Junos® OS

RIP Feature Guide

Modified: 2017-11-21

Copyright © 2017, Juniper Networks, Inc.

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

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

Junos® OS RIP Feature Guide

Copyright © 2017 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

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

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at http://www.juniper.net/support/eula/. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

About the Documentation xiii
Documentation and Release Notes xiii
Supported Platforms xiii
Using the Examples in This Manual xiv
Merging a Full Example xiv
Merging a Snippet xv
Documentation Conventions xv
Documentation Feedback xvii
Requesting Technical Support xviii
Self-Help Online Tools and Resources xviii
Opening a Case with JTAC xviii

Part 1 Overview

Chapter 1 **Introduction to RIP 3**

RIP Overview 3
Distance-Vector Routing Protocols 3
RIP Protocol Overview 4
RIP Packets 5
Maximizing Hop Count 6
Split Horizon and Poison Reverse Efficiency Techniques 6
Limitations of Unidirectional Connectivity 7
Supported RIP and RIPng Standards 8
Routing Protocol Process Overview 8

Chapter 2 **Introduction to RIPng 11**

RIPng Overview 11
RIPng Protocol Overview 11
RIPng Standards 12
RIPng Packets 12
Supported RIP and RIPng Standards 13

Part 2 Configuring RIP

Chapter 3 **Configuring a Basic RIP Network 17**

RIP Configuration Overview 17
Understanding Basic RIP Routing 18
Example: Configuring a Basic RIP Network 18
Configuring a RIP Network for EX Series Switches (J-Web Procedure) 25

Copyright © 2017, Juniper Networks, Inc.

Chapter 4 Configuring Authentication for RIP Routes 29
Understanding RIP Authentication 29
Enabling Authentication with MD5 Authentication (CLI Procedure) 30
Enabling Authentication with Plain-Text Passwords (CLI Procedure) 30
Example: Configuring Route Authentication for RIP 31
Chapter 5 Configuring RIP Timers and Demand Circuits 37
Understanding RIP Timers 37
Example: Configuring RIP Timers 38
RIP Demand Circuits Overview 44
RIP Demand Circuit Packets 45
Timers Used by RIP Demand Circuits 46
Example: Configuring RIP Demand Circuits 47
Chapter 6 Configuring BFD for RIP 51
Understanding BFD for RIP 51
Example: Configuring BFD for RIP 52
Understanding BFD Authentication for RIP 57
BFD Authentication Algorithms 58
Security Authentication Keychains 59
Strict Versus Loose Authentication 59
Example: Configuring BFD Authentication for RIP 59
Chapter 7 Using Metrics to Control Traffic in a RIP Network 67
Understanding Traffic Control with Metrics in a RIP Network 67
Example: Controlling Traffic in a RIP Network with an Incoming Metric 68
Example: Controlling Traffic in a RIP Network with an Outgoing Metric 70
Example: Configuring the Metric Value Added to Imported RIP Routes 72
Chapter 8 Configuring Point-to-Multipoint RIP Networks 77
Configuring Point-to-Multipoint RIP Networks Overview 77
Example: Configuring Point-to-Multipoint RIP Networks 78
Chapter 9 Configuring RIP Import Policy 85
Understanding RIP Import Policy 85
Example: Applying Policies to RIP Routes Imported from Neighbors 85
Chapter 10 Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets 93
Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets 93
Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets . . . 93
Verifying the Exchange of RIP Messages 97
Chapter 11 Redistributing Routes Between Two RIP Instances 99
Understanding Route Redistribution Among RIP Instances 99
Example: Redistributing Routes Between Two RIP Instances 100
Chapter 12 Tracing RIP Protocol Traffic 107
Understanding RIP Trace Operations 107
Example: Tracing RIP Protocol Traffic 108

Copyright © 2017, Juniper Networks, Inc.

Part 3 Configuring RIPng

Chapter 13 Configuring a Basic RIPng Network . . . 115

Understanding Basic RIPng Routing . . . 115
Example: Configuring a Basic RIPng Network . . . 115

Chapter 14 Applying Policies to RIPng Routes . . . 123

Understanding RIPng Import Policies to Filter Routes . . . 123
Example: Applying Policies to RIPng Routes Imported from Neighbors . . . 123
Example: Testing a Routing Policy with Complex Regular Expressions . . . 129

Chapter 15 Configuring Traffic Control with Metrics in a RIPng Network . . . 137

Understanding RIPng Traffic Control with Metrics for Optimizing the Path Cost . . . 137
Example: Configuring the Metric Value Added to Imported RIPng Routes to Control the Route Selection Process . . . 138

Chapter 16 Configuring RIPng Timers . . . 145

Understanding RIP Timers . . . 145
Example: Configuring RIPng Update Interval . . . 146

Chapter 17 Tracing RIPng Protocol Traffic . . . 153

Understanding RIPng Protocol Traffic Trace Operations . . . 153
Example: Tracing Global Routing Protocol Operations . . . 154
Example: Tracing RIPng Protocol Traffic . . . 158

Part 4 Troubleshooting

Chapter 18 Troubleshooting Network Issues . . . 165

Working with Problems on Your Network . . . 165
Isolating a Broken Network Connection . . . 166
Identifying the Symptoms of a Broken Network Connection . . . 167
Isolating the Causes of a Network Problem . . . 168
Taking Appropriate Action for Resolving the Network Problem . . . 169
Evaluating the Solution to Check Whether the Network Problem Is Resolved . . . 170

Chapter 19 Monitoring and Troubleshooting a RIP Network . . . 173

Monitoring RIP Routing Information . . . 173
Verifying a RIP Configuration . . . 174
Verifying the RIP-Enabled Interfaces . . . 174
Verifying Reachability of All Hosts in the RIP Network . . . 175

Part 5 Configuration Statements and Operational Commands

Chapter 20 RIP Configuration Statements . . . 179

any-sender . . . 180
authentication-key (Protocols RIP) . . . 181
authentication-type (Protocols RIP) . . . 182
bfd-liveness-detection (Protocols RIP) . . . 183
check-zero . . . 186
demand-circuit (Protocols RIP) . . . 187
dynamic-peers . . . 188

Copyright © 2017, Juniper Networks, Inc.

export . . . 189
graceful-restart (Protocols RIP) . . . 190
group (Protocols RIP) . . . 191
holddown (Protocols RIP) . . . 193
import (Protocols RIP) . . . 194
interface-type (Protocols RIP) . . . 195
max-retrans-time . . . 196
message-size . . . 197
metric-in (Protocols RIP) . . . 198
metric-out . . . 199
neighbor . . . 200
peer (Protocols RIP) . . . 201
preference (Protocols RIP) . . . 202
receive (Protocols RIP) . . . 203
rib-group (Protocols RIP) . . . 204
rip . . . 205
route-timeout (Protocols RIP) . . . 206
routing-instances (Multiple Routing Entities) . . . 207
send (Protocols RIP) . . . 209
traceoptions (Protocols RIP) . . . 210
update-interval (Protocols RIP) . . . 213

Chapter 21 RIPng Configuration Statements . . . 215

export (Protocols RIPng) . . . 216
graceful-restart (Protocols RIPng) . . . 217
group (Protocols RIPng) . . . 218
holddown (Protocols RIPng) . . . 219
import (Protocols RIPng) . . . 220
metric-in (Protocols RIPng) . . . 221
metric-out (Protocols RIPng) . . . 222
neighbor (Protocols RIPng) . . . 223
preference (Protocols RIPng) . . . 224
receive (Protocols RIPng) . . . 225
ripng . . . 226
route-timeout (Protocols RIPng) . . . 227
routing-instances (Multiple Routing Entities) . . . 228
send (Protocols RIPng) . . . 229
traceoptions (Protocols RIPng) . . . 230
update-interval (Protocols RIPng) . . . 232

Chapter 22 RIP Operational Commands . . . 233

clear rip general-statistics . . . 235
clear rip statistics . . . 236
restart . . . 238
show policy . . . 248
show policy conditions . . . 251
show rip general-statistics . . . 253
show rip neighbor . . . 255
show rip statistics . . . 257
show route . . . 260

Copyright © 2017, Juniper Networks, Inc.

show route active-path . . . 281
show route advertising-protocol . . . 286
show route all . . . 292
show route best . . . 294
show route brief . . . 297
show route detail . . . 299
show route exact . . . 319
show route export . . . 321
show route extensive . . . 324
show route forwarding-table . . . 342
show route hidden . . . 360
show route inactive-path . . . 363
show route instance . . . 367
show route next-hop . . . 375
show route output . . . 381
show route protocol . . . 386
show route receive-protocol . . . 399
show route table . . . 409
show route terse . . . 447
test policy . . . 450

Chapter 23 RIPng Operational Commands . . . 453

clear ripng general-statistics . . . 455
clear ripng statistics . . . 456
restart . . . 457
show policy . . . 467
show policy conditions . . . 470
show ripng general-statistics . . . 472
show ripng neighbor . . . 474
show ripng statistics . . . 476
show route . . . 479
show route active-path . . . 500
show route advertising-protocol . . . 505
show route all . . . 511
show route best . . . 513
show route brief . . . 516
show route detail . . . 518
show route exact . . . 538
show route export . . . 540
show route extensive . . . 543
show route forwarding-table . . . 561
show route hidden . . . 579
show route inactive-path . . . 582
show route instance . . . 586
show route next-hop . . . 594
show route output . . . 600
show route protocol . . . 605
show route receive-protocol . . . 618
show route table . . . 628

Copyright © 2017, Juniper Networks, Inc.

show route terse 666
test policy 669

Copyright © 2017, Juniper Networks, Inc.

List of Figures

Part 1 Overview

Chapter 1 Introduction to RIP . . . 3

Figure 1: Distance-Vector Protocol . . . 4
Figure 2: Split Horizon Example . . . 6
Figure 3: Poison Reverse Example . . . 7
Figure 4: Limitations of Unidirectional Connectivity . . . 7

Part 2 Configuring RIP

Chapter 3 Configuring a Basic RIP Network . . . 17

Figure 5: Sample RIP Network Topology . . . 19

Chapter 4 Configuring Authentication for RIP Routes . . . 29

Figure 6: RIP Authentication Network Topology . . . 32

Chapter 5 Configuring RIP Timers and Demand Circuits . . . 37

Figure 7: RIP Timers Network Topology . . . 39

Chapter 6 Configuring BFD for RIP . . . 51

Figure 8: RIP BFD Network Topology . . . 54
Figure 9: RIP BFD Authentication Network Topology . . . 60

Chapter 7 Using Metrics to Control Traffic in a RIP Network . . . 67

Figure 10: Controlling Traffic in a RIP Network with the Incoming Metric . . . 69
Figure 11: Controlling Traffic in a RIP Network with the Outgoing Metric . . . 71
Figure 12: RIP Incoming Metrics Network Topology . . . 72

Chapter 8 Configuring Point-to-Multipoint RIP Networks . . . 77

Figure 13: Configuring a Point-to-Multipoint RIP Network . . . 79

Chapter 9 Configuring RIP Import Policy . . . 85

Figure 14: RIP Import Policy Network Topology . . . 86

Chapter 10 Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets . . . 93

Figure 15: Sending and Receiving RIPv1 and RIPv2 Packets Network Topology . . . 94

Chapter 11 Redistributing Routes Between Two RIP Instances . . . 99

Figure 16: Redistributing Routes Between RIP Instances Network Topology . . . 101

Chapter 12 Tracing RIP Protocol Traffic . . . 107

Figure 17: RIP Trace Operations Network Topology . . . 109

Copyright © 2017, Juniper Networks, Inc.

Part 3 Configuring RIPng

Chapter 13 Configuring a Basic RIPng Network 115

Figure 18: Sample RIPng Network Topology 116

Chapter 14 Applying Policies to RIPng Routes 123

Figure 19: RIPng Import Policy Network Topology 124
Figure 20: Routing Policy Test for Complex Regular Expressions 131

Chapter 15 Configuring Traffic Control with Metrics in a RIPng Network 137

Figure 21: RIPng Incoming Metrics Network Topology 139

Chapter 16 Configuring RIPng Timers 145

Figure 22: RIPng Timers Network Topology 147

Chapter 17 Tracing RIPng Protocol Traffic 153

Figure 23: RIPng Trace Operations Network Topology 159

Part 4 Troubleshooting

Chapter 18 Troubleshooting Network Issues 165

Figure 24: Process for Diagnosing Problems in Your Network 166
Figure 25: Network with a Problem 166

Copyright © 2017, Juniper Networks, Inc.

List of Tables

About the Documentation xiii

Table 1: Notice Icons xvi
Table 2: Text and Syntax Conventions xvi

Part 2 Configuring RIP

Chapter 3 Configuring a Basic RIP Network 17

Table 3: RIP Routing Configuration Summary 25
Table 4: Edit RIP Global Settings 26

Chapter 4 Configuring Authentication for RIP Routes 29

Table 5: Configuring MD5 RIP Authentication 30
Table 6: Configuring Simple RIP Authentication 31

Chapter 5 Configuring RIP Timers and Demand Circuits 37

Table 7: RIP Demand Circuit Packet Types 45

Part 4 Troubleshooting

Chapter 18 Troubleshooting Network Issues 165

Table 8: Checklist for Working with Problems on Your Network 165

Chapter 19 Monitoring and Troubleshooting a RIP Network 173

Table 9: Summary of Key RIP Routing Output Fields 173

Part 5 Configuration Statements and Operational Commands

Chapter 22 RIP Operational Commands 233

Table 10: show policy Output Fields 249
Table 11: show policy conditions Output Fields 251
Table 12: show rip general-statistics Output Fields 253
Table 13: show rip neighbor Output Fields 256
Table 14: show rip statistics Output Fields 258
Table 15: show route Output Fields 261
Table 16: show route advertising-protocol Output Fields 287
Table 17: show route detail Output Fields 299
Table 18: Next-hop Types Output Field Values 305
Table 19: State Output Field Values 307
Table 20: Communities Output Field Values 309
Table 21: show route export Output Fields 322
Table 22: show route extensive Output Fields 324
Table 23: show route forwarding-table Output Fields 345

Copyright © 2017, Juniper Networks, Inc.

Table 24: show route instance Output Fields 368
Table 25: show route receive-protocol Output Fields 400
Table 26: show route table Output Fields 410
Table 27: Next-hop Types Output Field Values 416
Table 28: State Output Field Values 417
Table 29: Communities Output Field Values 419
Table 30: show route terse Output Fields 447

Chapter 23 RIPng Operational Commands 453

Table 31: show policy Output Fields 468
Table 32: show policy conditions Output Fields 470
Table 33: show ripng general-statistics Output Fields 472
Table 34: show ripng neighbor Output Fields 474
Table 35: show ripng statistics Output Fields 477
Table 36: show route Output Fields 480
Table 37: show route advertising-protocol Output Fields 506
Table 38: show route detail Output Fields 518
Table 39: Next-hop Types Output Field Values 524
Table 40: State Output Field Values 526
Table 41: Communities Output Field Values 528
Table 42: show route export Output Fields 541
Table 43: show route extensive Output Fields 543
Table 44: show route forwarding-table Output Fields 564
Table 45: show route instance Output Fields 587
Table 46: show route receive-protocol Output Fields 619
Table 47: show route table Output Fields 629
Table 48: Next-hop Types Output Field Values 635
Table 49: State Output Field Values 636
Table 50: Communities Output Field Values 638
Table 51: show route terse Output Fields 666

Copyright © 2017, Juniper Networks, Inc.

About the Documentation

- Documentation and Release Notes on page xiii
- Supported Platforms on page xiii
- Using the Examples in This Manual on page xiv
- Documentation Conventions on page xv
- Documentation Feedback on page xvii
- Requesting Technical Support on page xviii

Documentation and Release Notes

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at http://www.juniper.net/techpubs/.

If the information in the latest release notes differs from the information in the documentation, follow the product Release Notes.

Juniper Networks Books publishes books by Juniper Networks engineers and subject matter experts. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration. The current list can be viewed at http://www.juniper.net/books.

Supported Platforms

For the features described in this document, the following platforms are supported:

- ACX Series
- T Series
- MX Series
- M Series
- SRX Series
- PTX Series

Copyright © 2017, Juniper Networks, Inc.

Using the Examples in This Manual

If you want to use the examples in this manual, you can use the **load merge** or the **load merge relative** command. These commands cause the software to merge the incoming configuration into the current candidate configuration. The example does not become active until you commit the candidate configuration.

If the example configuration contains the top level of the hierarchy (or multiple hierarchies), the example is a *full example*. In this case, use the **load merge** command.

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

Merging a Full Example

To merge a full example, follow these steps:

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

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

   ```
   system {
     scripts {
       commit {
         file ex-script.xsl;
       }
     }
   }
   interfaces {
     fxp0 {
       disable;
       unit 0 {
         family inet {
           address 10.0.0.1/24;
         }
       }
     }
   }
   ```

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

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


Copyright © 2017, Juniper Networks, Inc.

Merging a Snippet

To merge a snippet, follow these steps:

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

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

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

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

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

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

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

For more information about the **load** command, see CLI Explorer.

Documentation Conventions

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

Copyright © 2017, Juniper Networks, Inc.

Table 1: Notice Icons

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

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

Table 2: Text and Syntax Conventions

Convention	Description	Examples
Bold text like this	Represents text that you type.	To enter configuration mode, type the **configure** command: user@host> **configure**
`Fixed-width text like this`	Represents output that appears on the terminal screen.	`user@host>` **show chassis alarms** `No alarms currently active`
Italic text like this	• Introduces or emphasizes important new terms. • Identifies guide names. • Identifies RFC and Internet draft titles.	• A policy *term* is a named structure that defines match conditions and actions. • *Junos OS CLI User Guide* • RFC 1997, *BGP Communities Attribute*
Italic text like this	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** *domain-name*

Copyright © 2017, Juniper Networks, Inc.

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

Convention	Description	Examples
Text like this	Represents names of configuration statements, commands, files, and directories; configuration hierarchy levels; or labels on routing platform components.	• To configure a stub area, include the **stub** statement at the **[edit protocols ospf area area-id]** hierarchy level. • The console port is labeled **CONSOLE**.
< > (angle brackets)	Encloses optional keywords or variables.	**stub <default-metric *metric*>;**
| (pipe symbol)	Indicates a choice between the mutually exclusive keywords or variables on either side of the symbol. The set of choices is often enclosed in parentheses for clarity.	**broadcast | multicast** **(*string1* | *string2* | *string3*)**
# (pound sign)	Indicates a comment specified on the same line as the configuration statement to which it applies.	**rsvp { # Required for dynamic MPLS only**
[ ] (square brackets)	Encloses a variable for which you can substitute one or more values.	**community name members [*community-ids*]**

Indention and braces ({ })	Identifies a level in the configuration hierarchy.	[edit]
		routing-options {
		  static {
		  route default {
		  nexthop *address*;
		  retain;
		  }
		  }
		}

; (semicolon)	Identifies a leaf statement at a configuration hierarchy level.	
GUI Conventions		
Bold text like this	Represents graphical user interface (GUI) items you click or select.	• In the Logical Interfaces box, select **All Interfaces**. • To cancel the configuration, click **Cancel**.
> (bold right angle bracket)	Separates levels in a hierarchy of menu selections.	In the configuration editor hierarchy, select **Protocols>Ospf**.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page of the Juniper Networks TechLibrary site at http://www.juniper.net/techpubs/index.html, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at http://www.juniper.net/techpubs/feedback/.

Copyright © 2017, Juniper Networks, Inc.

- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit http://www.juniper.net/support/warranty/.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.juniper.net/
- Find product documentation: http://www.juniper.net/documentation/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: http://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://entitlementsearch.juniper.net/entitlementsearch/

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

Copyright © 2017, Juniper Networks, Inc.

For international or direct-dial options in countries without toll-free numbers, see http://www.juniper.net/support/requesting-support.html.

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

PART 1

Overview

- Introduction to RIP on page 3
- Introduction to RIPng on page 11

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 1

Introduction to RIP

- RIP Overview on page 3
- Supported RIP and RIPng Standards on page 8
- Routing Protocol Process Overview on page 8

RIP Overview

RIP is an interior gateway protocol (IGP) that uses a distance-vector algorithm to determine the best route to a destination, using the hop count as the metric.

In a 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.

NOTE: In general, the term *RIP* refers to RIP version 1 and RIP version 2.

This topic contains the following sections:

- Distance-Vector Routing Protocols on page 3
- RIP Protocol Overview on page 4
- RIP Packets on page 5
- Maximizing Hop Count on page 6
- Split Horizon and Poison Reverse Efficiency Techniques on page 6
- Limitations of Unidirectional Connectivity on page 7

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 1 on page 4 shows how distance-vector routing works.

Copyright © 2017, Juniper Networks, Inc.

