See</i> IKE                                      |     |
| security associations.....                              | 252 |
| tunnels <i>See</i> IPSec tunnels                        |     |
| verifying tunnels.....                                  | 265 |
| IPSec security associations .....                       | 252 |

- See also* IKE
- IPSec tunnels
- IKE key (configuration editor) ..... 259
  - IKE key (Quick Configuration) ..... 254
  - incoming traffic filters ..... 252
  - IPSec rule (configuration editor) ..... 260
  - local endpoint (Quick Configuration) ..... 254
  - NAT pools (configuration editor) ..... 262
  - outgoing traffic filters ..... 252
  - overview ..... 251
  - private addresses (Quick Configuration) ..... 254
  - Quick Configuration ..... 252
  - remote endpoint (Quick Configuration) ..... 254
  - requirements ..... 252
  - services interfaces (configuration editor) ..... 255
  - services sets (configuration editor) ..... 256
  - stateful firewall filter (configuration editor) ..... 260
  - verifying ..... 265
- IPSec Tunnels page ..... 253
- field summary ..... 254
- IPv6 support ..... 97
- ITU Annex B non-UR-2 operating mode ..... 68
- ITU Annex B UR-2 operating mode ..... 68
- ITU DMT operating mode ..... 68
- ## J
- J-series
- BGP routing ..... 177
  - configuration tools ..... 3
  - CoS overview ..... 307
  - CoS with DiffServ ..... 371
  - firewall filter overview ..... 298
  - firewall filters ..... 331
  - IPSec tunnels ..... 251
  - MPLS for VPNs overview ..... 197
  - MPLS traffic engineering ..... 213
  - multicast ..... 279
  - multicast overview ..... 269
  - NAT ..... 331
  - network interfaces ..... 41
  - OSPF routing ..... 155
  - PPPoE ..... 77
  - release notes, URL ..... xvii
  - RIP routing ..... 139
  - routing policies ..... 317
  - routing policy overview ..... 293
  - routing protocols overview ..... 97
  - static routing ..... 127
  - VPNs ..... 227
- J-Web configuration editor
- BGP ..... 181
  - clickable configuration, committing ..... 12
  - clickable configuration, discarding changes ..... 11
  - clickable configuration, editing ..... 8
  - committing a text file, with caution ..... 13
  - configuration text, viewing ..... 12
  - editing a text file, with caution ..... 13
  - IPSec tunnels ..... 254
  - managing files ..... 15
  - MPLS traffic engineering ..... 215
  - network interfaces ..... 64
  - OSPF ..... 160
  - PPPoE ..... 82
  - PPPoE over ATM-for-ADSL ..... 82
  - RIP ..... 143
  - static routes ..... 132
  - uploading a file ..... 14
  - VPNs ..... 230
- 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 ..... 310
  - CoS functions ..... 309
  - DiffServ implementation ..... 309
  - release notes, URL ..... xvii
- junos-algs-outbound group, for stateful firewall filters ..... 300
- ## K
- keepalive interval, for LDP-signaled LSPs ..... 217
  - keepalive messages, for session hold time ..... 118
- ## L
- Label Distribution Protocol *See* LDP
- label switching ..... 200
- label-switched paths *See* LSPs
- label-switching routers (LSRs) ..... 201
- labels, MPLS ..... 202
- label operations ..... 202
  - PHP ..... 203
- Layer 2 circuits
- AS number ..... 236
  - basic, description ..... 229
  - encapsulation ..... 232
  - IGPs ..... 237
  - MPLS ..... 233
  - neighbor address ..... 240
  - participating interfaces ..... 231
  - signaling protocols ..... 237
  - task overview ..... 230
  - verifying PE router connections ..... 249
  - verifying PE router interfaces ..... 249

|                                             |     |                                                          |          |
|---------------------------------------------|-----|----------------------------------------------------------|----------|
| virtual circuit ID.....                     | 240 | load patch command.....                                  | 35       |
| Layer 2 VPNs                                |     | load replace command.....                                | 35       |
| AS number.....                              | 236 | loading a configuration file                             |          |
| basic, description.....                     | 228 | CLI configuration editor.....                            | 34       |
| BGP.....                                    | 235 | downloading (J-Web).....                                 | 20       |
| encapsulation.....                          | 232 | rollback (J-Web).....                                    | 21       |
| export routing policies.....                | 245 | rollback command.....                                    | 32       |
| IGPs.....                                   | 237 | uploading (J-Web).....                                   | 14       |
| import routing policies.....                | 244 | without specifying full hierarchy.....                   | 35       |
| MPLS.....                                   | 233 | local preference                                         |          |
| overview.....                               | 211 | description.....                                         | 120      |
| participating interfaces.....               | 231 | high value preferred.....                                | 121      |
| routing instance.....                       | 241 | role in route selection.....                             | 119      |
| signaling protocols.....                    | 237 | local tunnel endpoint, IPSec.....                        | 254      |
| task overview.....                          | 230 | locked configuration.....                                | 23       |
| verifying PE router connections.....        | 249 | logical interfaces                                       |          |
| verifying PE router interfaces.....         | 249 | adaptive shaping for.....                                | 401      |
| Layer 3 VPNs                                |     | adding (configuration editor).....                       | 66       |
| AS number.....                              | 236 | ATM-for-ADSL.....                                        | 68       |
| basic, description.....                     | 229 | CoS rules for.....                                       | 397, 401 |
| BGP.....                                    | 235 | E1.....                                                  | 50       |
| IGPs.....                                   | 237 | Fast Ethernet.....                                       | 53       |
| overview.....                               | 211 | inside services interface, IPSec.....                    | 255      |
| participating interfaces.....               | 231 | outside services interface, IPSec.....                   | 255      |
| route target.....                           | 241 | serial.....                                              | 62       |
| routing instance.....                       | 241 | T1.....                                                  | 55       |
| routing policies.....                       | 247 | T3.....                                                  | 59       |
| signaling protocols.....                    | 237 | virtual channels for.....                                | 397      |
| task overview.....                          | 230 | logical units                                            |          |
| verifying PE router connections.....        | 249 | adding (configuration editor).....                       | 66       |
| LDP (Label Distribution Protocol)           |     | E1 interface.....                                        | 50       |
| and OSPF for VPNs.....                      | 237 | Fast Ethernet interface.....                             | 53       |
| LDP-signaled LSPs.....                      | 215 | number in interface name.....                            | 47       |
| messages.....                               | 204 | pp0 interface.....                                       | 85       |
| operation.....                              | 204 | PPPoE encapsulation.....                                 | 83       |
| overview.....                               | 214 | PPPoE over ATM-for-ADSL encapsulation.....               | 83       |
| requirements.....                           | 214 | serial interface.....                                    | 62       |
| verifying LSPs.....                         | 222 | T1 interface.....                                        | 55       |
| verifying neighbors.....                    | 220 | T3 interface.....                                        | 59       |
| verifying sessions.....                     | 221 | long buildout <i>See</i> line buildout                   |          |
| verifying traffic forwarding.....           | 222 | longer route list match type.....                        | 320      |
| LDP neighbors, verifying.....               | 220 | loopback address, for PE routers in VPNs.....            | 237      |
| LDP-signaled LSP <i>See</i> LDP             |     | loopback interfaces, applying stateless firewall filters |          |
| leaves.....                                 | 272 | to (configuration editor).....                           | 358      |
| <i>See also</i> multicast                   |     | loose hops, RSVP.....                                    | 206      |
| line buildout                               |     | loss priority, CoS.....                                  | 310      |
| T1.....                                     | 57  | LSAs (link-state advertisements)                         |          |
| T3.....                                     | 60  | description.....                                         | 112      |
| line speed, serial interface.....           | 63  | three-way handshake.....                                 | 112      |
| link coloring, for MPLS path selection..... | 207 | LSPs (label-switched paths)                              |          |
| link states, verifying.....                 | 70  | bandwidth.....                                           | 220      |
| link-state advertisements <i>See</i> LSAs   |     | description.....                                         | 200      |
| load command.....                           | 34  | disabling CSPF.....                                      | 220      |
| load merge command.....                     | 35  | dynamic LSPs.....                                        | 203      |
| load override command.....                  | 35  | for RSVP in a VPN.....                                   | 234      |

|                                                         |      |
|---------------------------------------------------------|------|
| keepalive interval for LDP link .....                   | 217  |
| label operations .....                                  | 202  |
| label switching .....                                   | 200  |
| labels .....                                            | 202  |
| LDP .....                                               | 204  |
| LDP-signaled LSPs .....                                 | 215  |
| LSR types .....                                         | 201  |
| overview .....                                          | 213  |
| PHP .....                                               | 203  |
| RSVP .....                                              | 204  |
| RSVP-signaled LSPs .....                                | 217  |
| static LSPs .....                                       | 203  |
| verifying LDP-signaled LSPs .....                       | 220  |
| verifying RSVP-signaled LSPs .....                      | 223  |
| LSRs (label-switching routers) .....                    | 201  |
| <b>M</b>                                                |      |
| management interfaces, disabling PIM on .....           | 283  |
| managing files <i>See</i> file management               |      |
| manuals                                                 |      |
| comments on .....                                       | xxii |
| mapping, CoS forwarding classes to schedulers .....     | 395  |
| match conditions                                        |      |
| routing policy .....                                    | 294  |
| routing policy, summary of .....                        | 294  |
| stateful firewall filter and NAT .....                  | 301  |
| stateless firewall filters .....                        | 303  |
| stateless firewall filters, summary of .....            | 304  |
| match types .....                                       | 320  |
| maximum hop count, RIP .....                            | 108  |
| maximum transmission unit <i>See</i> MTU                |      |
| MED (multiple exit discriminator)                       |      |
| description .....                                       | 122  |
| role in route selection .....                           | 119  |
| merging a configuration file .....                      | 35   |
| example .....                                           | 37   |
| messages, LDP .....                                     | 204  |
| metrics <i>See</i> path cost metrics                    |      |
| MF classifier .....                                     | 374  |
| MPLS (Multiprotocol Label Switching) .....              | 208  |
| dynamic LSPs .....                                      | 203  |
| label operations .....                                  | 202  |
| label switching .....                                   | 200  |
| labels .....                                            | 202  |
| Layer 2 VPNs and Layer 2 circuits .....                 | 233  |
| LDP .....                                               | 204  |
| LSP for RSVP in a VPN .....                             | 234  |
| LSPs .....                                              | 200  |
| LSR types .....                                         | 201  |
| overview .....                                          | 197  |
| PHP .....                                               | 203  |
| RSVP .....                                              | 204  |
| static LSPs .....                                       | 203  |
| traffic engineering <i>See</i> MPLS traffic engineering |      |
| verifying .....                                         | 220  |

|                                                           |        |
|-----------------------------------------------------------|--------|
| <i>See also</i> VPNs                                      |        |
| MPLS traffic engineering                                  |        |
| LDP signaling .....                                       | 214    |
| LDP-signaled LSPs .....                                   | 215    |
| overview .....                                            | 213    |
| requirements .....                                        | 214    |
| RSVP signaling .....                                      | 214    |
| RSVP-signaled LSPs .....                                  | 217    |
| signaling protocols overview .....                        | 204    |
| verifying LDP neighbors .....                             | 220    |
| verifying LDP sessions .....                              | 221    |
| verifying LDP-signaled LSPs .....                         | 222    |
| verifying RSVP neighbors .....                            | 223    |
| verifying RSVP sessions .....                             | 224    |
| verifying RSVP-signaled LSPs .....                        | 224    |
| verifying traffic forwarding over LDP-signaled LSPs ..... | 222    |
| MSDP (Multicast Source Discovery Protocol) .....          | 277    |
| MTU (maximum transmission unit)                           |        |
| E1 .....                                                  | 51     |
| T1 .....                                                  | 56     |
| T3 .....                                                  | 60, 62 |
| multiarea network, OSPF .....                             | 162    |
| multicast                                                 |        |
| administrative scoping .....                              | 275    |
| architecture .....                                        | 272    |
| Auto-RP .....                                             | 276    |
| BSR .....                                                 | 276    |
| downstream interface .....                                | 272    |
| DVMRP .....                                               | 275    |
| forwarding state notation .....                           | 273    |
| *,G notation .....                                        | 273    |
| IGMP <i>See</i> IGMP                                      |        |
| IP address ranges .....                                   | 273    |
| MSDP .....                                                | 277    |
| network elements .....                                    | 273    |
| overview .....                                            | 269    |
| PGM .....                                                 | 277    |
| PIM dense mode <i>See</i> PIM                             |        |
| PIM source-specific multicast (SSM) .....                 | 276    |
| PIM sparse mode <i>See</i> PIM                            |        |
| preparation .....                                         | 280    |
| preventing routing loops .....                            | 274    |
| protocols .....                                           | 275    |
| reverse-path forwarding (RPF) .....                       | 274    |
| routing modes <i>See</i> multicast routing modes          |        |
| S,G notation .....                                        | 273    |
| SAP and SDP <i>See</i> SAP; SDP                           |        |
| session announcements .....                               | 280    |
| shortest-path tree (SPT) .....                            | 275    |
| static RP .....                                           | 282    |
| <i>See also</i> RP                                        |        |
| subnetwork leaves and branches .....                      | 272    |
| upstream interface .....                                  | 272    |
| verifying IGMP versions .....                             | 286    |

|                                                   |     |
|---------------------------------------------------|-----|
| verifying PIM mode and interface configuration .. | 287 |
| verifying PIM RPF routing table .....             | 288 |
| verifying PIM RPs .....                           | 287 |
| verifying SAP and SDP configuration .....         | 286 |
| multicast routing modes .....                     |     |
| dense mode .....                                  | 274 |
| dense mode, caution for use .....                 | 274 |
| sparse mode .....                                 | 274 |
| Multicast Source Discovery Protocol .....         | 277 |
| multifield classifier .....                       | 374 |
| multiple exit discriminator <i>See</i> MED        |     |
| multiple push label operation .....               | 203 |
| Multiprotocol Label Switching <i>See</i> MPLS     |     |