Figure 1: Distance-Vector Protocol

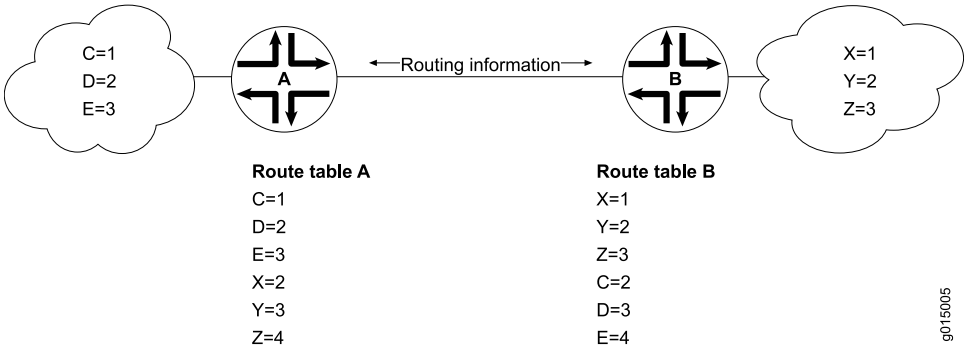

In Figure 1 on page 4, Routers A and B have RIP enabled on adjacent interfaces. Router A has known RIP neighbors Routers C, D, and E, which are 1, 2, and 3 hops away, respectively. Router B has known RIP neighbors Routers X, Y, and Z, which are 1, 2, and 3 hops away, respectively. Every 30 seconds, each 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.

RIP Protocol Overview

The RIP IGP uses the Bellman-Ford, or *distance-vector*, algorithm to determine the best route to a destination. RIP uses the hop count as the metric. RIP enables hosts and routers to exchange information for computing routes through an IP-based network. RIP is intended to be used as an IGP in reasonably homogeneous networks of moderate size.

The Junos® operating system (Junos OS) supports RIP versions 1 and 2.

NOTE: **RIP is not supported for multipoint interfaces.**

RIP version 1 packets contain the minimal information necessary to route packets through a network. However, this version of RIP does not support authentication or subnetting.

RIP uses User Datagram Protocol (UDP) port 520.

RIP has the following architectural limitations:

Copyright © 2017, Juniper Networks, Inc.

- The longest network path cannot exceed 15 hops (assuming that each network, or hop, has a cost of 1).
- RIP depends on counting to infinity to resolve certain unusual situations—When the network consists of several hundred routers, and when a routing loop has formed, the amount of time and network bandwidth required to resolve a next hop might be great.
- RIP uses only a fixed metric to select a route. Other IGPs use additional parameters, such as measured delay, reliability, and load.

RIP Packets

RIP packets contain the following fields:

- Command—Indicates whether the packet is a request or response message. Request messages seek information for the router's routing table. Response messages are sent periodically and also when a request message is received. Periodic response messages are called *update messages*. Update messages contain the command and version fields and 25 destinations (by default), each of which includes the destination IP address and the metric to reach that destination.

NOTE: Beginning with Junos OS Release 11.1, three additional command field types are available to support RIP demand circuits. When you configure an interface for RIP demand circuits, the command field indicates whether the packet is an update request, update response, or update acknowledge message. Neighbor interfaces send updates on demand, not periodically. These command field types are only valid on interfaces configured for RIP demand circuits. For more detailed information, see "RIP Demand Circuits Overview" on page 44.

- Version number—Version of RIP that the originating router is running.
- Address family identifier—Address family used by the originating router. The family is always IP.
- Address—IP address included in the packet.
- Metric—Value of the metric advertised for the address.
- Mask—Mask associated with the IP address (RIP version 2 only).
- Next hop—IP address of the next-hop router (RIP version 2 only).

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.

Copyright © 2017, Juniper Networks, Inc.

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.

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 (30 seconds per hop). 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*.

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 2 on page 6 shows an example of the split horizon technique.

Figure 2: Split Horizon Example

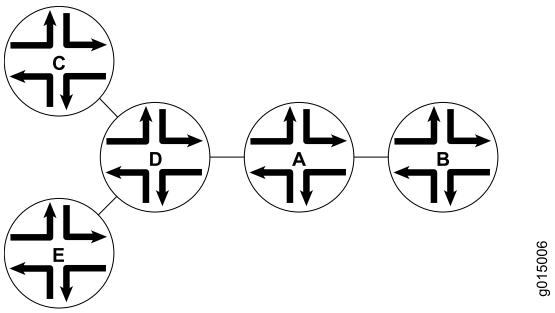

Copyright © 2017, Juniper Networks, Inc.

In Figure 2 on page 6, 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, Router 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, Router 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 3 on page 7 shows an example of the poison reverse technique.

Figure 3: Poison Reverse Example

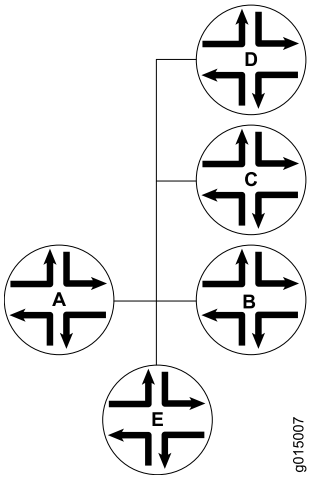

In Figure 3 on page 7, Router A learns through one of its interfaces that routes to Routers C, D, and E are unreachable. Router A readvertises those routes out the same interface as unreachable. The advertisement informs Router B that Routers C, D, and E are definitely not reachable through Router A.

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 4 on page 7 shows, RIP networks are limited by their unidirectional connectivity.

Figure 4: Limitations of Unidirectional Connectivity

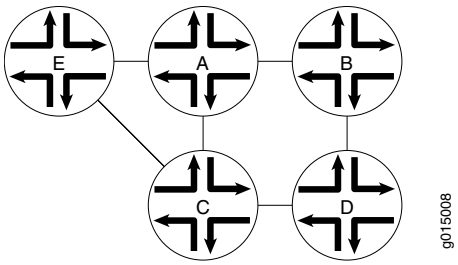

Copyright © 2017, Juniper Networks, Inc.

In Figure 4 on page 7, 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 because of 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.

Related Documentation

- RIP Configuration Overview on page 17
- *Example: Configuring RIP*

Supported RIP and RIPng Standards

Junos OS substantially supports the following RFCs, which define standards for RIP (for IP version 4 [IPv4]) and RIP next generation (RIPng, for IP version 6 [IPv6]).

Junos OS supports authentication for all RIP protocol exchanges (MD5 or simple authentication).

- RFC 1058, *Routing Information Protocol*
- RFC 2080, *RIPng for IPv6*
- RFC 2082, *RIP-2 MD5 Authentication*

  Multiple keys using distinct key IDs are not supported.
- RFC 2453, *RIP Version 2*

The following RFC does not define a standard, but provides information about RIPng. The IETF classifies it as “Informational.”

- RFC 2081, *RIPng Protocol Applicability Statement*

Related Documentation

- *Supported IPv4, TCP, and UDP Standards*
- *Supported IPv6 Standards*
- *Accessing Standards Documents on the Internet*

Routing Protocol Process Overview

Junos OS is based on the FreeBSD Unix operating system. The open source software is modified and hardened to operate in the device’s specialized environment. For example, some executables have been deleted, while other utilities were de-emphasized. Additionally, certain software processes were added to enhance the routing functionality. The result of this transformation is the kernel, the heart of the Junos OS software.

The kernel is responsible for operating multiple processes that perform the actual functions of the device. Each process operates in its own protected memory space, while

Copyright © 2017, Juniper Networks, Inc.

the communication among all the processes is still controlled by the kernel. This separation provides isolation between the processes, and resiliency in the event of a process failure. This is important in a core routing platform because a single process failure does not cause the entire device to cease functioning.

Some of the common software processes include the routing protocol process (rpd) that controls the device's protocols, the device control process (dcd) that controls the device's interfaces, the management process (mgd) that controls user access to the device, the chassis process (chassisd) that controls the device's properties itself, and the Packet Forwarding Engine process (pfed) that controls the communication between the device's Packet Forwarding Engine and the Routing Engine. The kernel also generates specialized processes as needed for additional functionality, such as SNMP, the Virtual Router Redundancy Protocol (VRRP), and Class of Service (CoS).

The routing protocol process is a software process within the Routing Engine software, which controls the routing protocols that run on the device. Its functionality includes all protocol messages, routing table updates, and implementation of routing policies.

The routing protocol process starts all configured routing protocols and handles all routing messages. It maintains one or more routing tables, which consolidate the routing information learned from all routing protocols. From this routing information, the routing protocol process determines the active routes to network destinations and installs these routes into the Routing Engine's forwarding table. Finally, it implements routing policy, which allows you to control the routing information that is transferred between the routing protocols and the routing table. Using routing policy, you can filter and limit the transfer of information as well as set properties associated with specific routes.

Related Documentation

- *show system processes*
- *show task*
- *show task memory*

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 2

Introduction to RIPng

- RIPng Overview on page 11
- Supported RIP and RIPng Standards on page 13

RIPng Overview

The Routing Information Protocol next generation (RIPng) is an interior gateway protocol (IGP) that uses a distance-vector algorithm to determine the best route to a destination, using hop count as the metric. RIPng exchanges routing information used to compute routes and is intended for IP version 6 (IPv6)-based networks. RIPng is disabled by default.

On devices in secure context, IPv6 is disabled. You must enable IPv6 to use RIPng. For instructions, see the *Junos OS Interfaces Configuration Guide for Security Devices*.

This topic contains the following sections:

- RIPng Protocol Overview on page 11
- RIPng Standards on page 12
- RIPng Packets on page 12

RIPng Protocol Overview

The RIPng IGP uses the Bellman-Ford distance-vector algorithm to determine the best route to a destination, using hop count as the metric. RIPng allows hosts and routers to exchange information for computing routes through an IP-based network. RIPng is intended to act as an IGP for moderately-sized autonomous systems.

RIPng is a distinct routing protocol from RIPv2. The Junos OS implementation of RIPng is similar to RIPv2, but has the following differences:

- RIPng does not need to implement authentication on packets.
- Junos OS does not support multiple instances of RIPng.
- Junos OS does not support RIPng routing table groups.

RIPng is a UDP-based protocol and uses UDP port 521.

RIPng has the following architectural limitations:

Copyright © 2017, Juniper Networks, Inc.

- The longest network path cannot exceed 15 hops (assuming that each network, or hop, has a cost of 1).
- RIPng is prone to routing loops when the routing tables are reconstructed. Especially when RIPng is implemented in large networks that consist of several hundred routers, RIPng might take an extremely long time to resolve routing loops.
- RIPng uses only a fixed metric to select a route. Other IGPs use additional parameters, such as measured delay, reliability, and load.

RIPng Standards

RIPng is defined in the following documents:

- RFC 2080, *RIPng for IPv6*
- RFC 2081, *RIPng Protocol Applicability Statement*

To access Internet Requests for Comments (RFCs) and drafts, see the Internet Engineering Task Force (IETF) website.

RIPng Packets

A RIPng packet header contains the following fields:

- Command—Indicates whether the packet is a request or response message. Request messages seek information for the router's routing table. Response messages are sent periodically or when a request message is received. Periodic response messages are called update messages. Update messages contain the command and version fields and a set of destinations and metrics.
- Version number—Specifies the version of RIPng that the originating router is running. This is currently set to Version 1.

The rest of the RIPng packet contains a list of routing table entries consisting of the following fields:

- Destination prefix—128-bit IPv6 address prefix for the destination.
- Prefix length—Number of significant bits in the prefix.
- Metric—Value of the metric advertised for the address.
- Route tag—A route attribute that must be advertised and redistributed with the route. Primarily, the route tag distinguishes external RIPng routes from internal RIPng routes when routes must be redistributed across an exterior gateway protocol (EGP).

Related Documentation

- RIP Overview on page 3
- Example: Configuring a Basic RIPng Network on page 115

Copyright © 2017, Juniper Networks, Inc.

Supported RIP and RIPng Standards

Junos OS substantially supports the following RFCs, which define standards for RIP (for IP version 4 [IPv4]) and RIP next generation (RIPng, for IP version 6 [IPv6]).

Junos OS supports authentication for all RIP protocol exchanges (MD5 or simple authentication).

- RFC 1058, *Routing Information Protocol*
- RFC 2080, *RIPng for IPv6*
- RFC 2082, *RIP-2 MD5 Authentication*

  Multiple keys using distinct key IDs are not supported.
- RFC 2453, *RIP Version 2*

The following RFC does not define a standard, but provides information about RIPng. The IETF classifies it as "Informational."

- RFC 2081, *RIPng Protocol Applicability Statement*

Related Documentation

- *Supported IPv4, TCP, and UDP Standards*
- *Supported IPv6 Standards*
- *Accessing Standards Documents on the Internet*

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

PART 2

Configuring RIP

- Configuring a Basic RIP Network on page 17
- Configuring Authentication for RIP Routes on page 29
- Configuring RIP Timers and Demand Circuits on page 37
- Configuring BFD for RIP on page 51
- Using Metrics to Control Traffic in a RIP Network on page 67
- Configuring Point-to-Multipoint RIP Networks on page 77
- Configuring RIP Import Policy on page 85
- Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 93
- Redistributing Routes Between Two RIP Instances on page 99
- Tracing RIP Protocol Traffic on page 107

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 3

Configuring a Basic RIP Network

- RIP Configuration Overview on page 17
- Understanding Basic RIP Routing on page 18
- Example: Configuring a Basic RIP Network on page 18
- Configuring a RIP Network for EX Series Switches (J-Web Procedure) on page 25

RIP Configuration 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, as described in the steps that follow.

To configure a RIP network:

1. Configure network interfaces. See the *Junos OS Interfaces Configuration Guide for Security Devices*.

2. Define RIP groups, which are logical groupings of interfaces, and add interfaces to the groups. Then, configure a routing policy to export directly connected routes and routes learned through RIP routing exchanges. See "Example: Configuring a Basic RIP Network" on page 18.

3. (Optional) Configure metrics to control traffic through the RIP network. See "Example: Controlling Traffic in a RIP Network with an Incoming Metric" on page 68 and "Example: Controlling Traffic in a RIP Network with an Outgoing Metric" on page 70.

4. (Optional) Configure authentication to ensure that only trusted routers participate in the autonomous system's routing. See "Enabling Authentication with Plain-Text Passwords (CLI Procedure)" on page 30 and "Enabling Authentication with MD5 Authentication (CLI Procedure)" on page 30.

Related Documentation

- *Junos OS Feature Support Reference for SRX Series and J Series Devices*
- RIP Overview on page 3
- Verifying a RIP Configuration on page 174

Copyright © 2017, Juniper Networks, Inc.

Understanding Basic RIP Routing

RIP is an interior gateway protocol (IGP) that routes packets within a single autonomous system (AS). By default, RIP does not advertise the subnets that are directly connected through the device'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.

Related Documentation

- RIP Overview on page 3
- Example: Configuring a Basic RIP Network on page 18

Example: Configuring a Basic RIP Network

This example shows how to configure a basic RIP network.

- Requirements on page 18
- Overview on page 18
- Configuration on page 19
- Verification on page 21

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, you configure a basic RIP network, create a RIP group called **rip-group**, and add the directly connected interfaces to the RIP group. Then you configure a routing policy to advertise direct routes using policy statement **advertise-routes-through-rip**.

By default, Junos OS does not advertise RIP routes, not even routes that are learned through RIP. To advertise RIP routes, you must configure and apply an export routing policy that advertises RIP-learned and direct routes.

In Junos OS, you do not need to configure the RIP version. RIP version 2 is used by default.

To use RIP on the device, you must configure RIP on all of the RIP interfaces within the network. Figure 5 on page 19 shows the topology used in this example.

Copyright © 2017, Juniper Networks, Inc.

Figure 5: Sample RIP Network Topology

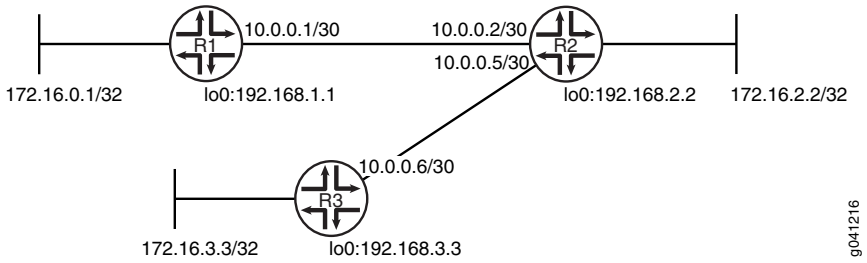

"CLI Quick Configuration" on page 19 shows the configuration for all of the devices in Figure 5 on page 19. The section "Step-by-Step Procedure" on page 20 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter commit from configuration mode.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
```

Copyright © 2017, Juniper Networks, Inc.

```
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a basic RIP network:

1. Configure the network interfaces.

   This example shows multiple loopback interface addresses to simulate attached networks.

   ```
   [edit interfaces]
   user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

   user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
   user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
   ```

2. Create the RIP group and add the interface.

   To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

   ```
   [edit protocols rip group rip-group]
   user@R1# set neighbor fe-1/2/0.1
   ```

3. Create the routing policy to advertise both direct and RIP-learned routes.

   ```
   [edit policy-options policy-statement advertise-routes-through-rip term 1]
   user@R1# set from protocol direct
   user@R1# set from protocol rip
   user@R1# set then accept
   ```

4. Apply the routing policy.

   In Junos OS, you can only apply RIP export policies at the group level.

   ```
   [edit protocols rip group rip-group]
   user@R1# set export advertise-routes-through-rip
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
```


Copyright © 2017, Juniper Networks, Inc.

```
        }
      }
    }
    lo0 {
      unit 1 {
        family inet {
          address 172.16.0.1/32;
          address 192.168.1.1/32;
        }
      }
    }
```

```
user@R1# show protocols
rip {
  group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1;
  }
}
```

```
user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Checking the Routing Table on page 21
- Looking at the Routes That Device R1 Is Advertising to Device R2 on page 22
- Looking at the Routes That Device R1 Is Receiving from Device R2 on page 22
- Verifying the RIP-Enabled Interfaces on page 23
- Verifying the Exchange of RIP Messages on page 23
- Verifying Reachability of All Hosts in the RIP Network on page 24

Checking the Routing Table

Purpose Verify that the routing table is populated with the expected routes..

Action From operational mode, enter the **show route protocol rip** command.

```
user@R1> show route protocol rip
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30        *[RIP/100] 00:59:15, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
```

Copyright © 2017, Juniper Networks, Inc.

```
172.16.2.2/32      *[RIP/100] 02:52:48, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
172.16.3.3/32      *[RIP/100] 00:45:05, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32     *[RIP/100] 02:52:48, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32     *[RIP/100] 00:45:05, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
224.0.0.9/32       *[RIP/100] 00:45:09, metric 1
                          MultiRecv
```

Meaning The output shows that the routes have been learned from Device R2 and Device R3.

If you were to delete the **from protocol rip** condition in the routing policy on Device R2, the remote routes from Device R3 would not be learned on Device R1.

Looking at the Routes That Device R1 Is Advertising to Device R2

Purpose Verify that Device R1 is sending the expected routes.

Action From operational mode, enter the **show route advertising-protocol rip** command.

```
user@R1> show route advertising-protocol rip 10.0.0.1
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.1/32      *[Direct/0] 05:18:26
                    >    via lo0.1
192.168.1.1/32     *[Direct/0] 05:18:25
                    >    via lo0.1
```

Meaning Device R1 is sending routes to its directly connected networks.

Looking at the Routes That Device R1 Is Receiving from Device R2

Purpose Verify that Device R1 is receiving the expected routes.

Action From operational mode, enter the **show route receive-protocol rip** command.

```
user@R1> show route receive-protocol rip 10.0.0.2
inet.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30        *[RIP/100] 02:31:22, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
172.16.2.2/32      *[RIP/100] 04:24:55, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
172.16.3.3/32      *[RIP/100] 02:17:12, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32     *[RIP/100] 04:24:55, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
```


Copyright © 2017, Juniper Networks, Inc.

```
192.168.3.3/32     *[RIP/100] 02:17:12, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
```

Meaning Device R1 is receiving from Device R2 all of Device R2's directly connected networks. Device R1 is also receiving from Device R2 all of Device R3's directly connected networks, which Device R2 learned from Device R3 through RIP.

Verifying the RIP-Enabled Interfaces

Purpose Verify that all RIP-enabled Interfaces are available and active.

Action From operational mode, enter the **show rip neighbor** command.

```
user@R1> show rip neighbor
                    Local  Source          Destination     Send   Receive   In
Neighbor            State  Address         Address         Mode   Mode     Met
--------            -----  -------         -----------     ----   -------  ---
fe-1/2/0.1             Up 10.0.0.1        224.0.0.9       mcast  both     1
```

Meaning The output shows that the RIP-enabled interface on Device R1 is operational.

In general for this command, the output shows a list of the RIP neighbors that are configured on the device. 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 **Local 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 operational mode, enter the **show rip statistics** command.

```
user@R1> show rip statistics
RIPv2 info: port 520; holddown 120s.
    rts learned  rts held down  rqsts dropped  resps dropped
              5              0              0              0

fe-1/2/0.1:  5 routes learned; 2 routes advertised; timeout 180s; update interval
 30s
Counter                         Total   Last 5 min  Last minute
-------                     -----------  ----------  -----------
Updates Sent                       2669          10            2
Triggered Updates Sent                2           0            0
Responses Sent                        0           0            0
Bad Messages                          0           0            0
```

Copyright © 2017, Juniper Networks, Inc.