## N

|                                                               |          |
|---------------------------------------------------------------|----------|
| names, of network interfaces .....                            | 45       |
| NAPT .....                                                    | 299      |
| NAT (Network Address Translation) .....                       |          |
| actions .....                                                 | 301      |
| applying to an interface (configuration editor) ..            | 341      |
| configuration editor .....                                    | 336, 338 |
| description .....                                             | 298      |
| enabling (Quick Configuration) .....                          | 335      |
| match conditions .....                                        | 301      |
| pools for IPsec tunnels (configuration editor) ..             | 262      |
| preparation .....                                             | 332      |
| Quick Configuration .....                                     | 332      |
| sample rules .....                                            | 337      |
| verifying .....                                               | 364      |
| neighbors <i>See</i> BGP peers; OSPF neighbors; RIP neighbors |          |
| Network Address Port Translation (NAPT) .....                 | 299      |
| Network Address Translation <i>See</i> NAT                    |          |
| network interfaces .....                                      |          |
| adding .....                                                  | 64       |
| ATM-for-ADSL configuration .....                              | 66       |
| deleting .....                                                | 69       |
| DS3 configuration .....                                       | 57       |
| E1 configuration .....                                        | 49       |
| enabling PIM on .....                                         | 283      |
| enabling RIP on .....                                         | 142      |
| Fast Ethernet configuration .....                             | 52       |
| multicast, upstream and downstream .....                      | 272      |
| naming conventions .....                                      | 45       |
| overview .....                                                | 44       |
| preparation .....                                             | 47       |
| serial configuration .....                                    | 60       |
| supported .....                                               | 44       |
| T1 configuration .....                                        | 53       |
| T3 configuration .....                                        | 57       |
| verifying ATM-for-ADSL properties .....                       | 72       |
| verifying link states .....                                   | 70       |
| verifying PIM on .....                                        | 287      |
| verifying properties .....                                    | 71       |
| verifying RIP message exchange .....                          | 151      |
| verifying RIP on .....                                        | 151      |
| VPN configuration .....                                       | 231      |
| networks .....                                                | 228      |
| description .....                                             | 102      |
| designated router <i>See</i> designated router, OSPF          |          |
| path cost metrics <i>See</i> path cost metrics                |          |
| PPPoE session on an ATM-for-ADSL loop .....                   | 80       |
| PPPoE session on an Ethernet loop .....                       | 79       |
| sample BGP AS path .....                                      | 121      |
| sample BGP confederation .....                                | 189      |
| sample BGP confederations .....                               | 126      |
| sample BGP external and internal links .....                  | 184      |
| sample BGP local preference use .....                         | 120      |
| sample BGP MED use .....                                      | 122      |
| sample BGP peer network .....                                 | 182      |
| sample BGP peer session .....                                 | 117      |
| sample BGP route reflector (one cluster) ...                  | 123, 186 |
| sample BGP route reflectors (cluster of clusters) ..          | 125      |
| sample BGP route reflectors (multiple clusters) ..            | 124      |
| sample distance-vector routing .....                          | 108      |
| sample LSP topology .....                                     | 201      |
| sample multiarea OSPF routing .....                           | 114      |
| sample OSPF backbone area .....                               | 115      |
| sample OSPF multiarea network .....                           | 162      |
| sample OSPF network with stubs and NSSAs .....                | 116      |
| sample OSPF single-area network .....                         | 161      |
| sample OSPF stub areas and NSSAs .....                        | 165      |
| sample OSPF topology .....                                    | 173      |
| sample poison reverse routing .....                           | 110      |
| sample RIP network with incoming metric .....                 | 146      |
| sample RIP network with outgoing metric .....                 | 148      |
| sample RIP topology .....                                     | 143      |
| sample route advertisement .....                              | 105      |
| sample route aggregation .....                                | 106      |
| sample routing topology .....                                 | 103      |
| sample RSVP topology .....                                    | 206      |
| sample split horizon routing .....                            | 109      |
| sample static route, preferred path .....                     | 134      |
| sample stub network for static routes .....                   | 132      |
| sample unidirectional routing .....                           | 111      |
| sample VPN topology .....                                     | 228      |
| static routing .....                                          | 104      |
| trusted .....                                                 | 298      |
| untrusted .....                                               | 298      |
| <i>See also</i> VPNs                                          |          |
| next hop .....                                                |          |
| address for static routes .....                               | 131      |
| defining for static routes .....                              | 133      |
| qualified, defining for static routes .....                   | 135      |
| qualified, for static routes .....                            | 128      |
| service set, for IPsec tunnels .....                          | 256      |
| non-UR-2 operating mode .....                                 | 68       |
| not-so-stubby areas <i>See</i> NSSAs                          |          |
| notice icons .....                                            | xix      |



|                                       |     |
|---------------------------------------|-----|
| NSSAs (not-so-stubby areas)           |     |
| area ID (configuration editor) .....  | 163 |
| area ID (Quick Configuration) .....   | 158 |
| area type (Quick Configuration) ..... | 159 |
| creating (configuration editor) ..... | 165 |
| description .....                     | 115 |
| example .....                         | 116 |
| sample topology .....                 | 165 |
| numeric range match conditions .....  | 304 |

## O

|                                                       |          |
|-------------------------------------------------------|----------|
| OK button                                             |          |
| J-Web configuration editor .....                      | 11       |
| Quick Configuration .....                             | 8        |
| Open Shortest Path First protocol <i>See</i> OSPF     |          |
| operational mode, entering during configuration ..... | 33       |
| origin, of BGP route .....                            | 121      |
| orlonger route list match type .....                  | 321      |
| OSPF (Open Shortest Path First)                       |          |
| and LDP for VPNs .....                                | 238      |
| and RSVP for VPNs .....                               | 239      |
| area border routers <i>See</i> area border routers    |          |
| area type (Quick Configuration) .....                 | 159      |
| areas .....                                           | 113, 156 |
| <i>See also</i> area border routers; backbone area;   |          |
| NSSAs; stub areas                                     |          |
| authenticating exchanges (OSPFv2 only) .....          | 169      |
| backbone area <i>See</i> backbone area                |          |
| controlling designated router election .....          | 170      |
| controlling route cost .....                          | 168      |
| designated router <i>See</i> designated router, OSPF  |          |
| designating OSPF interfaces (configuration            |          |
| editor) .....                                         | 162–163  |
| designating OSPF interfaces (Quick                    |          |
| Configuration) .....                                  | 159      |
| enabling (Quick Configuration) .....                  | 158      |
| enabling, description .....                           | 155      |
| ensuring efficient operation .....                    | 167      |
| injecting OSPF routes into BGP .....                  | 322      |
| LSAs .....                                            | 112      |
| multiarea network (configuration editor) .....        | 162      |
| NSSAs <i>See</i> NSSAs                                |          |
| overview .....                                        | 111, 155 |
| path cost metrics <i>See</i> path cost metrics        |          |
| Quick Configuration .....                             | 156      |
| requirements .....                                    | 156      |
| route preferences .....                               | 167      |
| router ID (configuration editor) .....                | 160      |
| router ID (Quick Configuration) .....                 | 158      |
| sample multiarea network .....                        | 162      |
| sample network topology .....                         | 173      |
| sample NSSAs .....                                    | 165      |
| sample single-area network .....                      | 161      |
| sample stub areas .....                               | 165      |
| single-area network (configuration editor) .....      | 161      |

|                                         |         |
|-----------------------------------------|---------|
| stub areas <i>See</i> stub areas        |         |
| supported versions .....                | 112     |
| three-way handshake .....               | 112     |
| tuning an OSPF network .....            | 167     |
| verifying host reachability .....       | 174     |
| verifying neighbors .....               | 172     |
| verifying RIP-enabled interfaces .....  | 171     |
| verifying routes .....                  | 173     |
| OSPF interfaces                         |         |
| enabling .....                          | 159     |
| enabling (configuration editor) .....   | 162–163 |
| enabling, for area border routers ..... | 165     |
| verifying .....                         | 171     |
| OSPF neighbors, verifying .....         | 172     |
| OSPF page .....                         | 157     |
| field summary .....                     | 158     |
| outbound router, in an LSP .....        | 201     |
| outgoing metric (RIP)                   |         |
| description .....                       | 140     |
| modifying .....                         | 149     |
| output queues                           |         |
| assigning forwarding classes .....      | 379     |
| sample assignments .....                | 378     |
| overriding a configuration file .....   | 35      |
| example .....                           | 36      |

## P

|                                                    |     |
|----------------------------------------------------|-----|
| P routers <i>See</i> provider routers              |     |
| packet encapsulation                               |     |
| E1 interfaces .....                                | 50  |
| Layer 2 circuits .....                             | 232 |
| Layer 2 VPNs .....                                 | 232 |
| serial interfaces .....                            | 62  |
| T1 interfaces .....                                | 55  |
| T3 interfaces .....                                | 59  |
| packets                                            |     |
| applying CoS scheduling rules .....                | 397 |
| handling packet fragments .....                    | 343 |
| handling packet fragments (configuration           |     |
| editor) .....                                      | 355 |
| PADI .....                                         | 80  |
| PADO .....                                         | 81  |
| PADR .....                                         | 81  |
| PADS .....                                         | 81  |
| PADT .....                                         | 81  |
| PPPoE discovery .....                              | 80  |
| RIP, description .....                             | 109 |
| PADI packets .....                                 | 80  |
| PADO packets .....                                 | 81  |
| PADR packets .....                                 | 81  |
| PADS packets .....                                 | 81  |
| PADT packets .....                                 | 81  |
| parentheses, in syntax descriptions .....          | xx  |
| passive routes, rejection, in static routing ..... | 129 |

|                                                        |          |
|--------------------------------------------------------|----------|
| password                                               |          |
| for OSPFv2 authentication                              | 170      |
| for RIPv2 authentication                               | 149      |
| patching a configuration file                          | 35       |
| path cost metrics                                      |          |
| for OSPF routes, description                           | 113, 156 |
| for OSPF routes, modifying                             | 168      |
| for RIP routes, description                            | 139      |
| for RIP routes, modifying                              | 146      |
| path selection <i>See</i> traffic engineering database |          |
| path-vector protocol <i>See</i> BGP                    |          |
| PE (provider edge) routers                             | 228      |
| description                                            | 209      |
| route distinguishers                                   | 241      |
| verifying Layer 2 circuit connections                  | 249      |
| verifying Layer 2 circuit interfaces                   | 249      |
| verifying Layer 2 VPN connections                      | 249      |
| verifying Layer 2 VPN interfaces                       | 249      |
| verifying Layer 3 VPN connections                      | 249      |
| VPN task overview                                      | 230      |
| VPN topology                                           | 228      |
| <i>See also</i> VPNs                                   |          |
| peering sessions <i>See</i> BGP peers; BGP sessions    |          |
| penultimate hop popping (PHP)                          | 203      |
| penultimate router, in an LSP                          | 201      |
| permanent routes, adding                               | 127      |
| PGM (Pragmatic General Multicast)                      | 277      |
| PHP (penultimate hop popping)                          | 203      |
| Physical Interface Modules, number in interface        |          |
| name                                                   | 46       |
| PIM (Protocol Independent Multicast)                   |          |
| dense mode                                             | 276      |
| disabling on the network management                    |          |
| interface                                              | 282      |
| RPF routing table group                                | 284      |
| source-specific multicast (SSM)                        | 276      |
| sparse mode                                            | 276      |
| static RP router                                       | 282      |
| supported versions                                     | 279      |
| verifying the mode                                     | 287      |
| verifying the RP                                       | 287      |
| PIMs (Physical Interface Modules), number in interface |          |
| name                                                   | 46       |
| ping                                                   |          |
| verifying link states                                  | 70       |
| VPN connection                                         | 248      |
| ping command                                           | 369      |
| explanation                                            | 369      |
| Ping Host page, output for BGP                         | 193      |
| ping mpls l2circuit interface command                  | 249      |
| ping mpls l2circuit virtual-circuit command            | 249      |
| ping mpls l2vpn instance                               | 249      |
| ping mpls l2vpn interface command                      | 249      |
| ping mpls l3vpn command                                | 249      |
| ping trusted-nw-trusted-host                           | 364      |
| explanation                                            | 365      |
| ping untrusted-nw-untrusted-host command               | 364      |
| explanation                                            | 365      |
| Point-to-Point Protocol over Ethernet <i>See</i> PPPoE |          |
| poison reverse technique                               | 109      |
| policers                                               |          |
| description                                            | 311      |
| for firewall filter                                    | 373      |
| for stateless firewall filters                         | 349      |
| policy <i>See</i> routing policies                     |          |
| pop label operation                                    | 202      |
| ports                                                  |          |
| DS1 <i>See</i> T1 ports                                |          |
| DS3 <i>See</i> T3 ports                                |          |
| E1 <i>See</i> E1 ports                                 |          |
| number in interface name                               | 47       |
| T1 <i>See</i> T1 ports                                 |          |
| T3 <i>See</i> T3 ports                                 |          |
| pp0                                                    |          |
| creating                                               | 85       |
| enabling CHAP                                          | 87       |
| information about                                      | 90       |
| logical Ethernet interface on                          | 85       |
| PPP over ATM AAL5 multiplex encapsulation              | 69       |
| PPP over ATM-for-ADSL <i>See</i> PPPoA                 |          |
| PPP over Ethernet <i>See</i> PPPoE                     |          |
| PPPoA (PPP over ATM-for-ADSL)                          |          |
| logical encapsulation                                  | 69       |
| physical encapsulation                                 | 68       |
| PPPoE (Point-to-Point Protocol over Ethernet)          | 79       |
| CHAP                                                   | 87       |
| client and server                                      | 78       |
| creating the pp0 interface                             | 85       |
| discovery packets                                      | 80       |
| encapsulation on an Ethernet interface                 | 83       |
| interfaces                                             | 79       |
| overview                                               | 78       |
| preparation                                            | 82       |
| sample topology                                        | 79       |
| service type                                           | 86       |
| session limit                                          | 81       |
| session overview                                       | 81       |
| session reconnection time                              | 86       |
| verifying interfaces                                   | 90       |
| verifying sessions                                     | 91       |
| verifying statistics                                   | 92       |
| verifying version information                          | 92       |
| <i>See also</i> PPPoE over ATM-for-ADSL                |          |
| PPPoE Active Discovery Initiation (PADI) packets       | 80       |
| PPPoE Active Discovery Offer (PADO) packets            | 81       |
| PPPoE Active Discovery Request (PADR) packets          | 81       |
| PPPoE Active Discovery Session-Confirmation (PADS)     |          |
| packets                                                | 81       |
| PPPoE Active Discovery Termination (PADT) packets      | 81       |

|                                               |     |
|-----------------------------------------------|-----|
| PPPoE over ATM LLC encapsulation .....        | 69  |
| PPPoE over ATM-for-ADSL .....                 | 79  |
| CHAP .....                                    | 87  |
| creating the pp0 interface .....              | 85  |
| encapsulation .....                           | 83  |
| preparation .....                             | 82  |
| sample topology .....                         | 79  |
| verifying configuration .....                 | 89  |
| <i>See also</i> PPPoE                         |     |
| PPPoEoA <i>See</i> PPPoE over ATM-for-ADSL    |     |
| Pragmatic General Multicast .....             | 277 |
| preferences                                   |     |
| for OSPF routes .....                         | 167 |
| for static routes .....                       | 128 |
| setting for static routes .....               | 135 |
| prefix-length-range match type .....          | 321 |
| propagation, suppressing .....                | 327 |
| properties, verifying                         |     |
| for ATM-for-ADSL network interfaces .....     | 72  |
| for network interfaces .....                  | 71  |
| Protocol Independent Multicast <i>See</i> PIM |     |
| protocols                                     |     |
| Auto-RP .....                                 | 276 |
| BGP <i>See</i> BGP                            |     |
| distance vector <i>See</i> RIP                |     |
| DVMRP .....                                   | 275 |
| EGPs .....                                    | 102 |
| IGMP <i>See</i> IGMP                          |     |
| IGPs .....                                    | 102 |
| IPSec <i>See</i> IPSec                        |     |
| LDP <i>See</i> LDP                            |     |
| MPLS <i>See</i> MPLS                          |     |
| MSDP .....                                    | 277 |
| multicast <i>See</i> multicast                |     |
| NAT <i>See</i> NAT                            |     |
| OSPF <i>See</i> OSPF                          |     |
| overview .....                                | 97  |
| path vector <i>See</i> BGP                    |     |
| PGM .....                                     | 277 |
| PIM dense mode <i>See</i> PIM                 |     |
| PIM source-specific multicast (SSM) .....     | 276 |
| PIM sparse mode <i>See</i> PIM                |     |
| PPPoE <i>See</i> PPPoE                        |     |
| RIP <i>See</i> RIP                            |     |
| RSVP <i>See</i> RSVP                          |     |
| SAP and SDP <i>See</i> SAP; SDP               |     |
| provider edge routers <i>See</i> PE routers   |     |
| provider routers .....                        | 228 |
| description .....                             | 209 |
| VPN task overview .....                       | 230 |
| VPN topology .....                            | 228 |
| <i>See also</i> VPNs                          |     |
| push label operation .....                    | 202 |