```
RIPv1 Updates Received            0          0          0
RIPv1 Bad Route Entries           0          0          0
RIPv1 Updates Ignored             0          0          0
RIPv2 Updates Received         2675         11          2
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
none                              0          0          0
```

Meaning The output shows the number of RIP routes learned. It also shows the number of RIP updates sent and received on the RIP-enabled interfaces. Verify the following information:

- The number of RIP routes learned matches the number of expected routes learned. Subnets learned by direct connectivity through an outgoing interface are not listed as RIP routes.
- RIP updates are being sent on each RIP-enabled interface. If no updates are being sent, the routing policy might not be configured to export routes.
- RIP updates are being received on each RIP-enabled interface. If no updates are being received, the routing policy might not be configured to export routes on the host connected to that subnet. The lack of updates might also indicate an authentication error.

Verifying Reachability of All Hosts in the RIP Network

Purpose Use the **traceroute** command on each loopback address in the network to verify that all hosts in the RIP network are reachable from each Juniper Networks device.

Action From operational mode, enter the **traceroute** command.

```
user@R1> traceroute 192.168.3.3
traceroute to 192.168.3.3 (192.168.3.3), 30 hops max, 40 byte packets
 1  10.0.0.2 (10.0.0.2)  1.094 ms  1.028 ms  0.957 ms
 2  192.168.3.3 (192.168.3.3)  1.344 ms  2.245 ms  2.125 ms
```

Meaning Each numbered row in the output indicates a routing hop in the path to the host. The three-time increments indicate the round-trip time (RTT) between the device 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.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Understanding Basic RIP Routing on page 18
- RIP Configuration Overview on page 17

Configuring a RIP Network for EX Series Switches (J-Web Procedure)

NOTE: This topic applies only to the J-Web Application package.

You can use the J-Web interface to create RIP networks.

To configure a RIP network:

1. Select **Configure** > **Routing** > **RIP**.

   NOTE: After you make changes to the configuration on this page, you must commit the changes for them to take effect. To commit all changes to the active configuration, select **Commit Options** > **Commit**. See Using the Commit Options to Commit Configuration Changes for details about all commit options.

2. Click one of the following options:

   - **Add**—Configures a RIP instance. Enter information into the RIP Configuration page as described in Table 3 on page 25.
   - **Edit**—Modifies an existing RIP instance. Enter information into the configuration page for RIP as described in Table 3 on page 25.
   - **Delete**—Deletes an existing RIP instance.

4. To modify RIP global settings, click **Edit**. Enter information in the configuration as described in Table 4 on page 26.

Table 3: RIP Routing Configuration Summary

Field	Function	Your Action
General tab		
Routing instance name	Specifies a name for the routing instance.	Type or select and edit the name.
Preference	Specifies the preference of external routes learned by RIP as compared to those learned from other routing protocols.	Type or select and edit the value.
Metric Out	Specifies the metric value to add to routes transmitted to the neighbor.	Type or select and edit the value.

Copyright © 2017, Juniper Networks, Inc.

Table 3: RIP Routing Configuration Summary *(continued)*

Field	Function	Your Action
Update interval	Specifies an update time interval to periodically send out routes learned by RIP to neighbors.	Type or select and edit the value.
Route timeout	Specifies the route timeout interval for RIP.	Type or select and edit the value.
Policies tab		
Import Policy	Applies one or more policies to routes being imported into the local routing device from the neighbors.	Click **Add** to add an import policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an import policy.
Export Policy	Applies a policy to routes being exported to the neighbors.	Click **Add** to add an export policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an export policy.
Neighbors tab		
RIP-Enabled Interfaces	Selects the interfaces to be associated with the RIP instance.	To enable RIP on an interface, click the check box next to the interface name. Click Edit if you want to modify an interface's settings.

Table 4: Edit RIP Global Settings

Field	Function	Your Action
General tab		
Send	Specifies RIP send options.	Select a value.
Receive	Configure RIP receive options.	Select a value.
Route timeout (sec)	Specifies the route timeout interval for RIP.	Type a value.
Update interval (sec)	Specifies the update time interval to periodically send out routes learned by RIP to neighbors.	Type or select and edit the value.
Hold timeout (sec)	Specifies the time period the expired route is retained in the routing table before being removed.	Type or select and edit the value.
Metric in	Specifies the metric to add to incoming routes when advertising into RIP routes that were learned from other protocols.	Type or select and edit the value.
RIB Group	Specifies a routing table group to install RIP routes into multiple routing tables.	Select and edit the name of the routing table group.

Copyright © 2017, Juniper Networks, Inc.

Table 4: Edit RIP Global Settings *(continued)*

Field	Function	Your Action
Message size	Specifies the number of route entries to be included in every RIP update message.	Type or select and edit the value.
Check Zero	Specifies whether the reserved fields in a RIP packet are zero. Options are: • **check-zero**—Discard version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications. • **no-check-zero**—Receive RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero. This is in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453.	Select a value.
Graceful switchover	Configures graceful switchover for OSPF.	To disable graceful restart, select **Disable**. Type or select and edit the estimated time for the restart to finish, in seconds.
Authentication Type	Specifies the type of authentication for RIP route queries received on an interface. Options are: • None • MD5 • Simple	Select the authentication type. Enter the authentication key for MD5.
Policies tab		
Import Policy	Applies one or more policies to routes being imported into the local routing device from the neighbors.	Click **Add** to add an import policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an import policy.
Export Policy	Applies a policy to routes being exported to the neighbors.	Click **Add** to add an export policy. Click **Move up** or **Move down** to move the selected policy up or down the list of policies. Click **Remove** to remove an export policy.
Trace Options tab		
File Name	Specifies the name of the file to receive the output of the tracing operation.	Type or select and edit the name.
Number of Files	Specifies the maximum number of trace files.	Type or select and edit the name.

Copyright © 2017, Juniper Networks, Inc.

Table 4: Edit RIP Global Settings *(continued)*

Field	Function	Your Action
File Size	Specifies the maximum size for each trace file.	Type or select and edit the name.
World Readable	Specifies whether the trace file can be read by any user or not.	Select **True** to allow any user to read the file. Select **False** to disallow all users being able to read the file.
Flags	Specifies the tracing operation to perform.	Select a value from the list.

Related Documentation

- Monitoring RIP Routing Information on page 173
- *Layer 3 Protocols Supported on EX Series Switches*

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 4

Configuring Authentication for RIP Routes

- Understanding RIP Authentication on page 29
- Enabling Authentication with MD5 Authentication (CLI Procedure) on page 30
- Enabling Authentication with Plain-Text Passwords (CLI Procedure) on page 30
- Example: Configuring Route Authentication for RIP on page 31

Understanding RIP Authentication

RIPv2 provides authentication support so that RIP links can require authentication keys (passwords) before they become active. Authentication provides an additional layer of security on the network beyond the other security features. By default, this authentication is disabled.

Authentication keys can be specified in either plain-text or MD5 form. Authentication requires all routers within the RIP network or subnetwork to have the same authentication type and key (password) configured.

This type of authentication is not supported on RIPv1 networks.

Note that the RIPv2 authentication described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

Release History Table

Release	Description
15.1X49	Note that the RIPv2 authentication described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

Related Documentation

- RIP Overview on page 3
- Enabling Authentication with Plain-Text Passwords (CLI Procedure) on page 30
- Enabling Authentication with MD5 Authentication (CLI Procedure) on page 30

Copyright © 2017, Juniper Networks, Inc.

Enabling Authentication with MD5 Authentication (CLI Procedure)

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 devices in the network:

1. Navigate to the top of the configuration hierarchy.
2. Perform the configuration tasks described in Table 5 on page 30.
3. If you are finished configuring the router, commit the configuration.

Table 5: Configuring MD5 RIP Authentication

Task	CLI Configuration Editor
Navigate to **Rip** level in the configuration hierarchy.	From the **[edit]** hierarchy level, enter **edit protocols rip**
Set the authentication type to **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.	Set the MD5 authentication key: **set authentication-key *password***

Related Documentation

- Understanding RIP Authentication on page 29
- RIP Configuration Overview on page 17
- Enabling Authentication with Plain-Text Passwords (CLI Procedure) on page 30

Enabling Authentication with Plain-Text Passwords (CLI Procedure)

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 devices in the network:

1. Navigate to the top of the configuration hierarchy.
2. Perform the configuration tasks described in Table 6 on page 31.
3. If you are finished configuring the router, commit the configuration.

Copyright © 2017, Juniper Networks, Inc.

Table 6: Configuring Simple RIP Authentication

Task	CLI Configuration Editor
Navigate to **Rip** level in the configuration hierarchy.	From the **[edit]** hierarchy level, enter **edit protocols rip**
Set the authentication type to **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.	Set the authentication key to a simple-text password: **set authentication-key *password***

Related Documentation

- Understanding RIP Authentication on page 29
- RIP Configuration Overview on page 17
- Enabling Authentication with MD5 Authentication (CLI Procedure) on page 30

Example: Configuring Route Authentication for RIP

This example shows how to configure authentication for a RIP network.

- Requirements on page 31
- Overview on page 31
- Configuration on page 32
- Verification on page 35

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

You can configure the router to authenticate RIP route queries. By default, authentication is disabled. You can use one of the following authentication methods:

- Simple authentication—Uses a text password that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet.
- MD5 authentication—Creates an encoded checksum that is included in the transmitted packet. The receiving router uses an authentication key (password) to verify the packet's MD5 checksum.

This example shows MD5 authentication.

Figure 6 on page 32 shows the topology used in this example.

Copyright © 2017, Juniper Networks, Inc.

Figure 6: RIP Authentication Network Topology

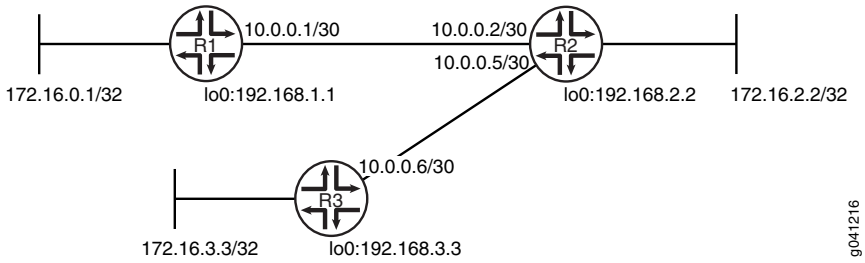

"CLI Quick Configuration" on page 32 shows the configuration for all of the devices in Figure 6 on page 32. The section "Step-by-Step Procedure" on page 33 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip authentication-type md5
set protocols rip authentication-key "$ABC123$ABC123"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip authentication-type md5
set protocols rip authentication-key "$ABC123$ABC123"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
```


Copyright © 2017, Juniper Networks, Inc.

```
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip authentication-type md5
set protocols rip authentication-key "$ABC123$ABC123"
set protocols rip traceoptions file rip-authentication-messages
set protocols rip traceoptions flag auth
set protocols rip traceoptions flag packets
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure RIP authentication:

1. Configure the network interfaces.

   This example shows multiple loopback interface addresses to simulate attached networks.

   ```
   [edit interfaces]
   user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

   user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
   user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
   ```

2. Create the RIP group and add the interface.

   To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

   ```
   [edit protocols rip group rip-group]
   user@R1# set neighbor fe-1/2/0.1
   ```

3. Create the routing policy to advertise both direct and RIP-learned routes.

   ```
   [edit policy-options policy-statement advertise-routes-through-rip term 1]
   user@R1# set from protocol direct
   user@R1# set from protocol rip
   user@R1# set then accept
   ```

4. Apply the routing policy.

   In Junos OS, you can only apply RIP export policies at the group level.

Copyright © 2017, Juniper Networks, Inc.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

5. Require MD5 authentication for RIP route queries received on an interface.

   The passwords must match on neighboring RIP routers. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.

   Do not enter the password as shown here. The password shown here is the encrypted password that is displayed in the configuration after the actual password is already configured.

   ```
   [edit protocols rip]
   user@R1# set authentication-type md5
   user@R1# set authentication-key "$ABC123$ABC123"
   ```

6. Configure tracing operations to track authentication.

   ```
   [edit protocols rip traceoptions]
   user@R1# set file rip-authentication-messages
   user@R1# set flag auth
   user@R1# set flag packets
   ```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
    unit 1 {
        family inet {
            address 10.0.0.1/30;
        }
    }
}
lo0 {
    unit 1 {
        family inet {
            address 172.16.0.1/32;
            address 192.168.1.1/32;
        }
    }
}
```

```
user@R1# show protocols
rip {
    traceoptions {
        file rip-authentication-messages;
        flag auth;
        flag packets;
    }
    authentication-type md5;
```


Copyright © 2017, Juniper Networks, Inc.

```
    authentication-key $ABC123$ABC123; ## SECRET-DATA
    group rip-group {
      export advertise-routes-through-rip;
      neighbor fe-1/2/0.1;
    }
  }
```

```
user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Checking for Authentication Failures on page 35
- Verifying That MD5 Authentication Is Enabled in RIP Update Packets on page 36

Checking for Authentication Failures

Purpose Verify that there are no authentication failures.

Action From operational mode, enter the **show rip statistics** command.

```
user@R1> show rip statistics
RIPv2 info: port 520; holddown 120s.
    rts learned  rts held down  rqsts dropped  resps dropped
              5              0              0              0

fe-1/2/0.1:  5 routes learned; 2 routes advertised; timeout 180s; update interval
 30s
Counter                         Total   Last 5 min  Last minute
-------                     ----------  ----------  ----------
Updates Sent                      2669          10           2
Triggered Updates Sent               2           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            2675          11           2
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
none                                 0           0           0
```

Meaning The output shows that there are no authentication failures.

Verifying That MD5 Authentication Is Enabled in RIP Update Packets

Purpose Use tracing operations to verify that MD5 authentication is enabled in RIP updates.

Action From operational mode, enter the **show log** command.

```
user@R1> show log rip-authentication-messages | match md5
Feb 15 15:45:13.969462          sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:45:43.229867          sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:46:13.174410          sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:46:42.716566          sending msg 0xb9a8c04, 3 rtes (needs MD5)
Feb 15 15:47:11.425076          sending msg 0xb9a8c04, 3 rtes (needs MD5)
...
```

Meaning The **(needs MD5)** output shows that all route updates require MD5 authentication.

Related Documentation

- Understanding Basic RIP Routing on page 18

 Copyright © 2017, Juniper Networks, Inc.

CHAPTER 5

Configuring RIP Timers and Demand Circuits

- Understanding RIP Timers on page 37
- Example: Configuring RIP Timers on page 38
- RIP Demand Circuits Overview on page 44
- Example: Configuring RIP Demand Circuits on page 47

Understanding RIP Timers

RIP uses several timers to regulate its operation.

The update interval is the interval at which routes that are learned by RIP are advertised to neighbors. This timer controls the interval between routing updates. The update interval is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can occur if all routing devices update their neighbors simultaneously.

To configure the update time interval, include the **update-interval** statement:

```
update-interval seconds;
```

seconds can be a value from 10 through 60.

You can set a route timeout interval. If a route is not refreshed after being installed in the routing table by the specified time interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.

To configure the route timeout for RIP, include the **route-timeout** statement:

```
route-timeout seconds;
```

seconds can be a value from 30 through 360. The default value is 180 seconds.

RIP routes expire when either a route timeout limit is met or a route metric reaches infinity, and the route is no longer valid. However, the expired route is retained in the routing table for a specified period so that neighbors can be notified that the route has been dropped. This time period is set by configuring the hold-down timer. Upon expiration of the hold-down timer, the route is removed from the routing table.

Copyright © 2017, Juniper Networks, Inc.

To configure the hold-down timer for RIP, include the **holddown** statement:

holddown *seconds*;

seconds can be a value from 10 through 180. The default value is 120 seconds.

NOTE: **In Junos OS Release 11.1 and later, a retransmission timer is available for RIP demand circuits.**

Generally, we recommend against changing the RIP timers, unless the effects of a change are well understood. The route timeout should be at least three times the update interval. Normally, the default values are best left in effect for standard operations.

Release History Table

Release	Description
11.1	In Junos OS Release 11.1 and later, a retransmission timer is available for RIP demand circuits.

Related Documentation

- Example: Configuring RIP Timers on page 38
- *Example: Configuring RIP Demand Circuits*

Example: Configuring RIP Timers

This example shows how to configure the RIP update interval and how to monitor the impact of the change.

- Requirements on page 38
- Overview on page 38
- Configuration on page 39
- Verification on page 42

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, Device R2 has an update interval of 60 seconds for its neighbor, Device R1, and an update interval of 10 seconds for its neighbor, Device R3.

This example is not necessarily practical, but it is shown for demonstration purposes. Generally, we recommend against changing the RIP timers, unless the effects of a change are well understood. Normally, the default values are best left in effect for standard operations.

Copyright © 2017, Juniper Networks, Inc.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

Figure 7 on page 39 shows the topology used in this example.

Figure 7: RIP Timers Network Topology

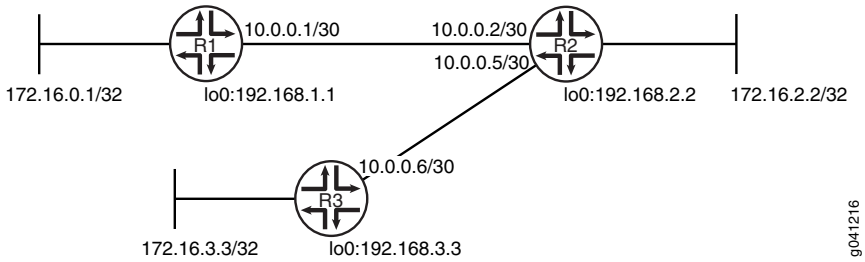

"CLI Quick Configuration" on page 39 shows the configuration for all of the devices in Figure 7 on page 39. The section "Step-by-Step Procedure" on page 40 describes the steps on Device R2.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2 update-interval 60
set protocols rip group rip-group neighbor fe-1/2/1.5 update-interval 10
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
```

Copyright © 2017, Juniper Networks, Inc.

```
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the RIP update interval:

1. Configure the network interfaces.

   This example shows multiple loopback interface addresses to simulate attached networks.

   ```
   [edit interfaces]
   user@R2# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30

   user@R2# set fe-1/2/1 unit 5 family inet address 10.0.0.5/30

   user@R2# set lo0 unit 2 family inet address 192.168.2.2/32
   user@R2# set lo0 unit 2 family inet address 172.16.2.2/32
   ```

2. Configure different update intervals for the two RIP neighbors.

   To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

   ```
   [edit protocols rip group rip-group]
   user@R2# set neighbor fe-1/2/0.2 update-interval 60
   user@R2# set neighbor fe-1/2/1.5 update-interval 10
   ```

3. Create the routing policy to advertise both direct and RIP-learned routes.

   ```
   [edit policy-options policy-statement advertise-routes-through-rip term 1]
   user@R2# set from protocol direct
   user@R2# set from protocol rip
   user@R2# set then accept
   ```

4. Apply the routing policy.

   In Junos OS, you can only apply RIP export policies at the group level.

   ```
   [edit protocols rip group rip-group]
   user@R2# set export advertise-routes-through-rip
   ```

Results

From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the

Copyright © 2017, Juniper Networks, Inc.

intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R2# show interfaces
fe-1/2/0 {
    unit 2 {
        family inet {
            address 10.0.0.2/30;
        }
    }
}
fe-1/2/1 {
    unit 5 {
        family inet {
            address 10.0.0.5/30;
        }
    }
}
lo0 {
    unit 2 {
        family inet {
            address 192.168.2.2/32;
            address 172.16.2.2/32;
        }
    }
}
```

```
user@R2# show protocols
rip {
    group rip-group {
        export advertise-routes-through-rip;
        neighbor fe-1/2/0.2 {
            update-interval 60;
        }
        neighbor fe-1/2/1.5 {
            update-interval 10;
        }
    }
}
```

```
user@R2# show policy-options
policy-statement advertise-routes-through-rip {
    term 1 {
        from protocol [ direct rip ];
        then accept;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2017, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Checking the RIP Updates Sent by Device R2 on page 42
- Checking the RIP Updates Received by Device R2 on page 43
- Checking the RIP Updates Received by Device R3 on page 43

Checking the RIP Updates Sent by Device R2

Purpose Make sure that the RIP update packets are sent at the expected interval.

Action From operational mode, enter the **show rip statistics** command.

```
user@R2> show rip statistics
RIPv2 info: port 520; holddown 120s.
    rts learned  rts held down  rqsts dropped  resps dropped
              4              2               0              0

fe-1/2/0.2:  2 routes learned; 5 routes advertised; timeout 180s; update interval
60s
Counter                       Total   Last 5 min  Last minute
-------                   -----------  ----------  -----------
Updates Sent                      123           5            1
Triggered Updates Sent              0           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            244          10            2
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
none                                0           0            0

fe-1/2/1.5:  2 routes learned; 5 routes advertised; timeout 180s; update interval
10s
Counter                       Total   Last 5 min  Last minute
-------                   -----------  ----------  -----------
Updates Sent                      734          32            6
Triggered Updates Sent              0           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            245          11            2
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
none                                0           0            0
```


Copyright © 2017, Juniper Networks, Inc.

Meaning The **update interval** field shows that the interval is 60 seconds for Neighbor R1 and 10 seconds for Neighbor R3. The **Updates Sent** field shows that Device R2 is sending updates to Device R1 at roughly 1/6 of the rate that it is sending updates to Device R3.

Checking the RIP Updates Received by Device R2

Purpose Make sure that the RIP update packets are sent at the expected interval.

Action From operational mode, enter the **show rip statistics** command.