## Q

|                                     |     |
|-------------------------------------|-----|
| queuing rules, CoS .....            | 397 |
| Quick Configuration                 |     |
| BGP page .....                      | 179 |
| buttons .....                       | 8   |
| E1 Interfaces page .....            | 49  |
| Fast Ethernet Interfaces page ..... | 52  |
| Interfaces page .....               | 48  |
| IPSec Tunnels page .....            | 253 |
| network interfaces .....            | 47  |
| OSPF page .....                     | 157 |
| overview .....                      | 7   |
| RIP page .....                      | 141 |
| serial Interfaces page .....        | 61  |
| Static Routes page .....            | 130 |
| Summary page .....                  | 7   |
| T1 Interfaces page .....            | 54  |
| T3 (DS3) Interfaces page .....      | 58  |

## R

|                                                        |       |
|--------------------------------------------------------|-------|
| radio buttons                                          |       |
| Delete Configuration Below This Point .....            | 11    |
| Discard All Changes .....                              | 11    |
| Discard Changes Below This Point .....                 | 11    |
| random early detection <i>See</i> RED drop profiles    |       |
| reactivate command .....                               | 29    |
| RED (random early detection) drop profiles .....       | 388   |
| samples .....                                          | 388   |
| redistributing routes .....                            | 323   |
| Refresh button .....                                   | 11    |
| rejecting invalid routes .....                         | 322   |
| relative option .....                                  | 35    |
| release notes, URL .....                               | xvii  |
| remote tunnel endpoint, IPSec .....                    | 254   |
| rename command .....                                   | 27    |
| renaming configuration identifiers .....               | 27    |
| 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 |
| setting (CLI configuration editor) .....               | 33    |
| setting (J-Web) .....                                  | 21    |
| viewing (CLI configuration editor) .....               | 33    |
| viewing (J-Web) .....                                  | 21–22 |
| reservation <i>See</i> RSVP                            |       |
| Resource Reservation Protocol <i>See</i> RSVP          |       |
| reverse-path forwarding <i>See</i> RPF                 |       |
| rewrite rules                                          |       |
| description .....                                      | 311   |
| replacing DSCPs .....                                  | 380   |
| sample rules .....                                     | 380   |

|                                                                 |          |                                                            |     |
|-----------------------------------------------------------------|----------|------------------------------------------------------------|-----|
| when applied .....                                              | 314      | route selection .....                                      | 119 |
| RIB <i>See</i> routing table .....                              |          | BGP process .....                                          | 119 |
| RIP (Routing Information Protocol) .....                        |          | BGP, determining by AS path .....                          | 121 |
| authentication (RIPv2 only) .....                               | 140      | BGP, determining by local preference .....                 | 120 |
| authentication (RIPv2 only), configuring .....                  | 149      | BGP, determining by MED metric .....                       | 122 |
| basic network (configuration editor) .....                      | 143      | BGP, lowest origin value preferred .....                   | 121 |
| designating RIP interfaces .....                                | 142      | static routes, defining .....                              | 133 |
| distance vector protocol .....                                  | 107      | route target, in a VPN routing instance .....              | 241 |
| efficiency techniques .....                                     | 109      | route targets, VPN .....                                   | 211 |
| enabling (Quick Configuration) .....                            | 142      | route-flap damping .....                                   | 327 |
| maximum hop count .....                                         | 108      | router <i>See</i> Services Router .....                    |     |
| overview .....                                                  | 107, 139 | routing .....                                              | 97  |
| packets .....                                                   | 109      | advertisements .....                                       | 105 |
| path cost metrics <i>See</i> path cost metrics .....            |          | aggregation .....                                          | 105 |
| poison reverse technique .....                                  | 109      | BGP <i>See</i> BGP .....                                   |     |
| Quick Configuration .....                                       | 140      | configuring PPPoE .....                                    | 77  |
| requirements .....                                              | 140      | configuring VPNs .....                                     | 227 |
| routing policy (configuration editor) .....                     | 143      | dynamic .....                                              | 104 |
| sample network with incoming metric .....                       | 146      | filtering and classifying routes .....                     | 291 |
| sample network with outgoing metric .....                       | 148      | filtering routes with policies .....                       | 317 |
| sample topology .....                                           | 143      | filtering traffic through a firewall .....                 | 331 |
| split horizon technique .....                                   | 109      | forwarding tables .....                                    | 103 |
| supported versions .....                                        | 107      | from one source to many destinations .....                 | 279 |
| traffic control with metrics <i>See</i> path cost metrics ..... |          | in multiple ASs with BGP .....                             | 177 |
| traffic control with metrics, configuring .....                 | 146      | in one AS with OSPF .....                                  | 155 |
| unidirectional limitations .....                                | 110      | in one AS with RIP .....                                   | 139 |
| verifying host reachability .....                               | 152      | MPLS for VPNs .....                                        | 197 |
| verifying RIP message exchange .....                            | 151      | MPLS traffic engineering .....                             | 213 |
| verifying RIP-enabled interfaces .....                          | 151      | multicast <i>See</i> multicast .....                       |     |
| RIP neighbors, verifying .....                                  | 151      | neighbors <i>See</i> BGP peers; OSPF neighbors; RIP .....  |     |
| RIP page .....                                                  | 141      | neighbors .....                                            |     |
| field summary .....                                             | 142      | OSPF <i>See</i> OSPF .....                                 |     |
| rollback ? command .....                                        | 33       | overriding default packet forwarding with CoS ..           | 371 |
| rollback command .....                                          | 32       | protecting local IP addresses with NAT .....               | 331 |
| rollback rescue command .....                                   | 32       | protocol overview .....                                    | 97  |
| rolling back a configuration file .....                         |          | RIP <i>See</i> RIP .....                                   |     |
| during configuration (CLI configuration editor) ....            | 32       | RIP statistics .....                                       | 151 |
| during configuration (J-Web) .....                              | 21       | routing tables .....                                       | 103 |
| route advertisements .....                                      |          | static <i>See</i> static routing .....                     |     |
| AS path in .....                                                | 121      | through IPsec tunnels .....                                | 251 |
| BGP, update messages .....                                      | 118      | VPNs .....                                                 | 227 |
| description .....                                               | 105      | <i>See also</i> protocols; routing policies; routing ..... |     |
| external, EBGp .....                                            | 118      | solutions .....                                            |     |
| internal, IBGP .....                                            | 119      | Routing Engine .....                                       |     |
| LSAs .....                                                      | 112      | handling packet fragments for (configuration .....         | 353 |
| stub areas and NSSAs, to control .....                          | 115      | editor) .....                                              |     |
| route aggregation .....                                         | 105      | protecting against DoS attacks (configuration .....        | 347 |
| route distinguishers .....                                      |          | editor) .....                                              |     |
| description .....                                               | 210      | protecting against untrusted services and .....            |     |
| formats for .....                                               | 241      | protocols (configuration editor) .....                     | 344 |
| route injection .....                                           | 322      | routing information base <i>See</i> routing table .....    |     |
| route list match types .....                                    | 320      | Routing Information Protocol <i>See</i> RIP .....          |     |
| route manipulation actions, routing policies .....              | 297      | routing instance .....                                     |     |
| route redistribution .....                                      | 322      | VPN configuration .....                                    | 241 |
| route reflectors <i>See</i> BGP route reflectors .....          |          | VPN route target .....                                     | 241 |

- VRF instances..... 210
  - VRF table..... 241
  - routing policies
    - actions ..... 296
    - applying ..... 298
    - BGP routing policy (configuration editor)..... 185
    - components..... 293
    - configuration tasks ..... 318
    - default actions ..... 298
    - export statement ..... 298
    - final actions ..... 298
    - forwarding class with source and destination.... 324
    - grouping source and destination prefixes ..... 324
    - import statement ..... 298
    - injecting routes from one protocol into another .. 322
    - Layer 2 VPN export policy..... 245
    - Layer 2 VPN import policy ..... 244
    - Layer 3 VPNs..... 247
    - making BGP routes less preferable ..... 325
    - match conditions ..... 294
    - overview ..... 293
    - policy name..... 319
    - preparation..... 318
    - prepending AS paths ..... 325
    - reducing update messages with flap damping... 327
    - rejecting invalid routes ..... 320
    - RIP routing policy (configuration editor) ..... 143
    - route redistribution ..... 322
    - route-flap damping ..... 327
    - terms ..... 294
    - terms, creating..... 319
    - VPN configuration ..... 243
  - routing protocols *See* protocols
  - routing solutions
    - BGP confederations, for scaling problems..... 188
    - BGP route reflectors, for scaling problems ..... 185
    - BGP scaling techniques..... 122
    - controlling designated router election ..... 170
    - controlling OSPF route cost ..... 168
    - controlling OSPF route selection..... 167
    - controlling RIP traffic with the incoming metric.. 146
    - controlling RIP traffic with the outgoing metric.. 147
    - CoS with DiffServ ..... 307, 371
    - designated router, to reduce flooding ..... 112
    - directing BGP traffic by local preference ..... 120
    - filtering unwanted services and protocols..... 344
    - firewall filters and NAT ..... 298, 331
    - handling packet fragments..... 343
    - handling packet fragments (configuration editor)..... 353
    - making BGP routes less preferable ..... 325
    - MPLS traffic engineering ..... 213
    - multicast administrative scoping ..... 275
    - multicast reverse-path forwarding (RPF) ..... 274
    - multicast shortest-path tree (SPT)..... 275
    - NSSAs, to control route advertisement ..... 115
    - path cost metrics, for packet flow control *See* path cost metrics
    - point-to-point sessions over Ethernet..... 77
    - poison reverse, for traffic reduction ..... 109
    - preventing multicast routing loops ..... 274
    - protecting against DoS attacks ..... 347
    - reducing update messages with flap damping ... 327
    - rejecting invalid routes ..... 320
    - routing policies ..... 293, 317
    - securing OSPF routing (OSPFv2 only) ..... 169
    - split horizon, for traffic reduction..... 109
    - static route control techniques ..... 128
    - stub areas, to control route advertisement ..... 115
    - VPNs ..... 227
  - routing table
    - controlling static routes in..... 128, 135
    - description ..... 103
    - displaying static routes in ..... 137
    - RPF group, for multicast ..... 284
    - sample distance-vector routing ..... 108
    - updates, limitations in RIP ..... 110
    - verifying for RPF..... 288
    - verifying LDP-signaled LSPs..... 222
    - verifying OSPF routes ..... 173
    - verifying RSVP-signaled LSPs ..... 224
  - RP (rendezvous point)
    - static..... 282
    - verifying..... 287
  - RPF (reverse-path forwarding)
    - description ..... 274
    - routing table group ..... 284
    - verifying the routing table..... 288
  - RSVP (Resource Reservation Protocol)
    - and OSPF for VPNs ..... 239
    - bandwidth reservation ..... 205
    - CSPF..... 207
    - disabling CSPF ..... 220
    - EROs ..... 205
    - fundamentals ..... 205
    - link coloring..... 207
    - overview ..... 214
    - requirements..... 214
    - RSVP-signaled LSPs..... 217
    - verifying LSPs..... 224
    - verifying neighbors ..... 223
    - verifying sessions..... 224
    - verifying the routing table on the entry router ... 224
  - RSVP neighbors, verifying..... 223
  - RSVP-signaled LSP *See* RSVP
  - run command..... 33
- S**
- S,G notation, for multicast forwarding states..... 273

|                                                             |     |                                                     |     |
|-------------------------------------------------------------|-----|-----------------------------------------------------|-----|
| samples .....                                               | 89  | CoS overview .....                                  | 307 |
| firewall filter configurations .....                        | 360 | CoS with DiffServ .....                             | 371 |
| PPPoE over ATM-for-ADSL configuration .....                 | 89  | CPE, with PPPoE .....                               | 77  |
| <i>See also</i> networks; topology                          |     | <i>See also</i> PPPoE                               |     |
| SAP (Session Announcement Protocol)                         |     | firewall filter overview .....                      | 298 |
| description .....                                           | 277 | firewall filters .....                              | 331 |
| session announcements .....                                 | 280 | IPSec tunnels .....                                 | 251 |
| verifying .....                                             | 286 | MPLS for VPNs overview .....                        | 197 |
| saving, configuration files .....                           | 37  | MPLS traffic engineering .....                      | 213 |
| scaling BGP <i>See</i> BGP confederations; BGP route        |     | multicast .....                                     | 279 |
| reflectors                                                  |     | multicast overview .....                            | 269 |
| schedulers                                                  |     | NAT .....                                           | 331 |
| assigning resources .....                                   | 391 | network interfaces .....                            | 41  |
| default settings .....                                      | 312 | OSPF routing .....                                  | 155 |
| description .....                                           | 311 | PPPoE .....                                         | 77  |
| mapping to forwarding classes .....                         | 395 | RIP routing .....                                   | 139 |
| sample mappings .....                                       | 395 | routing policies .....                              | 317 |
| sample schedulers .....                                     | 391 | routing policy overview .....                       | 293 |
| scheduling a commit .....                                   | 31  | routing protocols overview .....                    | 97  |
| scoping, administrative .....                               | 275 | static routing .....                                | 127 |
| SDP (Session Discovery Protocol)                            |     | VPNs .....                                          | 227 |
| description .....                                           | 277 | Session Announcement Protocol <i>See</i> SAP; SDP   |     |
| session announcements .....                                 | 280 | sessions                                            |     |
| verifying .....                                             | 286 | announcements, multicast .....                      | 280 |
| security                                                    |     | BGP session establishment .....                     | 118 |
| IPSec tunnels .....                                         | 251 | BGP session maintenance .....                       | 118 |
| MD5 authentication for OSPF .....                           | 170 | LDP, verifying .....                                | 221 |
| MD5 authentication for RIPv2 .....                          | 150 | limit on PPPoE sessions .....                       | 81  |
| password authentication for OSPFv2 .....                    | 170 | PPPoE .....                                         | 81  |
| password authentication for RIPv2 .....                     | 149 | PPPoE, reconnection time .....                      | 86  |
| security association <i>See</i> IPSec security associations |     | RSVP, verifying .....                               | 224 |
| serial ports                                                |     | shortest path first algorithm .....                 | 111 |
| CHAP .....                                                  | 62  | shortest-path tree .....                            | 275 |
| clock rate .....                                            | 63  | show bgp group command .....                        | 192 |
| clocking .....                                              | 63  | explanation .....                                   | 192 |
| clocking, inverting the transmit clock .....                | 63  | show bgp neighbor command .....                     | 190 |
| configuring .....                                           | 60  | explanation .....                                   | 191 |
| encapsulation type .....                                    | 62  | show bgp summary command .....                      | 192 |
| line speed .....                                            | 63  | explanation .....                                   | 193 |
| logical interfaces .....                                    | 62  | show class-of-service adaptive-shaper command ..... | 403 |
| service classes, corresponding DSCPs .....                  | 308 | show class-of-service interface command .....       | 403 |
| service sets                                                |     | show cli history command .....                      | 34  |
| for IPSec tunnels .....                                     | 256 | show command .....                                  | 25  |
| for NAT rules .....                                         | 341 | show firewall command .....                         | 360 |
| for stateful firewall filters .....                         | 341 | show firewall filter protect-RE command .....       | 366 |
| service types, naming for PPPoE .....                       | 86  | show firewall log command .....                     | 365 |
| services interfaces                                         |     | explanation .....                                   | 366 |
| applying a NAT rule to (configuration editor) .....         | 341 | show igmp interface command .....                   | 286 |
| applying a stateful firewall filter to (configuration       |     | explanation .....                                   | 287 |
| editor) .....                                               | 341 | show interfaces command .....                       | 89  |
| for IPSec tunnels .....                                     | 255 | show interfaces detail command .....                | 71  |
| Services Router                                             |     | show interfaces extensive command .....             | 72  |
| as a PPPoE client .....                                     | 78  | explanation, for ATM-for-ADSL interfaces .....      | 74  |
| BGP routing .....                                           | 177 | show interfaces lo0 command .....                   | 359 |
| configuration tools .....                                   | 3   | show interfaces ppo command .....                   | 90  |