```
user@R1> show rip statistics
RIPv2 info: port 520; holddown 120s.
    rts learned  rts held down  rqsts dropped  resps dropped
              5              0              0              0

fe-1/2/0.1:  5 routes learned; 2 routes advertised; timeout 180s; update interval
 30s
Counter                       Total   Last 5 min  Last minute
-------                   -----------  -----------  -----------
Updates Sent                      312           10            2
Triggered Updates Sent              2            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            181            5            1
RIPv2 Bad Route Entries             0            0            0
RIPv2 Updates Ignored               0            0            0
Authentication Failures             0            0            0
RIP Requests Received               1            0            0
RIP Requests Ignored                0            0            0
none                                0            0            0
```

Meaning The **RIPv2 Updates Received** field shows the number of updates received from Device R2.

Checking the RIP Updates Received by Device R3

Purpose Make sure that the RIP update packets are sent at the expected interval.

Action From operational mode, enter the **show rip statistics** command.

```
user@R3> show rip statistics
RIPv2 info: port 520; holddown 120s.
    rts learned  rts held down  rqsts dropped  resps dropped
              5              0              0              0

fe-1/2/0.6:  5 routes learned; 2 routes advertised; timeout 180s; update interval
 30s
```

```
Counter                       Total   Last 5 min  Last minute
-------                   -----------  ----------  -----------
Updates Sent                      314          11            2
Triggered Updates Sent              1           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            827          31            6
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
none                                0           0            0
```

Meaning The **RIPv2 Updates Received** field shows the number of updates received from Device R2.

RIP Demand Circuits Overview

RIP periodically sends routing information (RIP packets) to neighboring devices. These periodic broadcasts can consume bandwidth resources and interfere with network traffic by preventing WAN circuits from being closed. Demand circuits for RIP is defined in RFC 2091 and overcomes these issues by exchanging incremental updates on demand.

A demand circuit is a point-to-point connection between two neighboring interfaces configured for RIP. Demand circuits preserve bandwidth by establishing a link when data needs to be transferred, and terminating the link when the data transfer is complete. Demand circuits increase the efficiency of RIP on the configured interfaces by offering minimal network overhead in terms of messages passed between the demand circuit end points, conserving resources, and reducing costs.

By configuring RIP demand circuits, a specific event triggers the device to send an update, thereby eliminating the periodic transmission of RIP packets over the neighboring interface. To save overhead, the device sends RIP information only when changes occur in the routing database, such as:

- The device is first powered on
- The device receives a request for route update information
- A change occurs in the network
- The demand circuit goes down or comes up

The device sends update requests, update responses, and acknowledgments. In addition, the device retransmits updates and requests until valid acknowledgments are received. The device dynamically learns RIP neighbors. If the neighboring interface goes down, RIP flushes routes learned from the neighbor's IP address.

Copyright © 2017, Juniper Networks, Inc.

Routes learned from demand circuits do not age like other RIP entries because demand circuits are in a permanent state. Routes in a permanent state are only removed under the following conditions:

- A formerly reachable route changes to unreachable in an incoming response
- The demand circuit is down due to an excessive number of unacknowledged retransmissions

You can also set the RIP hold-down timer and the RIP demand circuit retransmission timer to regulate performance. The demand circuit uses these timers to determine if there is a change that requires update messages to be sent. There is also a database timer that runs only when RIP flushes learned routes from the routing table.

This topic includes the following sections:

- RIP Demand Circuit Packets on page 45
- Timers Used by RIP Demand Circuits on page 46

RIP Demand Circuit Packets

When you configure an interface for RIP demand circuits, the supported command field packet types are different than those for RIP version 1 and RIP version 2. RIP packets for RIP demand circuits contain three additional packet types and an extended 4-byte update header. Both RIP version 1 and RIP version 2 support the three packet types and the extended 4-byte header. Table 7 on page 45 describes the three packet types.

Table 7: RIP Demand Circuit Packet Types

Packet Type	Description
Update Request	Update request messages seek information for the device's routing table. This message is sent when the device is first powered on or when a down demand circuit comes up. The device sends this message every 5 seconds (by default) until an update response message is received.
Update Response	Update response messages are sent in response to an update request message, which occurs when the device is first powered on or when a down demand circuit comes up. Each update response message contains a sequence number that the neighbor uses to acknowledge the update request.
Update Acknowledge	Update acknowledge messages are sent in response to every update response message received by the neighbor.

Copyright © 2017, Juniper Networks, Inc.

NOTE: These packets are only valid on interfaces configured for RIP demand circuits. If a demand circuit receives a RIP packet that does not contain these packet types, it silently discards the packet and logs an error message similar to the following:

Ignoring RIP packet with invalid version 0 from neighbor 10.0.0.0 and source 10.0.0.1

Related Documentation

- RIP Packets on page 5
- demand-circuit on page 187

Timers Used by RIP Demand Circuits

RIP demand circuits use the RIP hold-down timer and the RIP demand circuit retransmission timer to regulate performance and to determine if there is a change in the network that requires the device to send update messages. The hold-down timer is a global RIP timer that affects the entire RIP configuration; whatever range you configure for RIP applies to RIP demand circuits. The retransmission timer affects only RIP demand circuits. In addition, there is a database timer that runs only when RIP flushes learned routes from the routing table.

- Hold-down timer (global RIP timer)—Use the hold-down timer to configure the number of seconds that RIP waits before updating the routing table. The value of the hold-down timer affects the entire RIP configuration, not just the demand circuit interfaces. The hold-down timer starts when a route timeout limit is met, when a formerly reachable route is unreachable, or when a demand circuit interface is down. When the hold-down timer is running, routes are advertised as unreachable on other interfaces. When the hold-down timer expires, the route is removed from the routing table if all destinations are aware that the route is unreachable or the remaining destinations are down. By default, RIP waits 120 seconds between routing table updates. The range is from 10 to 180 seconds.
- Retransmission timer (RIP demand circuit timer)—RIP demand circuits send update messages every 5 seconds to an unresponsive peer. Use the retransmission timer to limit the number of times a demand circuit resends update messages to an unresponsive peer. If the configured retransmission threshold is reached, routes from the next hop router are marked as unreachable and the hold-down timer starts. The value of the retransmission timer affects only the demand circuit interfaces. To determine the number of times to resend the update message, use the following calculation:

  5 seconds * number of retransmissions = retransmission seconds

  The retransmission range is from 5 through 180 seconds, which corresponds to sending an update message a minimum of 1 time (5 seconds) and a maximum of 36 times (180 seconds).
- Database timer (global timeout timer)—Routes learned from demand circuits do not age like other RIP entries because demand circuits are in a permanent state. On a RIP demand circuit, the database timer starts upon receipt of the update response message

Copyright © 2017, Juniper Networks, Inc.

with the flush flag sent from a RIP demand circuit peer. When the neighbor receives this message, all routes from that peer are flushed, and the database timer starts and runs for the configured route timeout interval. When the database timer is running, routes are still advertised as reachable on other interfaces. When the database timer expires, the device advertises all routes from its peer as unreachable.

Related Documentation

- Example: Configuring RIP Timers on page 38
- Example: Configuring RIP Demand Circuits on page 47
- **holddown on page 193**
- **max-retrans-time on page 196**

Example: Configuring RIP Demand Circuits

This example describes how to configure an interface as a RIP demand circuit.

- Requirements on page 47
- Overview on page 47
- Configuration on page 48
- Verification on page 49

Requirements

Before you begin, configure the device interfaces. See the *Junos OS Network Interfaces Library for Routing Devices* or the *Junos OS Interfaces Configuration Guide for Security Devices*.

Overview

A demand circuit is a point-to-point connection between two neighboring interfaces configured for RIP. Demand circuits increase the efficiency of RIP on the configured interfaces by eliminating the periodic transmission of RIP packets. Demand circuits preserve bandwidth by establishing a link when data needs to be transferred, and terminating the link when the data transfer is complete. In this example, two devices are connected using SONET/SDH interfaces.

NOTE: When you configure RIP demand circuits, any silent removal of the RIP configuration goes unnoticed by the RIP peer and leads to stale entries in the routing table. To clear the stale entries, deactivate and reactivate RIP on the neighboring devices.

In this example, you configure interface **so-0/1/0** with the following settings:

Copyright © 2017, Juniper Networks, Inc.

- **demand-circuit**—Configures the interface as a demand circuit. To complete the demand circuit, you must configure both ends of the pair as demand circuits.
- **max-retrans-time**—RIP demand circuits send update messages every 5 seconds to an unresponsive peer. Use the retransmission timer to limit the number of times a demand circuit resends update messages to an unresponsive peer. If the configured retransmission threshold is reached, routes from the next-hop router are marked as unreachable, and the hold-down timer starts. The value of the retransmission timer affects only the demand circuit interfaces. To determine the number of times to resend the update message, use the following calculation:

  5 seconds x retransmissions = retransmission seconds

  For example, if you want the demand circuit to send only two update messages to an unresponsive peer, the calculation is: 5 x 2 = 10. When you configure the retransmission timer, you enter 10 seconds.

  The retransmission range is from 5 through 180 seconds, which corresponds to sending an update message a minimum of 1 time (5 seconds) and a maximum of 36 times (180 seconds).

Configuration

In the following example, you configure a neighboring interface to be a RIP demand circuit and save the configuration.

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands in the CLI at the **[edit]** hierarchy level.

```
set interfaces so-0/1/0 unit 0 family inet address 192.0.2.0/24
set protocols rip group group1 neighbor so-0/1/0 demand-circuit
set protocols rip group group1 neighbor so-0/1/0 max-retrans-time 10
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIP demand circuit on one neighboring interface:

1. Configure the interface.

   ```
   [edit interfaces]
   user@host# set so-0/1/0 unit 0 family inet address 192.0.2.0/24
   ```

2. Configure the neighbor as a demand circuit.

   ```
   [edit protocols rip]
   user@host# set group group1 neighbor so-0/1/0 demand-circuit
   ```

3. Configure the demand circuit retransmission timer.

Copyright © 2017, Juniper Networks, Inc.

```
[edit protocols rip]
user@host# set group group1 neighbor so-0/1/0 max-retrans-time 10
```

4. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

NOTE: Repeat this entire configuration on the other neighboring interface.

Results

Confirm your configuration by entering the **show interfaces** and **show protocols** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show interfaces
so-0/1/0 {
  unit 0 {
    family inet {
      address 192.0.2.0/24;
    }
  }
}
```

```
user@host# show protocols
rip {
  group group1 {
    neighbor so-0/1/0 {
      demand-circuit;
      max-retrans-time 10;
    }
  }
}
```

Verification

Confirm that the configuration is working properly.

Verifying a Demand Circuit Configuration

Purpose Verify that the demand circuit configuration is working.

Action To verify that the demand circuit configuration is in effect, use the show rip neighbor operational mode command.

```
user@host> show rip neighbor
```

Copyright © 2017, Juniper Networks, Inc.

```
                          Source          Destination     Send   Receive  In
Neighbor          State   Address         Address         Mode   Mode     Met
--------          -----   -------         -----------     ----   -------  ---
so-0/1/0.0(DC)       Up   10.10.10.2      224.0.0.9       mcast  both       1
```

When you configure demand circuits, the **show rip neighbor** command displays a DC flag next to the neighboring interface configured for demand circuits.

NOTE: If you configure demand circuits at the **[edit protocols rip group *group-name* neighbor *neighbor-name*]** hierarchy level, the output shows only the neighboring interface that you specifically configured as a demand circuit. If you configure demand circuits at the **[edit protocols rip group *group-name*]** hierarchy level, all of the interfaces in the group are configured as demand circuits. Therefore, the output shows all of the interfaces in that group as demand circuits.

Related Documentation

- RIP Demand Circuits Overview on page 44

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 6

Configuring BFD for RIP

- Understanding BFD for RIP on page 51
- Example: Configuring BFD for RIP on page 52
- Understanding BFD Authentication for RIP on page 57
- Example: Configuring BFD Authentication for RIP on page 59

Understanding BFD for RIP

The Bidirectional Forwarding Detection (BFD) Protocol is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is detected when the routing device stops receiving a reply after a specified interval. BFD works with a wide variety of network environments and topologies. BFD failure detection times are shorter than RIP detection times, providing faster reaction times to various kinds of failures in the network. Instead of waiting for the routing protocol neighbor timeout, BFD provides rapid detection of link failures. BFD timers are adaptive and can be adjusted to be more or less aggressive. For example, a timer can adapt to a higher value if the adjacency fails, or a neighbor can negotiate a higher value for a timer than the one configured. Note that the functionality of configuring BFD for RIP described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

BFD enables quick failover between a primary and a secondary routed path. The protocol tests the operational status of the interface multiple times per second. BFD provides for configuration timers and thresholds for failure detection. For example, if the minimum interval is set for 50 milliseconds and the threshold uses the default value of three missed messages, a failure is detected on an interface within 200 milliseconds of the failure.

Intervening devices (for example, an Ethernet LAN switch) hide link-layer failures from routing protocol peers, such as when two routers are connected by way of a LAN switch, where the local interface status remains up even when a physical fault happens on the remote link. Link-layer failure detection times vary, depending on the physical media and the Layer 2 encapsulation. BFD can provide fast failure detection times for all media types, encapsulations, topologies, and routing protocols.

To enable BFD for RIP, both sides of the connection must receive an update message from the peer. By default, RIP does not export any routes. Therefore, you must enable update messages to be sent by configuring an export policy for routes before a BFD session is triggered.

Copyright © 2017, Juniper Networks, Inc.

Release History Table

Release	Description
15.1X49	Note that the functionality of configuring BFD for RIP described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

Related Documentation

- Example: Configuring BFD for RIP on page 52

Example: Configuring BFD for RIP

This example shows how to configure Bidirectional Forwarding Detection (BFD) for a RIP network.

- Requirements on page 52
- Overview on page 52
- Configuration on page 54
- Verification on page 56

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

To enable failure detection, include the **bfd-liveness-detection** statement:

```
bfd-liveness-detection {
    detection-time {
        threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    multiplier number;
    no-adaptation;
    transmit-interval {
        threshold milliseconds;
        minimum-interval milliseconds;
    }
    version (1 | automatic);
}
```

Optionally, you can specify the threshold for the adaptation of the detection time by including the **threshold** statement. When the BFD session detection time adapts to a value equal to or greater than the threshold, a single trap and a system log message are sent.

To specify the minimum transmit and receive interval for failure detection, include the **minimum-interval** statement. This value represents the minimum interval at which the local routing device transmits hello packets as well as the minimum interval at which

Copyright © 2017, Juniper Networks, Inc.

the routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds. This examples sets a minimum interval of 600 milliseconds.

NOTE: BFD is an intensive protocol that consumes system resources. Specifying a minimum interval for BFD of less than 100 ms for Routing Engine-based sessions and 10 ms for distributed BFD sessions can cause undesired BFD flapping.

Depending on your network environment, these additional recommendations might apply:

- For large-scale network deployments with a large number of BFD sessions, specify a minimum interval of 300 ms for Routing Engine-based sessions and 100 ms for distributed BFD sessions.
- For very large-scale network deployments with a large number of BFD sessions, contact Juniper Networks customer support for more information.
- For BFD sessions to remain up during a Routing Engine switchover event when nonstop active routing (NSR) is configured, specify a minimum interval of 2500 ms for Routing Engine-based sessions. For distributed BFD sessions with nonstop active routing configured, the minimum interval recommendations are unchanged and depend only on your network deployment.

You can optionally specify the minimum transmit and receive intervals separately.

To specify only the minimum receive interval for failure detection, include the **minimum-receive-interval** statement. This value represents the minimum interval at which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,00 milliseconds.

To specify only the minimum transmit interval for failure detection, include the **transmit-interval minimum-interval** statement. This value represents the minimum interval at which the local routing device transmits hello packets to the neighbor with which it has established a BFD session. You can configure a value in the range from 1 through 255,000 milliseconds.

To specify the number of hello packets not received by a neighbor that causes the originating interface to be declared down, include the **multiplier** statement. The default is 3, and you can configure a value in the range from 1 through 255.

To specify the threshold for detecting the adaptation of the transmit interval, include the **transmit-interval threshold** statement. The threshold value must be greater than the transmit interval.

To specify the BFD version used for detection, include the **version** statement. The default is to have the version detected automatically.

Copyright © 2017, Juniper Networks, Inc.

You can trace BFD operations by including the **traceoptions** statement at the **[edit protocols bfd]** hierarchy level.

In Junos OS Release 9.0 and later, you can configure BFD sessions not to adapt to changing network conditions. To disable BFD adaptation, include the **no-adaptation** statement. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

Figure 8 on page 54 shows the topology used in this example.

Figure 8: RIP BFD Network Topology

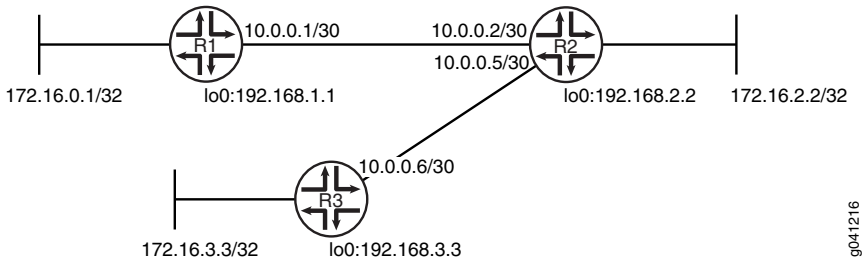

"CLI Quick Configuration" on page 54 shows the configuration for all of the devices in Figure 8 on page 54. The section "Step-by-Step Procedure" on page 55 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
```


Copyright © 2017, Juniper Networks, Inc.

```
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a BFD for a RIP network:

1. Configure the network interfaces.

   ```
   [edit interfaces]
   user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
   ```

2. Create the RIP group and add the interface.

   To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

   ```
   [edit protocols rip group rip-group]
   user@R1# set neighbor fe-1/2/0.1
   ```

3. Create the routing policy to advertise both direct and RIP-learned routes.

   ```
   [edit policy-options policy-statement advertise-routes-through-rip term 1]
   user@R1# set from protocol direct
   user@R1# set from protocol rip
   user@R1# set then accept
   ```

4. Apply the routing policy.

   In Junos OS, you can only apply RIP export policies at the group level.

   ```
   [edit protocols rip group rip-group]
   user@R1# set export advertise-routes-through-rip
   ```

5. Enable BFD.

   ```
   [edit protocols rip group rip-group]
   user@R1# set bfd-liveness-detection minimum-interval 600
   ```

6. Configure tracing operations to track BFD messages.

Copyright © 2017, Juniper Networks, Inc.

```
[edit protocols bfd traceoptions]
user@R1# set file bfd-trace
user@R1# set flag all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
    unit 1 {
        family inet {
            address 10.0.0.1/30;
        }
    }
}
```

```
user@R1# show protocols
bfd {
    traceoptions {
        file bfd-trace;
        flag all;
    }
}
rip {
    group rip-group {
        export advertise-routes-through-rip;
        bfd-liveness-detection {
            minimum-interval 600;
        }
        neighbor fe-1/2/0.1;
    }
}
```

```
user@R1# show policy-options
policy-statement advertise-routes-through-rip {
    term 1 {
        from protocol [ direct rip ];
        then accept;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying That the BFD Sessions Are Up on page 57
- Checking the BFD Trace File on page 57

 Copyright © 2017, Juniper Networks, Inc.

Verifying That the BFD Sessions Are Up

Purpose Make sure that the BFD sessions are operating.

Action From operational mode, enter the **show bfd session** command.

```
user@R1> show bfd session
                                                  Detect   Transmit
Address                  State     Interface      Time     Interval  Multiplier
10.0.0.2                 Up        fe-1/2/0.1     1.800     0.600        3

1 sessions, 1 clients
Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps
```

Meaning The output shows that there are no authentication failures.

Checking the BFD Trace File

Purpose Use tracing operations to verify that BFD packets are being exchanged.

Action From operational mode, enter the **show log** command.

```
user@R1> show log bfd-trace
Feb 16 10:26:32 PPM Trace: BFD periodic xmit to 10.0.0.2 (IFL 124, rtbl 53,
single-hop port)
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 86:
Feb 16 10:26:32    IfIndex (3) len 4: 0
Feb 16 10:26:32    Protocol (1) len 1: BFD
Feb 16 10:26:32    Data (9) len 61: (hex) 42 46 44 20 70 61 63 6b 65 74 20 66 72
 6f 6d 20 31 30 2e
Feb 16 10:26:32 PPM Trace: BFD packet from 10.0.0.1 (IFL 73, rtbl 56, ttl 255)
absorbed
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 60:
Feb 16 10:26:32    IfIndex (3) len 4: 0
Feb 16 10:26:32    Protocol (1) len 1: BFD
Feb 16 10:26:32    Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
 78 6d 69 74 20 6f
...
```

Meaning The output shows the normal functioning of BFD.

Related Documentation

- Understanding BFD for RIP on page 51

Understanding BFD Authentication for RIP

BFD enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when running BFD over

Copyright © 2017, Juniper Networks, Inc.

Network Layer protocols, the risk of service attacks can be significant. We strongly recommend using authentication if you are running BFD over multiple hops or through insecure tunnels. Beginning with Junos OS Release 9.6, Junos OS supports authentication for BFD sessions running over RIP. BFD authentication is only supported in the domestic image and is not available in the export image.

You authenticate BFD sessions by specifying an authentication algorithm and keychain, and then associating that configuration information with a security authentication keychain using the keychain name.

The following sections describe the supported authentication algorithms, security keychains, and the level of authentication that can be configured:

- BFD Authentication Algorithms on page 58
- Security Authentication Keychains on page 59
- Strict Versus Loose Authentication on page 59

BFD Authentication Algorithms

Junos OS supports the following algorithms for BFD authentication:

- **simple-password**—Plain-text password. One to 16 bytes of plain text are used to authenticate the BFD session. One or more passwords can be configured. This method is the least secure and should be used only when BFD sessions are not subject to packet interception.
- **keyed-md5**—Keyed Message Digest 5 hash algorithm for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed MD5 uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than or equal to the last sequence number received. Although more secure than a simple password, this method is vulnerable to replay attacks. Increasing the rate at which the sequence number is updated can reduce this risk.
- **meticulous-keyed-md5**—Meticulous keyed Message Digest 5 hash algorithm. This method works in the same manner as keyed MD5, but the sequence number is updated with every packet. Although more secure than keyed MD5 and simple passwords, this method might take additional time to authenticate the session.
- **keyed-sha-1**—Keyed Secure Hash Algorithm I for sessions with transmit and receive intervals greater than 100 ms. To authenticate the BFD session, keyed SHA uses one or more secret keys (generated by the algorithm) and a sequence number that is updated periodically. The key is not carried within the packets. With this method, packets are accepted at the receiving end of the session if one of the keys matches and the sequence number is greater than the last sequence number received.
- **meticulous-keyed-sha-1**—Meticulous keyed Secure Hash Algorithm I. This method works in the same manner as keyed SHA, but the sequence number is updated with every packet. Although more secure than keyed SHA and simple passwords, this method might take additional time to authenticate the session.

Copyright © 2017, Juniper Networks, Inc.

NOTE: Nonstop active routing is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

Security Authentication Keychains

The security authentication keychain defines the authentication attributes used for authentication key updates. When the security authentication keychain is configured and associated with a protocol through the keychain name, authentication key updates can occur without interrupting routing and signaling protocols.

The authentication keychain contains one or more keychains. Each keychain contains one or more keys. Each key holds the secret data and the time at which the key becomes valid. The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

BFD allows multiple clients per session, and each client can have its own keychain and algorithm defined. To avoid confusion, we recommend specifying only one security authentication keychain.

Strict Versus Loose Authentication

By default, strict authentication is enabled and authentication is checked at both ends of each BFD session. Optionally, to smooth migration from nonauthenticated sessions to authenticated sessions, you can configure *loose checking*. When loose checking is configured, packets are accepted without authentication being checked at each end of the session. This feature is intended for transitional periods only.

Related Documentation

- Example: Configuring BFD Authentication for RIP on page 59
- bfd-liveness-detection on page 183
- **authentication-key-chains** statement in the *Junos OS Administration Library*
- **show bfd session** command in the CLI Explorer
- Example: Configuring BFD for RIP on page 52

Example: Configuring BFD Authentication for RIP

This example shows how to configure Bidirectional Forwarding Detection (BFD) authentication for a RIP network.

- Requirements on page 60
- Overview on page 60
- Configuration on page 60
- Verification on page 64

Copyright © 2017, Juniper Networks, Inc.

Requirements

No special configuration beyond device initialization is required before configuring this example.

The devices must be running Junos OS Release 9.6 or later.

Overview

Only three steps are needed to configure authentication on a BFD session:

1. Specify the BFD authentication algorithm for the RIP protocol.
2. Associate the authentication keychain with the RIP protocol.
3. Configure the related security authentication keychain.

Figure 9 on page 60 shows the topology used in this example.

Figure 9: RIP BFD Authentication Network Topology

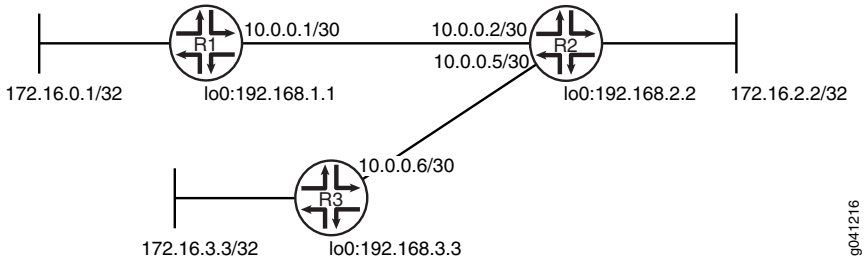

"CLI Quick Configuration" on page 60 shows the configuration for all of the devices in Figure 9 on page 60. The section "Step-by-Step Procedure" on page 61 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set protocols bfd traceoptions file bfd-trace
set protocols bfd traceoptions flag all
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
    keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
```


Copyright © 2017, Juniper Networks, Inc.

```
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret $ABC123$ABC123
set security authentication-key-chains key-chain bfd-rip key 53 start-time
  "2012-2-16.12:00:00 -0800"
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
  keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret $ABC123$ABC123
set security authentication-key-chains key-chain bfd-rip key 53 start-time
  "2012-2-16.12:00:00 -0800"
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group bfd-liveness-detection minimum-interval 600
set protocols rip group rip-group bfd-liveness-detection authentication key-chain bfd-rip
set protocols rip group rip-group bfd-liveness-detection authentication algorithm
  keyed-md5
set protocols rip group rip-group bfd-liveness-detection authentication loose-check
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set security authentication-key-chains key-chain bfd-rip key 53 secret $ABC123$ABC123
set security authentication-key-chains key-chain bfd-rip key 53 start-time
  "2012-2-16.12:00:00 -0800"
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a BFD authentication:

1. Configure the network interfaces.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

2. Create the RIP group and add the interface.

   To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

   ```
   [edit protocols rip group rip-group]
   user@R1# set neighbor fe-1/2/0.1
   ```

3. Create the routing policy to advertise both direct and RIP-learned routes.

   ```
   [edit policy-options policy-statement advertise-routes-through-rip term 1]
   user@R1# set from protocol direct
   user@R1# set from protocol rip
   user@R1# set then accept
   ```

4. Apply the routing policy.

   In Junos OS, you can only apply RIP export policies at the group level.

   ```
   [edit protocols rip group rip-group]
   user@R1# set export advertise-routes-through-rip
   ```

5. Enable BFD.

   ```
   [edit protocols rip group rip-group]
   user@R1# set bfd-liveness-detection minimum-interval 600
   ```

6. Specify the algorithm (**keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, **meticulous-keyed-sha-1**, or **simple-password**) to use.

   NOTE: Nonstop active routing is not supported with meticulous-keyed-md5 and meticulous-keyed-sha-1 authentication algorithms. BFD sessions using these algorithms might go down after a switchover.

   ```
   [edit protocols rip group rip-group]
   user@R1# set bfd-liveness-detection authentication algorithm keyed-md5
   ```

7. Specify the keychain to be used to associate BFD sessions on RIP with the unique security authentication keychain attributes.

   The keychain you specify must match a keychain name configured at the **[edit security authentication key-chains]** hierarchy level.

   The algorithm and keychain must be configured on both ends of the BFD session, and they must match. Any mismatch in configuration prevents the BFD session from being created.

   ```
   [edit protocols rip group rip-group]
   user@R1# set bfd-liveness-detection authentication key-chain bfd-rip
   ```


Copyright © 2017, Juniper Networks, Inc.

8. (Optional) Specify loose authentication checking if you are transitioning from nonauthenticated sessions to authenticated sessions.

```
[edit protocols rip group rip-group]
user@R1# set bfd-liveness-detection authentication loose-check
```

9. Specify the unique security authentication information for BFD sessions:

   - The matching keychain name as specified in Step 7.
   - At least one key, a unique integer between **0** and **63**. Creating multiple keys allows multiple clients to use the BFD session.
   - The secret data used to allow access to the session.
   - The time at which the authentication key becomes active, in the format *yyyy-mm-dd.hh:mm:ss*.

```
[edit security authentication-key-chains key-chain bfd-rip]
user@R1# set key 53 secret $ABC123$ABC123
user@R1# set key 53 start-time "2012-2-16.12:00:00 -0800"
```

10. Configure tracing operations to track BFD authentication.

```
[edit protocols bfd traceoptions]
user@R1# set file bfd-trace
user@R1# set flag all
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show security** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
   unit 1 {
      family inet {
         address 10.0.0.1/30;
      }
   }
}
```

```
user@R1# show protocols
bfd {
   traceoptions {
      file bfd-trace;
      flag all;
   }
}
rip {
   group rip-group {
      export advertise-routes-through-rip;
      bfd-liveness-detection {
         minimum-interval 600;
```

Copyright © 2017, Juniper Networks, Inc.

```
    }
    neighbor fe-1/2/0.1;
  }
}
```

```
user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
```

```
user@R1# show security
authentication-key-chains {
  key-chain bfd-rip {
    key 53 {
      secret $ABC123$ABC123
      start-time "2012-2-16.12:00:00 -0800";
    }
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying That the BFD Sessions Are Authenticated on page 64
- Viewing Extensive Information About the BFD Authentication on page 65
- Checking the BFD Trace File on page 65

Verifying That the BFD Sessions Are Authenticated

Purpose Make sure that the BFD sessions are authenticated.

Action From operational mode, enter the **show bfd session detail** command.

```
user@R1> show bfd session detail
                                                 Detect   Transmit
Address                 State     Interface      Time     Interval  Multiplier
10.0.0.2                Up        fe-1/2/0.1     1.800     0.600        3
 Client RIP, TX interval 0.600, RX interval 0.600, Authenticate
 Session up time 01:39:34
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Logical system 6, routing table index 53

1 sessions, 1 clients
Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps
```

Meaning **Authenticate** is displayed to indicate that BFD authentication is configured.

Copyright © 2017, Juniper Networks, Inc.

Viewing Extensive Information About the BFD Authentication

Purpose View the keychain name, the authentication algorithm and mode for each client in the session, and the BFD authentication configuration status.

Action From operational mode, enter the **show bfd session extensive** command.

```
user@R1> show bfd session extensive
                                                  Detect   Transmit
Address                 State     Interface       Time     Interval  Multiplier
10.0.0.2                Up        fe-1/2/0.1      1.800     0.600        3
 Client RIP, TX interval 0.600, RX interval 0.600, Authenticate
        keychain bfd-rip, algo keyed-md5, mode loose
 Session up time 01:46:29
 Local diagnostic None, remote diagnostic None
 Remote state Up, version 1
 Logical system 6, routing table index 53
 Min async interval 0.600, min slow interval 1.000
 Adaptive async TX interval 0.600, RX interval 0.600
 Local min TX interval 0.600, minimum RX interval 0.600, multiplier 3
 Remote min TX interval 0.600, min RX interval 0.600, multiplier 3
 Local discriminator 225, remote discriminator 226
 Echo mode disabled/inactive
 Authentication enabled/active, keychain bfd-rip, algo keyed-md5, mode loose
  Session ID: 0x300501

1 sessions, 1 clients
Cumulative transmit rate 1.7 pps, cumulative receive rate 1.7 pps
```

Meaning The output shows the keychain name, the authentication algorithm and mode for the client in the session, and the BFD authentication configuration status.

Checking the BFD Trace File

Purpose Use tracing operations to verify that BFD packets are being exchanged.

Action From operational mode, enter the **show log** command.

```
user@R1> show log bfd-trace
Feb 16 10:26:32 PPM Trace: BFD periodic xmit to 10.0.0.2 (IFL 124, rtbl 53,
single-hop port)
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 86:
Feb 16 10:26:32    IfIndex (3) len 4: 0
Feb 16 10:26:32    Protocol (1) len 1: BFD
Feb 16 10:26:32    Data (9) len 61: (hex) 42 46 44 20 70 61 63 6b 65 74 20 66 72
 6f 6d 20 31 30 2e
Feb 16 10:26:32 PPM Trace: BFD packet from 10.0.0.1 (IFL 73, rtbl 56, ttl 255)
absorbed
Feb 16 10:26:32 Received Downstream TraceMsg (24) len 60:
Feb 16 10:26:32    IfIndex (3) len 4: 0
Feb 16 10:26:32    Protocol (1) len 1: BFD
Feb 16 10:26:32    Data (9) len 35: (hex) 42 46 44 20 70 65 72 69 6f 64 69 63 20
```

Copyright © 2017, Juniper Networks, Inc.

```
 78 6d 69 74 20 6f
...
```

Meaning The output shows the normal functioning of BFD.

Related Documentation

- Understanding BFD Authentication for RIP on page 57

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 7

Using Metrics to Control Traffic in a RIP Network

- Understanding Traffic Control with Metrics in a RIP Network on page 67
- Example: Controlling Traffic in a RIP Network with an Incoming Metric on page 68
- Example: Controlling Traffic in a RIP Network with an Outgoing Metric on page 70
- Example: Configuring the Metric Value Added to Imported RIP Routes on page 72

Understanding Traffic Control with Metrics in a RIP Network

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.

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.

You might want to increase the metric of routes to decrease the likelihood that a particular route is selected and installed in the routing table. This process is sometimes referred to

Copyright © 2017, Juniper Networks, Inc.

as *route poisoning*. Some reasons that you might want to poison a route are that the route is relatively expensive to use, or it has relatively low bandwidth.

A route with a higher metric than another route becomes the active route only when the lower-metric route becomes unavailable. In this way, the higher-metric route serves as a backup path.

One way to increase the metric of imported routes is to configure an import policy. Another way is to include the **metric-in** statement in the RIP neighbor configuration. One way to increase the metric of export routes is to configure an export policy. Another way is to include the **metric-out** statement in the RIP neighbor configuration.

Related Documentation

- RIP Overview on page 3
- Example: Controlling Traffic in a RIP Network with an Incoming Metric on page 68
- Example: Controlling Traffic in a RIP Network with an Outgoing Metric on page 70

Example: Controlling Traffic in a RIP Network with an Incoming Metric

This example shows how to control traffic with an incoming metric.

- Requirements on page 68
- Overview on page 68
- Configuration on page 69
- Verification on page 69

Requirements

Before you begin, define RIP groups, and add interfaces to the groups. Then configure a routing policy to export directly connected routes and routes learned through the RIP routing exchanges. See "Example: Configuring a Basic RIP Network" on page 18.

Overview

In this example, routes to Router D are received by Router A across both of its RIP-enabled interfaces as shown in Figure 10 on page 69. 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 Router A through Router B to Router D.

Copyright © 2017, Juniper Networks, Inc.

Figure 10: Controlling Traffic in a RIP Network with the Incoming Metric

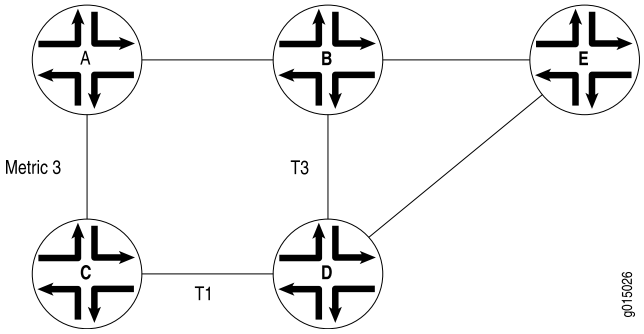

To force this flow, you can modify the route metrics as they are imported into Router A's routing table. By setting the incoming metric on the interface from Router A to Router C, you modify the metric on all routes received through that interface. Setting the incoming route metric on Router A changes 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.

This example uses a RIP group called **alpha 1** on interface **g3–0/0/0**.

Configuration

Step-by-Step Procedure

To control traffic with an incoming metric:

1. Enable RIP on the interface.

   ```
   [edit protocols rip]
   user@host# set group alpha1 neighbor ge-0/0/0
   ```

2. Set the incoming metric.

   ```
   [edit protocols rip]
   user@host# set metric-in 3
   ```

3. If you are done configuring the device, commit the configuration.

   ```
   [edit]
   user@host# commit
   ```

Verification

To verify that the configuration is working properly, enter the **show route protocols rip** command.

Copyright © 2017, Juniper Networks, Inc.

Related Documentation

- Understanding Traffic Control with Metrics in a RIP Network on page 67
- RIP Configuration Overview on page 17
- Example: Controlling Traffic in a RIP Network with an Outgoing Metric on page 70
- Verifying a RIP Configuration on page 174

Example: Controlling Traffic in a RIP Network with an Outgoing Metric

This example shows how to control traffic with an outgoing metric.

- Requirements on page 70
- Overview on page 70
- Configuration on page 71
- Verification on page 71

Requirements

Before you begin:

- Define RIP groups, and add interfaces to the groups. Then configure a routing policy to export directly connected routes and routes learned through RIP routing exchanges. See "Example: Configuring a Basic RIP Network" on page 18.
- Control traffic with an incoming metric. See "Example: Controlling Traffic in a RIP Network with an Incoming Metric" on page 68.

Overview

In this example, each route from Router A to Router D has two hops as shown in Figure 11 on page 71. However, because the link from Router A to Router B in the RIP group 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.

Copyright © 2017, Juniper Networks, Inc.

Figure 11: Controlling Traffic in a RIP Network with the Outgoing Metric

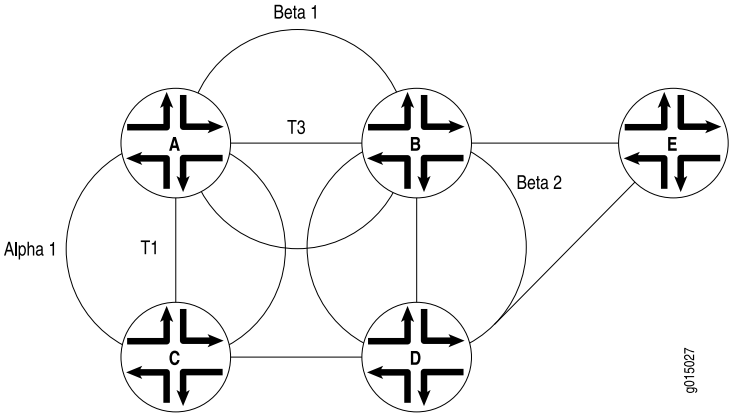

If the outgoing metric for the Alpha 1 RIP group—the A-to-C link—is changed to 3, Router D calculates the total path metric from Router A through Router C as 4. In contrast, the unchanged default total path metric to Router A through Router B in the RIP group is 2. The fact that Router A's interfaces 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 outgoing 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.

This example uses an outgoing metric of **3**.

Configuration

Step-by-Step Procedure

To control traffic with an outgoing metric:

1. Set the outgoing metric.

```
[edit protocols rip group alpha1]
user@host# set metric-out 3
```

2. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Verification

To verify that the configuration is working properly, enter the **show protocols rip** command.

Related Documentation

- Understanding Traffic Control with Metrics in a RIP Network on page 67
- RIP Configuration Overview on page 17
- Verifying a RIP Configuration on page 174

Copyright © 2017, Juniper Networks, Inc.

Example: Configuring the Metric Value Added to Imported RIP Routes

This example shows how to change the default metric to be added to incoming routes to control the route selection process.

- Requirements on page 72
- Overview on page 72
- Configuration on page 72
- Verification on page 75

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

Normally, when multiple routes are available, RIP selects the route with the lowest hop count. Changing the default metric enables you to control the route selection process such that a route with a higher hop count can be preferred over of a route with a lower hop count.

Figure 12 on page 72 shows the topology used in this example.

Figure 12: RIP Incoming Metrics Network Topology

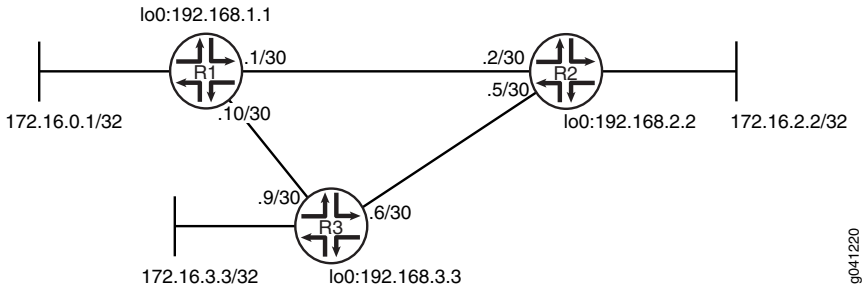

Device R1 has two potential paths to reach 172.16.2.2/32. The default behavior is to send traffic out the 0.1/30 interface facing Device R2. Suppose, though, that the path through Device R3 is less expensive to use or has higher bandwidth links. This example shows how to use the **metric-in** statement to ensure that Device R1 uses the path through Device R3 to reach 172.16.2.2/32. "CLI Quick Configuration" on page 72 shows the configuration for all of the devices in Figure 12 on page 72. The section "Step-by-Step Procedure" on page 73 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Copyright © 2017, Juniper Networks, Inc.

Device R1

```
set interfaces fe-1/2/0 unit 1 description to-R2
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces ge-1/2/1 unit 10 description to-R3
set interfaces ge-1/2/1 unit 10 family inet address 10.0.0.10/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group primary export advertise-routes-through-rip
set protocols rip group primary neighbor ge-1/2/1.10
set protocols rip group secondary export advertise-routes-through-rip
set protocols rip group secondary neighbor fe-1/2/0.1 metric-in 4
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces ge-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor ge-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces ge-1/2/1 unit 9 family inet address 10.0.0.9/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set protocols rip group rip-group neighbor ge-1/2/1.9
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIP metrics:

1. Configure the network interfaces.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 description to-R2
user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30
```

Copyright © 2017, Juniper Networks, Inc.