- show ldp neighbor command ..... 220
  - explanation ..... 221
- show ldp session detail command ..... 221
  - explanation ..... 222
- show multicast rpf command ..... 288
  - explanation ..... 288
- show ospf interface command ..... 171
  - explanation ..... 172
- show ospf neighbor command ..... 172
  - explanation ..... 173
- show ospf route command ..... 174
  - results ..... 174
- show pim interface command ..... 287
  - explanation ..... 287
- show pim rps command ..... 287
  - explanation ..... 288
- show pppoe interfaces command ..... 91
- show pppoe statistics command ..... 92
- show pppoe version command ..... 92
- show rip neighbor command ..... 151
  - explanation ..... 151
- show rip statistics command ..... 151
- show route summary command ..... 368–369
  - explanation ..... 368, 370
- show route table inet.3 command ..... 222, 224
  - explanation ..... 222, 225
- show route terse command ..... 137
  - explanation ..... 138
- show rsvp neighbor command ..... 223
  - explanation ..... 223
- show rsvp session detail command ..... 224
  - explanation ..... 224
- show sap listen command ..... 286, 402
  - explanation ..... 286, 403
- show services command ..... 360
- show services ipsec-vpn ipsec statistics command ... 265
  - explanation ..... 265
- show system reboot command ..... 34
- signaling protocols ..... 213
  - overview ..... 204
  - VPNs ..... 237
  - See also* LDP; MPLS traffic engineering; RSVP
- single-area network, OSPF ..... 161
- source-specific multicast ..... 276
- sp-0/0/0
  - for IPsec tunnels (configuration editor) ..... 255
  - no stateful firewall filters ..... 300
- sparse mode *See* multicast routing modes
- SPF (shortest path first) algorithm ..... 111
- split horizon technique ..... 109
- SPT (shortest-path tree) ..... 275
- ssh command ..... 367
  - explanation ..... 368
- stateful firewall filters
  - actions ..... 301
  - applying to an interface (configuration editor) ... 341
  - automatic discard rule ..... 299
  - configuration editor ..... 336, 338
  - configuration overview ..... 300
  - description ..... 299
  - do not apply to sp-0/0/0 ..... 300
  - enabling (Quick Configuration) ..... 335
  - for IPsec tunnels (configuration editor) ..... 260
  - junos-algs-outbound default group ..... 300
  - match conditions ..... 301
  - preparation ..... 332
  - Quick Configuration ..... 332
  - sample rules ..... 337
  - untrusted network ..... 300
  - verifying ..... 364
  - verifying actions ..... 367
- stateless firewall filters
  - actions and action modifiers ..... 306
  - applying to an interface (configuration editor) ... 358
  - automatic discard rule ..... 299, 302
  - bit-field logical operators ..... 306
  - description ..... 299
  - handling packet fragments ..... 343
  - handling packet fragments (configuration editor) ..... 353
  - match conditions ..... 303
  - planning ..... 302, 343
  - policers for ..... 349
  - preparation ..... 332
  - protecting the Routing Engine against ICMP floods (configuration editor) ..... 347
  - protecting the Routing Engine against TCP floods (configuration editor) ..... 347
  - protecting the Routing Engine against untrusted protocols (configuration editor) ..... 344
  - protecting the Routing Engine against untrusted services (configuration editor) ..... 344
  - sample terms, to filter fragments ..... 354
  - sample terms, to filter services and protocols ... 344
  - sample terms, to protect against DoS attacks ... 348
  - typical, planning ..... 343
- statements
  - adding or modifying ..... 26
  - copying ..... 27
  - deactivating ..... 29
  - deleting ..... 26
  - replacing ..... 35
- static LSPs ..... 203
- static routes
  - configuring basic routes (configuration editor) ... 132
  - controlling ..... 128
  - controlling in routing and forwarding tables .... 135
  - default properties ..... 129
  - default properties, setting ..... 136
  - defining route selection ..... 133

|                                                 |     |                                                     |        |
|-------------------------------------------------|-----|-----------------------------------------------------|--------|
| preferences .....                               | 128 | data inversion .....                                | 56     |
| preventing readvertisement .....                | 129 | encapsulation type .....                            | 55     |
| qualified next hops .....                       | 128 | fractional, channel number .....                    | 47     |
| Quick Configuration .....                       | 130 | frame checksum .....                                | 57     |
| rejecting passive traffic .....                 | 129 | framing .....                                       | 56     |
| requirements .....                              | 130 | logical interfaces .....                            | 55     |
| route retention .....                           | 129 | MTU .....                                           | 56     |
| sample preferred path .....                     | 134 | time slots .....                                    | 56     |
| sample stub network .....                       | 132 |                                                     |        |
| verifying .....                                 | 137 | T3 ports                                            |        |
| Static Routes page .....                        | 130 | C-bit parity .....                                  | 60     |
| field summary .....                             | 131 | cable length .....                                  | 60     |
| static routing                                  |     | CHAP .....                                          | 59     |
| default gateway .....                           | 131 | clocking .....                                      | 60     |
| description .....                               | 104 | configuring .....                                   | 57     |
| overview .....                                  | 127 | encapsulation type .....                            | 59     |
| <i>See also</i> static routes                   |     | frame checksum .....                                | 60     |
| static RP router .....                          | 282 | framing .....                                       | 60     |
| <i>See also</i> RP                              |     | logical interfaces .....                            | 59     |
| statistics                                      |     | MTU .....                                           | 60, 62 |
| ATM-for-ADSL interfaces .....                   | 75  | TCP policers .....                                  | 349    |
| firewall filters .....                          | 366 | technical support                                   |        |
| IPSec tunnels .....                             | 265 | contacting JTAC .....                               | xxii   |
| PPPoE .....                                     | 92  | TED <i>See</i> traffic engineering database         |        |
| RIP .....                                       | 151 | telnet command .....                                | 368    |
| status command .....                            | 22  | explanation .....                                   | 369    |
| status, link states, verifying .....            | 70  | terminology                                         |        |
| strict hops, RSVP .....                         | 206 | configuration .....                                 | 3      |
| stub areas                                      |     | CoS .....                                           | 291    |
| area ID (configuration editor) .....            | 163 | firewall filters .....                              | 291    |
| area ID (Quick Configuration) .....             | 158 | MPLS .....                                          | 197    |
| area type (Quick Configuration) .....           | 159 | multicast .....                                     | 269    |
| controlling OSPF route cost .....               | 169 | network interfaces .....                            | 41     |
| creating (configuration editor) .....           | 165 | PPPoE .....                                         | 77     |
| description .....                               | 115 | routing .....                                       | 97     |
| example .....                                   | 116 | routing policies .....                              | 291    |
| sample topology .....                           | 165 | VPNs .....                                          | 197    |
| sub-ASs, BGP .....                              | 125 | terms                                               |        |
| subautonomous systems, BGP .....                | 125 | firewall filter, for multifield classifier .....    | 375    |
| subnetworks                                     |     | in a routing policy .....                           | 294    |
| description .....                               | 102 | in a routing policy, creating .....                 | 319    |
| multicast leaves and branches .....             | 272 | three-way handshake .....                           | 112    |
| route aggregation .....                         | 106 | through route list match type .....                 | 321    |
| Summary Quick Configuration page .....          | 7   | time slots                                          |        |
| support, technical <i>See</i> technical support |     | E1 .....                                            | 51     |
| swap and push label operation .....             | 203 | number in interface name .....                      | 47     |
| swap label operation .....                      | 202 | T1 .....                                            | 56     |
| syntax conventions .....                        | xix | to statement, routing policy match conditions ..... | 294    |
|                                                 |     | top command .....                                   | 25     |
|                                                 |     | topology                                            |        |
| <b>T</b>                                        |     | PPPoE session on an ATM-for-ADSL loop .....         | 80     |
| T1 ports                                        |     | PPPoE session on an Ethernet loop .....             | 79     |
| cable length .....                              | 57  | sample BGP AS path .....                            | 121    |
| CHAP .....                                      | 55  | sample BGP confederation .....                      | 189    |
| clocking .....                                  | 56  | sample BGP confederations .....                     | 126    |
| configuring .....                               | 53  | sample BGP external and internal links .....        | 184    |