```
user@R1# set ge-1/2/1 unit 10 description to-R3
user@R1# set ge-1/2/1 unit 10 family inet address 10.0.0.10/30

user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
```

2. Create the RIP groups and add the interfaces.

   To configure RIP in Junos OS, you must configure one or more groups that contain the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

   For the interface that is facing Device R2, the **metric-in 4** setting causes this route to be less likely to be chosen as the active route.

   ```
   [edit protocols rip]
   user@R1# set group primary neighbor ge-1/2/1.10
   user@R1# set group secondary neighbor fe-1/2/0.1 metric-in 4
   ```

3. Create the routing policy to advertise both direct and RIP-learned routes.

   ```
   [edit policy-options policy-statement advertise-routes-through-rip term 1]
   user@R1# set from protocol direct
   user@R1# set from protocol rip
   user@R1# set then accept
   ```

4. Apply the routing policy.

   In Junos OS, you can only apply RIP export policies at the group level.

   ```
   [edit protocols rip]
   user@R1# set group primary export advertise-routes-through-rip
   user@R1# set group secondary export advertise-routes-through-rip
   ```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
    unit 1 {
        description to-R2;
        family inet {
            address 10.0.0.1/30;
        }
    }
}
ge-1/2/1 {
    unit 10 {
        description to-R3;
        family inet {
            address 10.0.0.10/30;
```


Copyright © 2017, Juniper Networks, Inc.

```
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}
```

```
user@R1# show protocols
rip {
  group primary {
    export advertise-routes-through-rip;
    neighbor ge-1/2/1.10;
  }
  group secondary {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1 {
      metric-in 4;
    }
  }
}
```

```
user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying That the Expected Route Is Active on page 75
- Removing the metric-in Statement on page 76

Verifying That the Expected Route Is Active

Purpose Make sure that to reach 172.16.2.2/32, Device R1 uses the path through Device R3.

Action From operational mode, enter the **show route 172.16.2.2** command.

```
user@R1> show route 172.16.2.2
inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.2.2/32      *[RIP/100] 00:15:46, metric 3, tag 0
                    > to 10.0.0.9 via ge-1/2/1.10
```

Copyright © 2017, Juniper Networks, Inc.

Meaning The **to 10.0.0.9 via ge-1/2/1.10** output shows that Device R1 uses the path through Device R3 to reach 172.16.2.2/32. The metric for this route is 3.

Removing the metric-in Statement

Purpose Delete or deactivate the **metric-in** statement to see what happens to the 172.16.2.2/32 route.

Action

1. From configuration mode, deactivate the **metric-in** statement.

   ```
   [edit protocols rip group secondary neighbor fe-1/2/0.1]
   user@R1# deactivate metric-in
   user@R1# commit
   ```

2. From operational mode, enter the **show route 172.16.2.2** command.

   ```
   user@R1> show route 172.16.2.2
   inet.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
   + = Active Route, - = Last Active, * = Both

   172.16.2.2/32      *[RIP/100] 00:00:06, metric 2, tag 0
                       > to 10.0.0.2 via fe-1/2/0.1
   ```

Meaning The **to 10.0.0.2 via fe-1/2/0.1** output shows that Device R1 uses the path through Device R2 to reach 172.16.2.2/32. The metric for this route is 2.

Related Documentation

- Understanding Traffic Control with Metrics in a RIP Network on page 67

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 8

Configuring Point-to-Multipoint RIP Networks

- Configuring Point-to-Multipoint RIP Networks Overview on page 77
- Example: Configuring Point-to-Multipoint RIP Networks on page 78

Configuring Point-to-Multipoint RIP Networks Overview

A point-to-multipoint RIP network consists of a device having two or more peers on a single interface. All the devices forming a point-to-multipoint connection are placed in a single broadcast domain.

In a RIP network, a device can have a single peer or multiple peers for an interface. However, the demand circuit feature implementation in a RIP network requires the use of a single RIP peer. When you configure the following statements, a RIP network with demand circuits can also be configured to have multiple peers on an interface:

- Configuring the interface type to be a multipoint interface by using the interface-type (Protocols RIP) **p2mp** statement.
- Enabling dynamic peer discovery by using the dynamic-peers statement (SRX Series devices only).

NOTE: Before configuring the dynamic-peers statement, IPsec must be configured and IPsec tunnels must be set up by configuring IPsec parameters. Without IPsec configuration, the remote peers have to be explicitly configured at the RIP protocol level by using the peer ***address*** statement. See *Configuring Security Associations for IPsec on an ES PIC* for more details.

- Configuring peers by using the peer ***address*** statement.

```
[edit]
protocols {
  rip {
    group red {
      neighbor fe-0/1/3 {
        interface-type (Protocols RIP) p2mp;
        peer address; (or use dynamic-peers;)
```

Copyright © 2017, Juniper Networks, Inc.

```
        }
      }
    }
  }
```

The show rip statistics peer ***address*** command can be used to display the RIP statistics at the peer level. The clear rip statistics peer ***address*** command can be used to clear the RIP statistics for a peer. Alternatively, you can use the show rip statistics peer **all** and clear rip statistics peer **all** command to display and clear RIP statistics for all peers.

Related Documentation

- Example: Configuring Point-to-Multipoint RIP Networks on page 78

Example: Configuring Point-to-Multipoint RIP Networks

This example shows how to configure a point-to-multipoint RIP network.

- Requirements on page 78
- Overview on page 78
- Configuration on page 79
- Verification on page 82

Requirements

This example uses the following hardware and software components:

- M Series routers, MX Series routers, T Series routers, or SRX Series devices
- Junos OS Release 12.1 or later

Overview

In a RIP network, a device can have a single peer or multiple peers for an interface. However, the demand circuit feature implementation in a RIP network requires the use of a single RIP peer.

When you include the following statements, the demand circuit implementation can have multiple peers for a given RIP neighbor.

- Configuring the interface type to be a multipoint interface by using the interface-type (Protocols RIP) **p2mp** statement.
- Enabling dynamic peer discovery by using the dynamic-peers statement (SRX Series devices only).

NOTE: To configure the dynamic-peers statement, IPsec tunnels must be set up by configuring IPsec parameters. See *Configuring Security Associations for IPsec on an ES PIC* for more details.

- Configuring peers by using the peer ***address*** statement.

Copyright © 2017, Juniper Networks, Inc.

```
[edit]
protocols {
  rip {
    group red {
      neighbor fe-0/1/3 {
        interface-type (Protocols RIP) p2mp;
        peer address; (or use dynamic-peers;)
      }
    }
  }
}
```

The show rip statistics **peer** command can be used to display the RIP statistics at the peer level.

Topology

In this example, Devices R1, R2, and R3 form a point-to-multipoint network. R1 is connected to R2 and to R3 as a point-to-multipoint connection through a switch that places all devices in the same broadcast domain. RIP demand circuits are configured on all three devices. The two peers to R1 are configured statically by using the peer ***address*** statement. The dynamic-peers statement is not used here.

Figure 13 on page 79 shows the topology used in this example.

Figure 13: Configuring a Point-to-Multipoint RIP Network

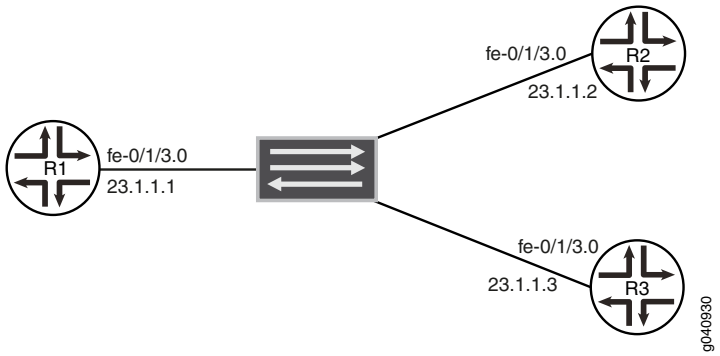

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-0/1/3 unit 0 family inet address 23.1.1.1/24
set policy-options policy-statement accept-rip-routes term from-direct from protocol
  direct
set policy-options policy-statement accept-rip-routes term from-direct then accept
set policy-options policy-statement accept-rip-routes term from-rip from protocol rip
set policy-options policy-statement accept-rip-routes term from-rip then accept
set protocols rip traceoptions file R1.log size 4m world-readable
```

Copyright © 2017, Juniper Networks, Inc.

```
set protocols rip traceoptions flag all detail
set protocols rip group red export accept-rip-routes
set protocols rip group red neighbor fe-0/1/3.0 interface-type p2mp
set protocols rip group red neighbor fe-0/1/3.0 peer 23.1.1.2
set protocols rip group red neighbor fe-0/1/3.0 peer 23.1.1.3
set protocols rip group red neighbor fe-0/1/3.0 demand-circuit
set protocols rip group red neighbor fe-0/1/3.0 max-retrans-time 10
```

Similarly, configure Devices R2 and R3, omitting the **peer *address*** configuration statement.

Configuring a Point-to-Multipoint RIP Network (with Demand Circuits)

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the point-to-multipoint feature across a RIP network:

1. Configure the device interface.

   ```
   [edit interfaces fe-0/1/3 unit 0]
   user@R1# set family inet address 23.1.1.1/24
   ```

2. Define a policy for exporting RIP routes from the routing table to the protocol for transmission through the network.

   ```
   [edit policy-options policy-statement accept-rip-routes]
   user@R1# set term from-direct from protocol direct
   user@R1# set term from-direct then accept
   user@R1# set term from-rip from protocol rip
   user@R1# set term from-rip then accept
   ```

3. Configure RIP and a RIP group with the defined export policy and point-to-multipoint configuration statements.

   ```
   [edit protocols rip]
   user@R1# set traceoptions file R1.log size 4m world-readable
   user@R1# set traceoptions flag all detail
   user@R1# set group red export accept-rip-routes
   user@R1# set group red neighbor fe-0/1/3.0 interface-type p2mp
   user@R1# set group red neighbor fe-0/1/3.0 peer 23.1.1.2
   user@R1# set group red neighbor fe-0/1/3.0 peer 23.1.1.3
   user@R1# set group red neighbor fe-0/1/3.0 demand-circuit
   user@R1# set group red neighbor fe-0/1/3.0 max-retrans-time 10
   ```

   Similarly, configure Devices R2 and R3, omitting the **peer *address*** configuration statement.

Copyright © 2017, Juniper Networks, Inc.

NOTE: Configuring max-retrans-time is optional. In the absence of this configuration statement, the default retransmission time of 180 seconds is configured.

The configuration used in this example is for a RIP network with demand circuits. To configure RIP for networks without demand circuits, exclude the *demand-circuit* and max-retrans-time statements from the configuration and check the resulting output. For more information about configuring RIP demand circuits, see "Example: Configuring RIP Demand Circuits" on page 47.

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show policy-options**, and **show protocols rip** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
[edit]
user@R1# show interfaces
fe-0/1/3 {
   unit 0 {
      family inet {
         address 23.1.1.1/24;
      }
   }
}
```

```
user@R1# show protocols rip
traceoptions {
   file R1.log size 4m world-readable;
   flag all detail;
}
group red {
   export accept-rip-routes;
   neighbor fe-0/1/3.0 {
      interface-type p2mp;
      peer 23.1.1.2;
      peer 23.1.1.3;
      demand-circuit;
      max-retrans-time 10;
   }
}
```

```
user@R1# show policy-options
policy-statement accept-rip-routes {
   term from-direct {
      from protocol direct;
      then accept;
   }
   term from-rip {
      from protocol rip;
```

Copyright © 2017, Juniper Networks, Inc.

```
        then accept;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Verifying the Point-to-Multipoint RIP Network

Purpose Verify that the RIP network is functional with the point-to-multipoint feature configured.

Copyright © 2017, Juniper Networks, Inc.

Action From operational mode, run the **show rip neighbor** command.

```
user@R1> show rip neighbor

                  Local  Source          Destination     Send   Receive   In
Neighbor          State  Address         Address         Mode   Mode     Met
--------          -----  -------         -----------     ----   -------  ---
fe-0/1/3.0(DC)       Up 23.1.1.1        23.1.1.2        unicast unicast    1
fe-0/1/3.0(DC)       Up 23.1.1.1        23.1.1.3        unicast unicast    1
```

From operational mode, run the **show rip statistics peer *address*** command.

```
user@R1> show rip statistics peer 23.1.1.2

RIPv2 info: port 520; holddown 120s.
    rts learned  rts held down  rqsts dropped  resps dropped
              3              0              0              0

fe-0/1/3.0 Peer-Ip 23.1.1.2:  2 routes learned; 3 routes advertised; timeout 180s;
 update interval 0s
Counter                       Total   Last 5 min  Last minute
-------                   -----------  -----------  -----------
Updates Sent                        0            0            0
Triggered Updates Sent              3            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              2            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
none                                3            0            0


user@R1> show rip statistics peer 23.1.1.3

RIPv2 info: port 520; holddown 120s.
    rts learned  rts held down  rqsts dropped  resps dropped
              3              0              0              0

fe-0/1/3.0 Peer-Ip 23.1.1.3:  2 routes learned; 3 routes advertised; timeout 180s;
 update interval 0s
Counter                       Total   Last 5 min  Last minute
-------                   -----------  -----------  -----------
Updates Sent                        0            0            0
Triggered Updates Sent              3            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              2            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
```

Copyright © 2017, Juniper Networks, Inc.

```
RIP Requests Ignored                0           0           0
none                                3           0           0
```

Meaning The RIP network is up and running with the point-to-multipoint feature configured.

Related Documentation

- demand-circuit on page 187
- dynamic-peers on page 188
- Example: Configuring RIP Demand Circuits on page 47
- interface-type (Protocols RIP) on page 195
- peer on page 201

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 9

Configuring RIP Import Policy

- Understanding RIP Import Policy on page 85
- Example: Applying Policies to RIP Routes Imported from Neighbors on page 85

Understanding RIP Import Policy

The default RIP import policy is to accept all received RIP routes that pass a sanity check. To filter routes being imported by the local routing device from its neighbors, include the **import** statement, and list the names of one or more policies to be evaluated. If you specify more than one policy, they are evaluated in order (first to last) and the first matching policy is applied to the route. If no match is found, the local routing device does not import any routes. Note that the functionality of applying policies to RIP routes imported from neighbors described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

Release History Table

Release	Description
15.1X49	Note that the functionality of applying policies to RIP routes imported from neighbors described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

Related Documentation

- Example: Applying Policies to RIP Routes Imported from Neighbors on page 85

Example: Applying Policies to RIP Routes Imported from Neighbors

This example shows how to configure an import policy in a RIP network.

- Requirements on page 86
- Overview on page 86
- Configuration on page 86
- Verification on page 89

Copyright © 2017, Juniper Networks, Inc.

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, Device R1 has an import policy that accepts the 10/8 and 192.168/16 RIP routes and rejects all other RIP routes. This means that the 172.16/16 RIP routes are excluded from Device R1's routing table.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

Figure 14 on page 86 shows the topology used in this example.

Figure 14: RIP Import Policy Network Topology

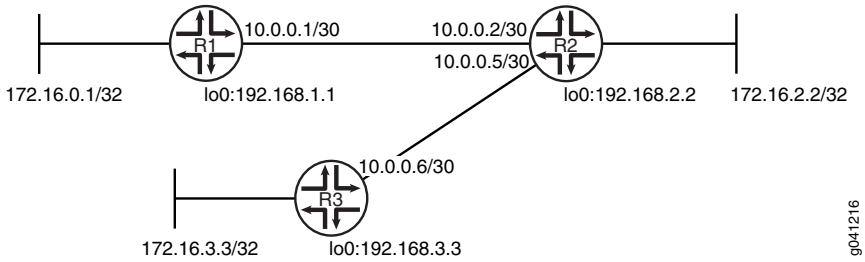

"CLI Quick Configuration" on page 86 shows the configuration for all of the devices in Figure 14 on page 86. The section "Step-by-Step Procedure" on page 87 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip import rip-import
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set policy-options policy-statement rip-import term 1 from protocol rip
set policy-options policy-statement rip-import term 1 from route-filter 10.0.0.0/8 orlonger
set policy-options policy-statement rip-import term 1 from route-filter 192.168.0.0/16
    orlonger
```


Copyright © 2017, Juniper Networks, Inc.

```
set policy-options policy-statement rip-import term 1 then accept
set policy-options policy-statement rip-import term 2 then reject
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIP import policy:

1. Configure the network interfaces.

   This example shows multiple loopback interface addresses to simulate attached networks.

   ```
   [edit interfaces]
   user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30


   user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
   user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
   ```

2. Create the RIP group and add the interface.

   To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled.

   You do not need to enable RIP on the loopback interface.

   ```
   [edit protocols rip group rip-group]
   user@R1# set neighbor fe-1/2/0.1
   ```

Copyright © 2017, Juniper Networks, Inc.

3. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

4. Apply the routing policy.

   In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

5. Configure the import policy.

```
[edit policy-options policy-statement rip-import]
user@R1# set term 1 from protocol rip
user@R1# set term 1 from route-filter 10.0.0.0/8 orlonger
user@R1# set term 1 from route-filter 192.168.0.0/16 orlonger
user@R1# set term 1 then accept
user@R1# set term 2 then reject
```

6. Apply the import policy.

```
[edit protocols rip]
user@R1# set import rip-import
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
    unit 1 {
        family inet {
            address 10.0.0.1/30;
        }
    }
}
lo0 {
    unit 1 {
        family inet {
            address 172.16.0.1/32;
            address 192.168.1.1/32;
        }
    }
}
```

```
user@R1# show protocols
rip {
    import rip-import;
    group rip-group {
```


Copyright © 2017, Juniper Networks, Inc.

```
      export advertise-routes-through-rip;
      neighbor fe-1/2/0.1;
    }
  }
```

```
user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
policy-statement rip-import {
  term 1 {
    from {
      protocol rip;
      route-filter 10.0.0.0/8 orlonger;
      route-filter 192.168.0.0/16 orlonger;
    }
    then accept;
  }
  term 2 {
    then reject;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Looking at the Routes That Device R2 Is Advertising to Device R1 on page 89
- Looking at the Routes That Device R1 Is Receiving from Device R2 on page 90
- Checking the Routing Table on page 90
- Testing the Import Policy on page 91

Looking at the Routes That Device R2 Is Advertising to Device R1

Purpose Verify that Device R2 is sending the expected routes.

Action From operational mode, enter the **show route advertising-protocol rip** command.

```
user@R2> show route advertising-protocol rip 10.0.0.2

inet.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30             *[Direct/0] 2d 01:17:44
                         >    via fe-1/2/0.5
172.16.2.2/32           *[Direct/0] 2d 04:09:52
                         >    via lo0.2
172.16.3.3/32           *[RIP/100] 23:40:02, metric 2, tag 0
                         > to 10.0.0.6 via fe-1/2/0.5
```

Copyright © 2017, Juniper Networks, Inc.

```
192.168.2.2/32     *[Direct/0] 2d 04:09:52
                    >    via lo0.2
192.168.3.3/32     *[RIP/100] 23:40:02, metric 2, tag 0
                    > to 10.0.0.6 via fe-1/2/0.5
```

Meaning Device R2 is sending 172.16/16 routes to Device R1.

Looking at the Routes That Device R1 Is Receiving from Device R2

Purpose Verify that Device R1 is receiving the expected routes.

Action From operational mode, enter the **show route receive-protocol rip** command.

```
user@R1> show route receive-protocol rip 10.0.0.2
inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30        *[RIP/100] 01:06:03, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32     *[RIP/100] 01:06:03, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32     *[RIP/100] 01:06:03, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
```

Meaning The output shows that the 172.16/16 routes are excluded.

Checking the Routing Table

Purpose Verify that the routing table is populated with the expected routes.

Action From operational mode, enter the **show route protocol rip** command.

```
user@R1> show route protocol rip

inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.4/30        *[RIP/100] 00:54:34, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.2.2/32     *[RIP/100] 00:54:34, metric 2, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
192.168.3.3/32     *[RIP/100] 00:54:34, metric 3, tag 0
                    > to 10.0.0.2 via fe-1/2/0.1
224.0.0.9/32       *[RIP/100] 00:49:00, metric 1
                          MultiRecv
```

Meaning The output shows that the routes have been learned from Device R2 and Device R3.

Copyright © 2017, Juniper Networks, Inc.

If you delete or deactivate the import policy, the routing table contains the 172.16/16 routes.

Testing the Import Policy

Purpose By using the **test policy** command, monitor the number of rejected prefixes.

Action From operational mode, enter the **test policy rip-import 172.16/16** command.

```
user@R1> test policy rip-import 172.16/16
Policy rip-import: 0 prefix accepted, 1 prefix rejected
```

Meaning The output shows that the policy rejected one prefix.

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 10

Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets

- Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets on page 93
- Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 93
- Verifying the Exchange of RIP Messages on page 97

Understanding the Sending and Receiving of RIPv1 and RIPv2 Packets

RIP version 1 (RIPv1) and RIP version 2 (RIPv2) can run simultaneously. This might make sense when you are migrating a RIPv1 network to a RIPv2 network. This also allows interoperation with a device that supports RIPv1 but not RIPv2.

By default, when RIP is enabled on an interface, Junos OS receives both RIPv1 and RIPv2 packets and sends only RIPv2 packets. You can configure this behavior by including the **send** and **receive** statements in the RIP configuration.Note that the functionality of configuring the sending and receiving of RIPv1 and RIPv2 packets described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

Release History Table

Release	Description
15.149	Note that the functionality of configuring the sending and receiving of RIPv1 and RIPv2 packets described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

Related Documentation

- Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 93

Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets

This example shows how to configure whether the RIP update messages conform to RIP version 1 (RIPv1) only, to RIP version 2 (RIPv2) only, or to both versions. You can also disable the sending or receiving of update messages.

- Requirements on page 94
- Overview on page 94

Copyright © 2017, Juniper Networks, Inc.

- Configuration on page 94
- Verification on page 96

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

By default, when RIP is enabled on an interface, Junos OS receives both RIPv1 and RIPv2 packets and sends only RIPv2 packets.

Figure 15 on page 94 shows the topology used in this example.

Figure 15: Sending and Receiving RIPv1 and RIPv2 Packets Network Topology

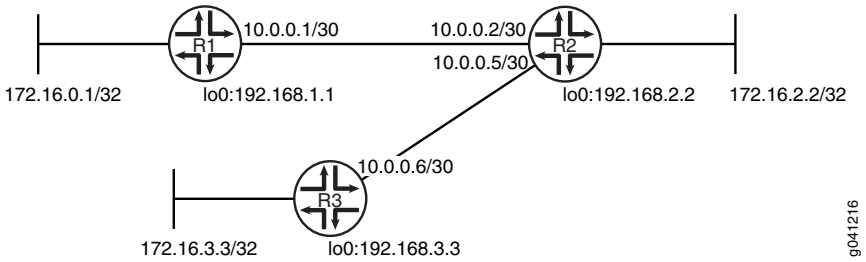

In this example, Device R1 is configured to receive only RIPv2 packets.

"CLI Quick Configuration" on page 94 shows the configuration for all of the devices in Figure 15 on page 94. The section "Step-by-Step Procedure" on page 95 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1 receive version-2
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
```


Copyright © 2017, Juniper Networks, Inc.

```
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIP packet versions that can be received:

1. Configure the network interfaces.

   ```
   [edit interfaces]
   user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

   user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
   user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
   ```

2. Create the RIP groups and add the interfaces.

   To configure RIP in Junos OS, you must configure one or more groups that contain the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

   For the interface that is facing Device R2, the **receive version-2** setting causes this interface to accept only RIPv2 packets.

   ```
   [edit protocols rip group rip-group]
   user@R1# set neighbor fe-1/2/0.1 receive version-2
   ```

3. Create the routing policy to advertise both direct and RIP-learned routes.

   ```
   [edit policy-options policy-statement advertise-routes-through-rip term 1]
   user@R1# set from protocol direct
   user@R1# set from protocol rip
   user@R1# set then accept
   ```

Copyright © 2017, Juniper Networks, Inc.

4. Apply the routing policy.

   In Junos OS, you can only apply RIP export policies at the group level.

   ```
   [edit protocols rip group rip-group]
   user@R1# set export advertise-routes-through-rip
   ```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    family inet {
      address 10.0.0.1/30;
    }
  }
}
lo0 {
  unit 1 {
    family inet {
      address 172.16.0.1/32;
      address 192.168.1.1/32;
    }
  }
}
```

```
user@R1# show protocols
rip {
  group rip-group {
    export advertise-routes-through-rip;
    neighbor fe-1/2/0.1 {
      receive version-2;
    }
  }
}
```

```
user@R1# show policy-options
policy-statement advertise-routes-through-rip {
  term 1 {
    from protocol [ direct rip ];
    then accept;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Copyright © 2017, Juniper Networks, Inc.

Verifying That the Receive Mode Is Set to RIPv2 Only

Purpose Make sure that the interfacing Device R2 is configured to receive only RIPv2 packets, instead of both RIPv1 and RIPv2 packets.

Action From operational mode, enter the **show rip neighbor** command.

```
user@R1> show rip neighbor

                    Local  Source          Destination     Send   Receive   In
Neighbor            State  Address         Address         Mode   Mode      Met
--------            -----  -------         -----------     ----   -------   ---
fe-1/2/0.1             Up 10.0.0.1        224.0.0.9       mcast  v2 only     1
```

Meaning In the output, the **Receive Mode** field displays **v2 only**. The default **Receive Mode** is **both**.

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; holddown 120s.
    rts learned  rts held down  rqsts dropped  resps dropped
              10              0              0              0

t1-0/0/2.0:  0 routes learned; 13 routes advertised; timeout 120s; update interval
 45s
Counter                         Total   Last 5 min  Last minute
-------                     -----------  ----------  -----------
Updates Sent                       2855          11            2
Triggered Updates Sent                5           0            0
Responses Sent                        0           0            0
Bad Messages                          0           0            0
RIPv1 Updates Received                0           0            0
RIPv1 Bad Route Entries               0           0            0
RIPv1 Updates Ignored                 0           0            0
RIPv2 Updates Received               41           0            0
RIPv2 Bad Route Entries               0           0            0
RIPv2 Updates Ignored                 0           0            0
Authentication Failures               0           0            0
RIP Requests Received                 0           0            0
RIP Requests Ignored                  0           0            0

ge-0/0/1.0:  10 routes learned; 3 routes advertised; timeout 180s; update interval
 30s
Counter                         Total   Last 5 min  Last minute
-------                     -----------  ----------  -----------
Updates Sent                       2855          11            2
Triggered Updates Sent                3           0            0
```

Copyright © 2017, Juniper Networks, Inc.

```
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
```

Meaning The output shows the number of RIP routes learned. It also shows the number of RIP updates sent and received on the RIP-enabled interfaces. Verify the following information:

- The number of RIP routes learned matches the number of expected routes learned. Subnets learned by direct connectivity through an outgoing interface are not listed as RIP routes.
- RIP updates are being sent on each RIP-enabled interface. If no updates are being sent, the routing policy might not be configured to export routes.
- RIP updates are being received on each RIP-enabled interface. If no updates are being received, the routing policy might not be configured to export routes on the host connected to that subnet. The lack of updates might also indicate an authentication error.

Related Documentation

- Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 93

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 11

Redistributing Routes Between Two RIP Instances

- Understanding Route Redistribution Among RIP Instances on page 99
- Example: Redistributing Routes Between Two RIP Instances on page 100

Understanding Route Redistribution Among RIP Instances

You can redistribute routes among RIP processes. Another way to say this is to export RIP routes from one RIP instance to other RIP instances.

In Junos OS, route redistribution among routing instances is accomplished by using routing table groups, also called RIB groups. Routing table groups allow you to import and export routes from a protocol within one routing table into another routing table. Note that the functionality of redistributing routes among RIP instances described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

NOTE: In contrast, the policy-based import and export functions allow you import and export routes between different protocols within the same routing table.

Consider the following partial example:

```
protocols {
   rip {
      rib-group inet-to-voice;
   }
}
routing-instances {
   voice {
      protocols {
         rip {
            rib-group voice-to-inet;
         }
      }
   }
}
routing-options {
   rib-groups {
```

Copyright © 2017, Juniper Networks, Inc.

```
        inet-to-voice {
          import-rib [ inet.0 voice.inet.0 ];
        }
        voice-to-inet {
          import-rib [ voice.inet.0 inet.0 ];
        }
      }
    }
```

The way to read the **import-rib** statement is as follows. Take the routes from the protocol (RIP, in this case), and import them into the primary (or local) routing table and also into any other routing tables listed after this. The primary routing table is the routing table where the routing table group is being used. That would be either **inet.0** if used in the main routing instance or **voice.inet.0** if used within the routing instance. In the **inet-to-voice** routing table group, **inet.0** is listed first because this routing table group is used in the main routing instance. In the **voice-to-inet** routing table group, **voice.inet.0** is listed first because this routing table group is used in the voice routing instance.

Release History Table

Release	Description
15.1X49	Note that the functionality of redistributing routes among RIP instances described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

Related Documentation

- Example: Redistributing Routes Between Two RIP Instances on page 100

Example: Redistributing Routes Between Two RIP Instances

This example shows how to configure a RIP routing instance and control the redistribution of RIP routes between the routing instance and the master instance.

- Requirements on page 100
- Overview on page 100
- Configuration on page 101
- Verification on page 104

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

When you create a routing instance called voice, Junos OS creates a routing table called **voice.inet.0**. The example shows how to install routes learned through the master RIP instance into the **voice.inet.0** routing table. The example also shows how to install routes learned through the voice routing instance into **inet.0**. This is done by configuring routing table groups. RIP routes are installed into each routing table that belongs to a routing table group.

Copyright © 2017, Juniper Networks, Inc.

Figure 16 on page 101 shows the topology used in this example.

Figure 16: Redistributing Routes Between RIP Instances Network Topology

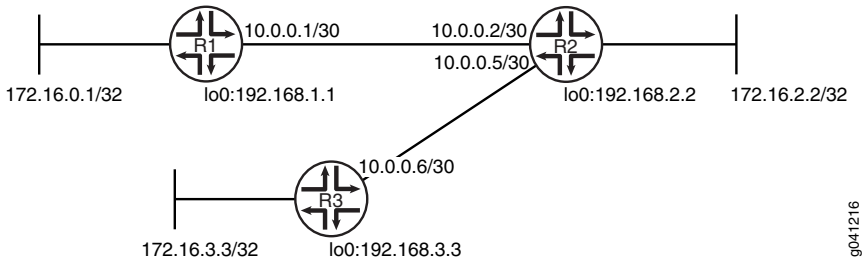

"CLI Quick Configuration" on page 101 shows the configuration for all of the devices in Figure 16 on page 101. The section "Step-by-Step Procedure" on page 102 describes the steps on Device R2.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, copy and paste the commands into the CLI at the **[edit]** hierarchy level, and then enter **commit** from configuration mode.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip group to-R2 export advertise-routes-through-rip
set protocols rip group to-R2 neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
set protocols rip rib-group inet-to-voice
set protocols rip group to-R3 export advertise-routes-through-rip
set protocols rip group to-R3 neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
set routing-instances voice protocols rip group to-R1 export advertise-routes-through-rip
set routing-instances voice interface fe-1/2/0.2
set routing-instances voice protocols rip rib-group voice-to-inet
set routing-instances voice protocols rip group to-R1 neighbor fe-1/2/0.2
set routing-options rib-groups inet-to-voice import-rib inet.0
set routing-options rib-groups inet-to-voice import-rib voice.inet.0
set routing-options rib-groups voice-to-inet import-rib voice.inet.0
```

Copyright © 2017, Juniper Networks, Inc.

```
set routing-options rib-groups voice-to-inet import-rib inet.0
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group to-R2 export advertise-routes-through-rip
set protocols rip group to-R2 neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To redistribute RIP routes between routing instances:

1. Configure the network interfaces.

```
[edit interfaces]
user@R2# set fe-1/2/0 unit 2 family inet address 10.0.0.2/30

user@R2# set fe-1/2/1 unit 5 family inet address 10.0.0.5/30

user@R2# set lo0 unit 2 family inet address 192.168.2.2/32
user@R2# set lo0 unit 2 family inet address 172.16.2.2/32
```

2. Create the routing instance, and add one or more interfaces to the routing instance.

```
[edit routing-instances voice]
user@R2# set interface fe-1/2/0.2
```

3. Create the RIP groups and add the interfaces.

```
[edit protocols rip group to-R3]
user@R2# set neighbor fe-1/2/1.5

[edit routing-instances voice protocols rip group to-R1]
user@R2# set neighbor fe-1/2/0.2
```

4. Create the routing table groups.

```
[edit routing-options rib-groups]
user@R2# set inet-to-voice import-rib inet.0
user@R2# set inet-to-voice import-rib voice.inet.0

user@R2# set voice-to-inet import-rib voice.inet.0
user@R2# set voice-to-inet import-rib inet.0
```

5. Apply the routing table groups.

Copyright © 2017, Juniper Networks, Inc.

```
[edit protocols rip]
user@R2# set rib-group inet-to-voice

[edit routing-instances voice protocols rip]
user@R2# set rib-group voice-to-inet
```

6. Create the routing policy to advertise both direct and RIP-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-rip term 1]
user@R2# set from protocol direct
user@R2# set from protocol rip
user@R2# set then accept
```

7. Apply the routing policy.

   In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group to-R3]
user@R2# set export advertise-routes-through-rip

[edit routing-instances voice protocols rip group to-R1]
user@R2# set export advertise-routes-through-rip
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, **show routing-instances**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R2# show interfaces
fe-1/2/0 {
    unit 2 {
        family inet {
            address 10.0.0.2/30;
        }
    }
}
fe-1/2/1 {
    unit 5 {
        family inet {
            address 10.0.0.5/30;
        }
    }
}
lo0 {
    unit 2 {
        family inet {
            address 192.168.2.2/32;
            address 172.16.2.2/32;
        }
    }
}

user@R2# show protocols
```

Copyright © 2017, Juniper Networks, Inc.

```
rip {
   rib-group inet-to-voice;
   group to-R3 {
      export advertise-routes-through-rip;
      neighbor fe-1/2/1.5;
   }
}
```

```
user@R2# show policy-options
policy-statement advertise-routes-through-rip {
   term 1 {
      from protocol [ direct rip ];
      then accept;
   }
}
```

```
user@R2# show routing-instances
voice {
   interface fe-1/2/0.2;
   protocols {
      rip {
         rib-group voice-to-inet;
         group to-R1 {
            export advertise-routes-through-rip;
            neighbor fe-1/2/0.2;
         }
      }
   }
}
```

```
user@R2# show routing-options
rib-groups {
   inet-to-voice {
      import-rib [ inet.0 voice.inet.0 ];
   }
   voice-to-inet {
      import-rib [ voice.inet.0 inet.0 ];
   }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Checking the Routing Tables

Purpose Make sure that the routing tables contain the expected routes.

Action From operational mode, enter the **show route protocol rip** command.

```
user@R2> show route protocol rip
inet.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
```


Copyright © 2017, Juniper Networks, Inc.

```
172.16.0.1/32      *[RIP/100] 01:58:14, metric 2, tag 0
                    > to 10.0.0.1 via fe-1/2/0.2
172.16.3.3/32      *[RIP/100] 02:06:03, metric 2, tag 0
                    > to 10.0.0.6 via fe-1/2/0.5
192.168.1.1/32     *[RIP/100] 01:58:14, metric 2, tag 0
                    > to 10.0.0.1 via fe-1/2/0.2
192.168.3.3/32     *[RIP/100] 02:06:03, metric 2, tag 0
                    > to 10.0.0.6 via fe-1/2/0.5
224.0.0.9/32       *[RIP/100] 01:44:13, metric 1
                          MultiRecv

voice.inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.1/32      *[RIP/100] 02:06:03, metric 2, tag 0
                    > to 10.0.0.1 via fe-1/2/0.2
172.16.3.3/32      *[RIP/100] 01:58:14, metric 2, tag 0
                    > to 10.0.0.6 via fe-1/2/0.5
192.168.1.1/32     *[RIP/100] 02:06:03, metric 2, tag 0
                    > to 10.0.0.1 via fe-1/2/0.2
192.168.3.3/32     *[RIP/100] 01:58:14, metric 2, tag 0
                    > to 10.0.0.6 via fe-1/2/0.5
224.0.0.9/32       *[RIP/100] 01:44:13, metric 1
                          MultiRecv
```

Meaning The output shows that both routing tables contain all of the RIP routes.

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 12

Tracing RIP Protocol Traffic

- Understanding RIP Trace Operations on page 107
- Example: Tracing RIP Protocol Traffic on page 108

Understanding RIP Trace Operations

You can trace various types of RIP protocol traffic to help debug RIP protocol issues. Note that the functionality of tracing RIP protocol traffic described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

To trace RIP protocol traffic, include the **traceoptions** statement at the **[edit protocols rip]** hierarchy level:

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

You can specify the following RIP protocol-specific trace options using the **flag** statement:

- **auth**—RIP authentication
- **error**—RIP error packets
- **expiration**—RIP route expiration processing
- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop active routing synchronization events
- **packets**—All RIP packets
- **request**—RIP information packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

You can optionally specify one or more of the following flag modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

Copyright © 2017, Juniper Networks, Inc.

NOTE: Use the **detail** flag modifier with caution as this may cause the CPU to become very busy.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the RIP protocol using the **traceoptions flag** statement included at the **[edit protocols rip]** hierarchy level:

- **all**—All tracing operations
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

NOTE: Use the **trace flag all** with caution because this may cause the CPU to become very busy.

Release History Table

Release	Description
15.1X49	Note that the functionality of tracing RIP protocol traffic described in this topic is not supported in Junos OS Releases 15.1X49, 15.1X49-D30, or 15.1X49-D40.

Related Documentation

- Example: Tracing RIP Protocol Traffic on page 108

Example: Tracing RIP Protocol Traffic

This example shows how to trace RIP protocol operations.

- Requirements on page 109
- Overview on page 109
- Configuration on page 109
- Verification on page 112

Copyright © 2017, Juniper Networks, Inc.

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, Device R1 is set to trace routing information updates.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIP.

Figure 17 on page 109 shows the topology used in this example.

Figure 17: RIP Trace Operations Network Topology

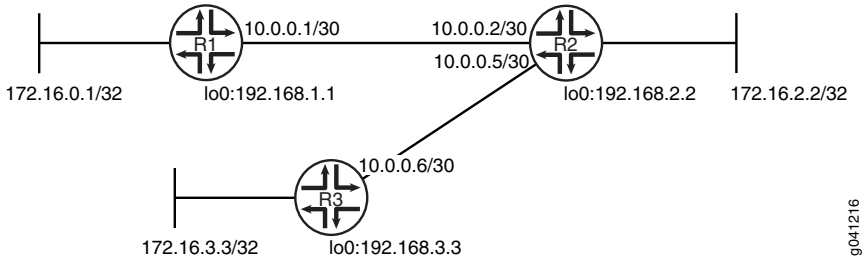

"CLI Quick Configuration" on page 109 shows the configuration for all of the devices in Figure 17 on page 109. The section "Step-by-Step Procedure" on page 110 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 family inet address 10.0.0.1/30
set interfaces lo0 unit 1 family inet address 172.16.0.1/32
set interfaces lo0 unit 1 family inet address 192.168.1.1/32
set protocols rip traceoptions file rip-trace-file
set protocols rip traceoptions flag route
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
    rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet address 10.0.0.2/30
set interfaces fe-1/2/1 unit 5 family inet address 10.0.0.5/30
set interfaces lo0 unit 2 family inet address 192.168.2.2/32
set interfaces lo0 unit 2 family inet address 172.16.2.2/32
```

Copyright © 2017, Juniper Networks, Inc.

```
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.2
set protocols rip group rip-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet address 10.0.0.6/30
set interfaces lo0 unit 3 family inet address 192.168.3.3/32
set interfaces lo0 unit 3 family inet address 172.16.3.3/32
set protocols rip group rip-group export advertise-routes-through-rip
set protocols rip group rip-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-rip term 1 from protocol
  rip
set policy-options policy-statement advertise-routes-through-rip term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the RIP update interval:

1. Configure the network interfaces.

   This example shows multiple loopback interface addresses to simulate attached networks.

   ```
   [edit interfaces]
   user@R1# set fe-1/2/0 unit 1 family inet address 10.0.0.1/30

   user@R1# set lo0 unit 1 family inet address 172.16.0.1/32
   user@R1# set lo0 unit 1 family inet address 192.168.1.1/32
   ```

2. Configure the RIP group, and add the interface to the group.

   To configure RIP in Junos OS, you must configure a group that contains the interfaces on which RIP is enabled. You do not need to enable RIP on the loopback interface.

   ```
   [edit protocols rip group rip-group]
   user@R1# set neighbor fe-1/2/0.1
   ```

3. Configure RIP tracing operations.

   ```
   [edit protocols rip traceoptions]
   user@R1# set file rip-trace-file
   user@R1# set flag route
   ```

4. Create the routing policy to advertise both direct and RIP-learned routes.

   ```
   [edit policy-options policy-statement advertise-routes-through-rip term 1]
   ```


Copyright © 2017, Juniper Networks, Inc.

```
user@R1# set from protocol direct
user@R1# set from protocol rip
user@R1# set then accept
```

5. Apply the routing policy.

   In Junos OS, you can only apply RIP export policies at the group level.

```
[edit protocols rip group rip-group]
user@R1# set export advertise-routes-through-rip
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
    unit 1 {
        family inet {
            address 10.0.0.1/30;
        }
    }
}
lo0 {
    unit 1 {
        family inet {
            address 172.16.0.1/32;
            address 192.168.1.1/32;
        }
    }
}
```

```
user@R1# show protocols
rip {
    traceoptions {
        file rip-trace-file;
        flag route;
    }
    group rip-group {
        export advertise-routes-through-rip;
        neighbor fe-1/2/0.1;
    }
}
```

```
user@R1# show policy-options
policy-statement advertise-routes-through-rip {
    term 1 {
        from protocol [ direct rip ];
        then accept;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2017, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

Checking the Log File

Purpose Make sure that the RIP route updates are logged in the configured log file.

Action

1. Deactivate the extra loopback interface address on Device R3.

   ```
   [edit interfaces lo0 unit 3 family inet]
   user@R3# deactivate address 172.16.3.3/32
   user@R3# commit
   ```

2. From operational mode on Device R1, enter the **show log rip-trace-file** command with the **| match 172.16.3.3** option.