- sample BGP local preference use ..... 120
  - sample BGP MED use ..... 122
  - sample BGP peer network ..... 182
  - sample BGP peer session ..... 117
  - sample BGP route reflector (one cluster) ... 123, 186
  - sample BGP route reflectors (cluster of clusters) .. 125
  - sample BGP route reflectors (multiple clusters) .. 124
  - sample distance-vector routing ..... 108
  - sample LSP network ..... 201
  - sample multiarea OSPF routing ..... 114
  - sample OSPF backbone area ..... 115
  - sample OSPF multiarea network ..... 162
  - sample OSPF network ..... 173
  - sample OSPF network with stubs and NSSAs ..... 116
  - sample OSPF single-area network ..... 161
  - sample OSPF stub areas and NSSAs ..... 165
  - sample poison reverse routing ..... 110
  - sample RIP network ..... 143
  - sample RIP network with incoming metric ..... 146
  - sample RIP network with outgoing metric ..... 148
  - sample route advertisement ..... 105
  - sample route aggregation ..... 106
  - sample router network ..... 103
  - sample RSVP-signaled LSP ..... 206
  - sample split horizon routing ..... 109
  - sample static route ..... 104
  - sample static route, preferred path ..... 134
  - sample stub network for static routes ..... 132
  - sample unidirectional routing ..... 111
  - sample VPN ..... 228
  - topology database, OSPF ..... 155
  - Traceroute page
    - results for OSPF ..... 175
    - results for RIP ..... 153
  - traceroute source bypass-routing gateway
    - command ..... 222
    - explanation ..... 223
  - traffic
    - controlling with incoming RIP metric ..... 146
    - controlling with outgoing RIP metric ..... 147
    - incoming, securing ..... 252
    - outgoing, securing ..... 252
  - traffic engineering *See* MPLS traffic engineering
  - traffic engineering database
    - CSPF constraints on path selection ..... 207
    - CSPF rules for path selection ..... 207
    - link coloring for CSPF path selection ..... 207
  - transit interfaces
    - LDP-signaled LSPs for ..... 215
    - RSVP-signaled LSPs for ..... 217
  - transit routers, in an LSP ..... 201
  - transmit clock source *See* clocking
  - trusted networks, firewall filter protection ..... 298
  - tunnels, through a public network *See* IPSec tunnels; VPNs
  - types, of network interfaces ..... 46
- U**
- untrusted networks, firewall filter actions on ..... 298
  - up command ..... 24
  - uploading a configuration file ..... 14
  - upstream interfaces ..... 272
    - See also* multicast
  - upto route list match type ..... 321
  - UR-2 operating mode ..... 68
  - URLs
    - release notes ..... xvii
- V**
- VCI (virtual channel identifier)
    - ATM-for-ADSL interfaces ..... 69
    - PPPoE over ATM-for-ADSL interfaces ..... 84
  - verification
    - adaptive shaping ..... 403
    - BGP configuration ..... 192
    - BGP groups ..... 191
    - BGP peer reachability ..... 193
    - BGP peers (neighbors) ..... 190
    - configuration syntax ..... 30
    - firewall filter actions ..... 367
    - firewall filter flood protection ..... 368
    - firewall filter handles fragments ..... 369
    - firewall filter operation ..... 365
    - firewall filters ..... 359
    - firewall statistics ..... 366
    - IGMP version ..... 286
    - IPSec tunnel operation ..... 265
    - LDP neighbors ..... 220
    - LDP sessions ..... 221
    - LDP-signaled LSP ..... 222
    - MPLS traffic engineering ..... 220
    - multicast SAP and SDP ..... 286
    - multicast session announcements ..... 402
    - network interfaces ..... 70
    - OSPF host reachability ..... 174
    - OSPF neighbors ..... 172
    - OSPF routes ..... 173
    - OSPF-enabled interfaces ..... 171
    - PIM mode and interface configuration ..... 287
    - PIM RP address ..... 287
    - PIM RPF routing table ..... 288
    - PPPoE interfaces ..... 90
    - PPPoE over ATM-for-ADSL configuration ..... 89
    - PPPoE sessions ..... 91
    - PPPoE statistics ..... 92
    - PPPoE version ..... 92
    - RIP host reachability ..... 152
    - RIP message exchange ..... 151
    - RIP-enabled interfaces ..... 151
    - RSVP neighbors ..... 223

|                                                              |     |                                                               |          |
|--------------------------------------------------------------|-----|---------------------------------------------------------------|----------|
| RSVP sessions .....                                          | 224 | configuration overview .....                                  | 227      |
| RSVP-signaled LSP .....                                      | 224 | configuration task overview .....                             | 230      |
| stateful firewall filters .....                              | 364 | IGPs .....                                                    | 237      |
| static routes in the routing table .....                     | 137 | Layer 2 circuit configuration .....                           | 240      |
| traffic forwarding over LDP-signaled LSPs .....              | 222 | LSP for RSVP .....                                            | 234      |
| VPNs .....                                                   | 248 | MPLS .....                                                    | 233      |
| verifying .....                                              |     | overview .....                                                | 197, 208 |
| ATM-for-ADSL interface properties .....                      | 72  | participating interfaces .....                                | 231      |
| version .....                                                |     | preparation .....                                             | 230      |
| OSPF, supported .....                                        | 112 | protocols for .....                                           | 233      |
| PPPoE, verifying .....                                       | 92  | route distinguishers .....                                    | 210, 241 |
| RIP, supported .....                                         | 107 | route target .....                                            | 241      |
| View Configuration Text page .....                           | 13  | route targets .....                                           | 211      |
| virtual channel identifier <i>See</i> VCI .....              |     | routing information .....                                     | 210      |
| virtual channels, applying CoS rules to logical .....        |     | routing instance .....                                        | 241      |
| interfaces .....                                             | 397 | routing policies .....                                        | 243      |
| virtual circuit ID, for Layer 2 circuits .....               | 240 | routing requirements .....                                    | 209      |
| virtual link, through the backbone area .....                | 114 | sample topology .....                                         | 228      |
| virtual path identifier (VPI), PPPoE over ATM-for-ADSL ..... |     | signaling protocols .....                                     | 237      |
| interfaces .....                                             | 84  | tunneling process .....                                       | 209      |
| virtual private networks <i>See</i> VPNs .....               |     | types .....                                                   | 211      |
| VPI, PPPoE over ATM-for-ADSL interfaces .....                | 84  | verifying connectivity .....                                  | 248      |
| VPN routing and forwarding (VRF) instances .....             | 210 | VRF instances .....                                           | 210      |
| VPN routing and forwarding table <i>See</i> VRF table .....  |     | VRF table <i>See</i> VRF table .....                          |          |
| VPNs (virtual private networks) .....                        | 227 | <i>See also</i> Layer 2 circuits; Layer 2 VPNs: Layer 3 ..... |          |
| AS number .....                                              | 236 | VPNs; MPLS .....                                              |          |
| basic Layer 2 circuit description .....                      | 229 | VRF (VPN routing and forwarding) table .....                  | 241      |
| basic Layer 2 VPN description .....                          | 228 | route targets .....                                           | 211      |
| basic Layer 3 VPN description .....                          | 229 | VRF instances .....                                           | 210      |
| BGP .....                                                    | 235 | VRF instances .....                                           | 210      |
| components .....                                             | 208 |                                                               |          |