   ```
   user@R1> show log rip-trace-file | match 172.16.3.3
   Mar  1 11:39:53.975192 Setting RIPv2 rtbit on route 172.16.3.3/32, tsi =
   0xbb69228
   Mar  1 11:39:59.847118  172.16.3.3/32: metric-in: 16, change: 3 -> 16; # gw:
   1, pkt_upd_src 10.0.0.2, inx: 0, rte_upd_src 10.0.0.2
   Mar  1 11:39:59.847568 CHANGE   172.16.3.3/32       nhid 591 gw 10.0.0.2
      RIP      pref 100/0 metric 3/0 fe-1/2/0.1 <Delete Int>
   Mar  1 11:39:59.847629 Best route to 172.16.3.3/32 got deleted. Doing route calculation
    on the stored rte-info
   ```

Meaning The output shows that the route to 172.16.3.3/32 was deleted.

Related Documentation

Copyright © 2017, Juniper Networks, Inc.

PART 3

Configuring RIPng

- Configuring a Basic RIPng Network on page 115
- Applying Policies to RIPng Routes on page 123
- Configuring Traffic Control with Metrics in a RIPng Network on page 137
- Configuring RIPng Timers on page 145
- Tracing RIPng Protocol Traffic on page 153

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 13

Configuring a Basic RIPng Network

- Understanding Basic RIPng Routing on page 115
- Example: Configuring a Basic RIPng Network on page 115

Understanding Basic RIPng Routing

By default, RIP next generation (RIPng) routes are not redistributed. You must configure export policy to redistribute RIPng routes.

To have a router exchange routes with other routers, you must configure RIPng groups and neighbors. RIPng routes received from routers not configured as RIPng neighbors are ignored. Likewise, RIPng routes are advertised only to routers configured as RIPng neighbors.

Related Documentation

- RIPng Overview on page 11
- Example: Configuring a Basic RIPng Network on page 115

Example: Configuring a Basic RIPng Network

This example shows how to configure a basic RIPng network.

- Requirements on page 115
- Overview on page 115
- Configuration on page 116
- Verification on page 118

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, you configure a basic RIPng network, create a RIPng group called **ripng-group**, and add the directly connected interfaces to the RIPng group. Then you configure a routing policy to advertise direct routes using the policy statement **advertise-routes-through-ripng**.

Copyright © 2017, Juniper Networks, Inc.

By default, Junos OS does not advertise RIPng routes, not even routes that are learned through RIPng. To advertise RIPng routes, you must configure and apply an export routing policy that advertises RIPng-learned and direct routes.

To use RIPng on the device, you must configure RIPng on all of the RIPng interfaces within the network. Figure 18 on page 116 shows the topology used in this example.

Figure 18: Sample RIPng Network Topology

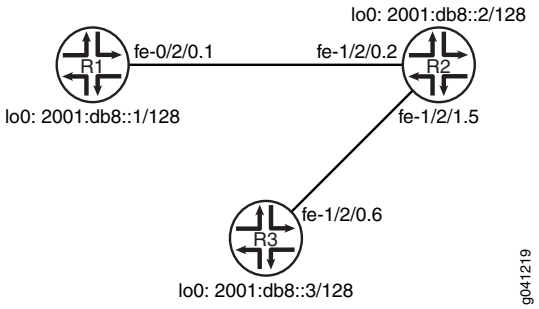

"CLI Quick Configuration" on page 116 shows the configuration for all of the devices in Figure 18 on page 116. The section "Step-by-Step Procedure" on page 117 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 description to-R2
set interfaces fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces lo0 unit 1 family inet6 address 2001:db8::1/128
set protocols ripng group ripng-group export advertise-routes-through-ripng
set protocols ripng group ripng-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 description to-R1
set interfaces fe-1/2/0 unit 2 family inet6 address 2001:db8:0:2::/64 eui-64
set interfaces fe-1/2/1 unit 5 description to-R3
set interfaces fe-1/2/1 unit 5 family inet6 address 2001:db8:0:3::/64 eui-64
set interfaces lo0 unit 2 family inet6 address 2001:db8::2/128
set protocols ripng group ripng-group export advertise-routes-through-ripng
set protocols ripng group ripng-group neighbor fe-1/2/0.2
set protocols ripng group ripng-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  ripng
```


Copyright © 2017, Juniper Networks, Inc.

```
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 description to-R2
set interfaces fe-1/2/0 unit 6 family inet6 address 2001:db8:0:4::/64 eui-64
set interfaces lo0 unit 3 family inet6 address 2001:db8::3/128
set protocols ripng group ripng-group export advertise-routes-through-ripng
set protocols ripng group ripng-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a basic RIPng network:

1. Configure the network interfaces.

   Use the **eui-64** statement to automatically generate the host portion of the interface address and the link-local address.

   For the loopback interface, you must assign a 128-bit address.

   ```
   [edit interfaces]
   user@R1# set fe-1/2/0 unit 1 description to-R2
   user@R1# set fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64

   user@R1# set lo0 unit 1 family inet6 address 2001:db8::1/128
   ```

2. Create the RIPng group and add the interface.

   To configure RIPng in Junos OS, you must configure a group that contains the interfaces on which RIPng is enabled. You do not need to enable RIPng on the loopback interface.

   ```
   [edit protocols ripng group ripng-group]
   user@R1# set neighbor fe-1/2/0.1
   ```

3. Create the routing policy to advertise both direct and RIPng-learned routes.

   ```
   [edit policy-options policy-statement advertise-routes-through-ripng term 1]
   user@R1# set from protocol direct
   user@R1# set from protocol ripng
   user@R1# set then accept
   ```

4. Apply the routing policy.

   In Junos OS, you can only apply RIPng export policies at the group level.

   ```
   [edit protocols ripng group ripng-group]
   user@R1# set export advertise-routes-through-ripng
   ```

Copyright © 2017, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
    unit 1 {
        description to-R2;
        family inet6 {
            address 2001:db8:0:1::/64 {
                eui-64;
            }
        }
    }
}
lo0 {
    unit 1 {
        family inet6 {
            address 2001:db8::1/128;
        }
    }
}
```

```
user@R1# show protocols
ripng {
    group ripng-group {
        export advertise-routes-through-ripng;
        neighbor fe-1/2/0.1;
    }
}
```

```
user@R1# show policy-options
policy-statement advertise-routes-through-ripng {
    term 1 {
        from protocol [ direct ripng ];
        then accept;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Checking the Routing Table on page 119
- Checking the Interface Addresses on page 119
- Looking at the Routes That Device R1 Is Advertising to Device R2 on page 120
- Verifying the RIPng-Enabled Interfaces on page 120
- Looking at the Routes That Device R1 Is Receiving from Device R2 on page 120
- Verifying the Exchange of RIPng Messages on page 121
- Verifying Reachability of All Hosts in the RIPng Network on page 122

Copyright © 2017, Juniper Networks, Inc.

Checking the Routing Table

Purpose Verify that the routing table is populated with the expected routes.

Action From operational mode, enter the **show route protocol ripng** command.

```
user@R1> show route protocol ripng
inet6.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::2/128    *[RIPng/100] 3d 19:24:43, metric 2, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
2001:db8::3/128    *[RIPng/100] 3d 19:24:40, metric 3, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
2001:db8:0:2::/64  *[RIPng/100] 3d 19:24:43, metric 2, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
2001:db8:0:3::/64  *[RIPng/100] 3d 19:24:43, metric 2, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
2001:db8:0:4::/64  *[RIPng/100] 3d 19:24:40, metric 3, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
ff02::9/128        *[RIPng/100] 3d 19:24:47, metric 1
                          MultiRecv
```

Meaning The output shows that the routes have been learned from Device R2 and Device R3.

If you were to delete the **from protocol ripng** condition in the routing policy on Device R2, the remote routes from Device R3 would not be learned on Device R1.

Checking the Interface Addresses

Purpose Verify that the **eui-64** statement automatically generated the host portion of the interface address and the link-local address.

Action From operational mode, enter the **show interfaces terse** command.

```
user@R1> show interfaces terse
Interface               Admin Link Proto    Local                 Remote
fe-1/2/0
fe-1/2/0.1              up    up   inet6    2001:db8:0:1:2a0:a514:0:14c/64
                                            fe80::2a0:a514:0:14c/64
lo0
lo0.1                   up    up   inet6    2001:db8::1
                                            fe80::2a0:a50f:fc56:14c
```

Meaning The output shows that the interface address on fe-1/2/0.1 includes both the network portion (2001:db8:0:1) and the host portion (2a0:a514:0:14c).

Also, link-local (fe80) addresses are assigned to interfaces fe-1/2/0.1 and lo0.1.

Copyright © 2017, Juniper Networks, Inc.

Looking at the Routes That Device R1 Is Advertising to Device R2

Purpose Verify that Device R1 is sending the expected routes.

Action From operational mode, enter the **show route advertising-protocol ripng** command, using Device R1's link-local address as the neighbor address.

```
user@R1> show route advertising-protocol ripng fe80::2a0:a514:0:14c
inet6.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::1/128    *[Direct/0] 3d 19:45:55
                    >    via lo0.1
2001:db8:0:1::/64  *[Direct/0] 3d 19:45:55
                    >    via fe-1/2/0.1
```

Meaning Device R1 is sending routes to its directly connected networks.

Verifying the RIPng-Enabled Interfaces

Purpose Verify that all RIPng-enabled Interfaces are available and active.

Action From operational mode, enter the **show ripng neighbor** command.

```
user@R1> show ripng neighbor
                    Source                                Dest                In
Neighbor     State  Address                               Address  Send Recv Met
--------     -----  -------                               -------  ---- ---- ---
fe-1/2/0.1      Up fe80::2a0:a514:0:14c                  ff02::9  yes  yes  1
```

Meaning The output shows that the RIPng-enabled interface on Device R1 is operational.

The output also shows the link-local address that is assigned to Device R2's directly connected link-local interface.

In general for this command, the output shows a list of the RIPng neighbors that are configured on the device. 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 **State** column. A state of **Up** indicates that the link is passing RIPng traffic. A state of **Dn** indicates that the link is not passing RIPng traffic. In a point-to-point link, this state generally means that either the end point is not configured for RIPng or the link is unavailable.

Looking at the Routes That Device R1 Is Receiving from Device R2

Purpose Verify that Device R1 is receiving the expected routes.

Copyright © 2017, Juniper Networks, Inc.

Action From operational mode, enter the **show route receive-protocol ripng** command, using Device R2's directly connected link-local interface address as the neighbor address.

```
user@R1> show route receive-protocol ripng fe80::2a0:a514:0:24c
inet6.0: 12 destinations, 12 routes (12 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::2/128    *[RIPng/100] 3d 19:58:09, metric 2, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
2001:db8::3/128    *[RIPng/100] 3d 19:58:06, metric 3, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
2001:db8:0:2::/64  *[RIPng/100] 3d 19:58:09, metric 2, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
2001:db8:0:3::/64  *[RIPng/100] 3d 19:58:09, metric 2, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
2001:db8:0:4::/64  *[RIPng/100] 3d 19:58:06, metric 3, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
```

Meaning Device R1 is receiving from Device R2 all of Device R2's directly connected networks. Device R1 is also receiving from Device R2 all of Device R3's directly connected networks, which Device R2 learned from Device R3 through RIPng.

Verifying the Exchange of RIPng Messages

Purpose Verify that RIPng messages are being sent and received on all RIPng-enabled interfaces.

Action From operational mode, enter the **show ripng statistics** command.

```
user@R1> show ripng statistics
RIPng info: port 521; holddown 120s.
    rts learned  rts held down  rqsts dropped  resps dropped
              5              0               0              0

fe-1/2/0.1:  5 routes learned; 2 routes advertised; timeout 180s; update interval
 30s
Counter                         Total   Last 5 min  Last minute
-------                     ----------  ----------  -----------
Updates Sent                     11632          10            2
Triggered Updates Sent               0           0            0
Responses Sent                       0           0            0
Bad Messages                         0           0            0
Updates Received                 11634          11            2
Bad Route Entries                    0           0            0
Updates Ignored                      0           0            0
RIPng Requests Received              1           0            0
RIPng Requests Ignored               0           0            0
```

Meaning The output shows the number of RIPng routes learned. It also shows the number of RIPng updates sent and received on the RIPng-enabled interfaces. Verify the following information:

Copyright © 2017, Juniper Networks, Inc.

- The number of RIPng routes learned matches the number of expected routes learned. Subnets learned by direct connectivity through an outgoing interface are not listed as RIPng routes.
- RIPng updates are being sent on each RIPng-enabled interface. If no updates are being sent, the routing policy might not be configured to export routes.
- RIPng updates are being received on each RIPng-enabled interface. If no updates are being received, the routing policy might not be configured to export routes on the host connected to that subnet. The lack of updates might also indicate an authentication error.

Verifying Reachability of All Hosts in the RIPng Network

Purpose By using the traceroute command on each loopback address in the network, verify that all hosts in the RIPng network are reachable from each Juniper Networks device.

Action From operational mode, enter the **traceroute** command.

```
user@R1> traceroute 2001:db8::3
traceroute6 to 2001:db8::3 (2001:db8::3) from 2001:db8:0:1:2a0:a514:0:14c, 64
hops max, 12 byte packets
 1  2001:db8:0:2:2a0:a514:0:24c (2001:db8:0:2:2a0:a514:0:24c)  8.881 ms  1.175
ms  1.101 ms
 2  2001:db8::3 (2001:db8::3)  1.544 ms  2.445 ms  2.043 ms
```

Meaning Each numbered row in the output indicates a routing hop in the path to the host. The three-time increments indicate the round-trip time (RTT) between the device and the hop for each traceroute packet.

To ensure that the RIPng 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.

Related Documentation

- Understanding Basic RIP Routing on page 18
- RIP Configuration Overview on page 17

 Copyright © 2017, Juniper Networks, Inc.

CHAPTER 14

Applying Policies to RIPng Routes

- Understanding RIPng Import Policies to Filter Routes on page 123
- Example: Applying Policies to RIPng Routes Imported from Neighbors on page 123
- Example: Testing a Routing Policy with Complex Regular Expressions on page 129

Understanding RIPng Import Policies to Filter Routes

The default RIPng import policy is to accept all received RIPng routes that pass a validity check. To filter routes being imported by the local routing device from its neighbors, include the **import** statement and list the names of one or more policies to be evaluated. If you specify more than one policy, they are evaluated in order (first to last) and the first matching policy is applied to the route. If no match is found, the local routing device does not import any routes.

Related Documentation

- Example: Applying Policies to RIPng Routes Imported from Neighbors on page 123

Example: Applying Policies to RIPng Routes Imported from Neighbors

This example shows how to configure an import policy in a RIPng network.

- Requirements on page 123
- Overview on page 123
- Configuration on page 124
- Verification on page 127

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, Device R2 has several extra loopback interface addresses configured to simulate additional networks.

Copyright © 2017, Juniper Networks, Inc.

Device R1 has an import policy that accepts the fe80::/64 and 2001:db8::/64 routes and rejects all other routes. This means that the extra networks advertised by Device R2 are not accepted into Device R1's routing table.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIPng.

Figure 19 on page 124 shows the topology used in this example.

Figure 19: RIPng Import Policy Network Topology

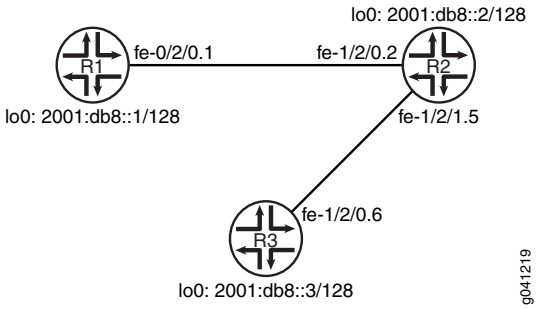

"CLI Quick Configuration" on page 124 shows the configuration for all of the devices in Figure 19 on page 124. The section "Step-by-Step Procedure" on page 125 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 description to-R2
set interfaces fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces lo0 unit 1 family inet6 address 2001:db8::1/128
set protocols ripng group ripng-group export advertise-routes-through-ripng
set protocols ripng group ripng-group neighbor fe-1/2/0.1 import ripng-import
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
    ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
set policy-options policy-statement ripng-import term 1 from route-filter fe80::/64
    orlonger
set policy-options policy-statement ripng-import term 1 from route-filter 2001:db8::/64
    orlonger
set policy-options policy-statement ripng-import term 1 then accept
set policy-options policy-statement ripng-import term 2 then reject
```

Device R2

```
set interfaces fe-1/2/0 unit 2 description to-R1
set interfaces fe-1/2/0 unit 2 family inet6 address 2001:db8:0:2::/64 eui-64
set interfaces fe-1/2/1 unit 5 description to-R3
set interfaces fe-1/2/1 unit 5 family inet6 address 2001:db8:0:3::/64 eui-64
```


Copyright © 2017, Juniper Networks, Inc.

```
set interfaces lo0 unit 2 family inet6 address 2001:db8::2/128
set interfaces lo0 unit 2 family inet6 address 2002:db8::2/128
set interfaces lo0 unit 2 family inet6 address 2002:db9::2/128
set interfaces lo0 unit 2 family inet6 address 2002:db7::2/128
set protocols ripng group ripng-group export advertise-routes-through-ripng
set protocols ripng group ripng-group neighbor fe-1/2/0.2
set protocols ripng group ripng-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 description to-R2
set interfaces fe-1/2/0 unit 6 family inet6 address 2001:db8:0:4::/64 eui-64
set interfaces lo0 unit 3 family inet6 address 2001:db8::3/128
set protocols ripng group ripng-group export advertise-routes-through-ripng
set protocols ripng group ripng-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIPng import policy:

1. Configure the network interfaces.

   This example shows multiple loopback interface addresses to simulate attached networks.

   ```
   [edit interfaces]
   user@R1# set fe-1/2/0 unit 1 description to-R2
   user@R1# set fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64


   user@R1# set lo0 unit 1 family inet6 address 2001:db8::1/128
   ```

2. Create the RIPng group and add the interface.

   To configure RIPng in Junos OS, you must configure a group that contains the interfaces on which RIPng is enabled. You do not need to enable RIPng on the loopback interface.

   ```
   [edit protocols ripng group ripng-group]
   user@R1# set neighbor fe-1/2/0.1
   ```

3. Create the routing policy to advertise both direct and RIPng-learned routes.

   ```
   [edit policy-options policy-statement advertise-routes-through-ripng term 1]
   user@R1# set from protocol direct
   ```

Copyright © 2017, Juniper Networks, Inc.

```
user@R1# set from protocol ripng
user@R1# set then accept
```

4. Apply the routing policy.

   In Junos OS, you can only apply RIPng export policies at the group level.

   ```
   [edit protocols ripng group ripng-group]
   user@R1# set export advertise-routes-through-ripng
   ```

5. Configure the import policy.

   ```
   [edit policy-options policy-statement ripng-import]
   user@R1# set term 1 from route-filter fe80::/64 orlonger
   user@R1# set term 1 from route-filter 2001:db8::/64 orlonger
   user@R1# set term 1 then accept
   user@R1# set term 2 then reject
   ```

6. Apply the import policy.

   ```
   [edit protocols ripng group ripng-group]
   user@R1# set neighbor fe-1/2/0.1 import ripng-import
   ```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
    unit 1 {
        description to-R2;
        family inet6 {
            address 2001:db8:0:1::/64 {
                eui-64;
            }
        }
    }
}
lo0 {
    unit 1 {
        family inet6 {
            address 2001:db8::1/128;
        }
    }
}
```

```
user@R1# show protocols
ripng {
    group ripng-group {
        export advertise-routes-through-ripng;
        neighbor fe-1/2/0.1 {
            import ripng-import;
        }
```


Copyright © 2017, Juniper Networks, Inc.

```
    }
  }
  user@R1# show policy-options
  policy-statement advertise-routes-through-ripng {
    term 1 {
      from protocol [ direct ripng ];
      then accept;
    }
  }
  policy-statement ripng-import {
    term 1 {
      from {
        route-filter fe80::/64 orlonger;
        route-filter 2001:db8::/64 orlonger;
      }
      then accept;
    }
    term 2 {
      then reject;
    }
  }
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Looking at the Neighbor Addresses for Device R2 on page 127
- Looking at the Routes That Device R2 Is Advertising to Device R1 on page 128
- Looking at the Routes That Device R1 Is Receiving from Device R2 on page 128
- Checking the Routing Table on page 128

Looking at the Neighbor Addresses for Device R2

Purpose Determine the neighbor address that Device R2 is using for Device R1.

Action From operational mode, enter the **show ripng neighbor** command.

```
user@R2> show ripng neighbor fe-1/2/0.2

                    Source                          Dest              In
Neighbor     State  Address                         Address  Send Recv Met
--------     -----  -------                         -------  ---- ---- ---
fe-1/2/0.2      Up fe80::2a0:a514:0:24c             ff02::9  yes  yes  1
```

Meaning Device R2 is using the fe80::2a0:a514:0:24c address to send routes to Device R1.

Copyright © 2017, Juniper Networks, Inc.

Looking at the Routes That Device R2 Is Advertising to Device R1

Purpose Verify that Device R2 is sending the expected routes.

Action From operational mode, enter the **show route advertising-protocol ripng** command.

```
user@R2> show route advertising-protocol ripng fe80::2a0:a514:0:24c
inet6.0: 17 destinations, 18 routes (17 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::2/128    *[Direct/0] 3d 22:00:34
                    >    via lo0.2
2001:db8::3/128    *[RIPng/100] 3d 21:47:00, metric 2, tag 0
                    > to fe80::2a0:a514:0:64c via fe-1/2/1.5
2001:db8:0:2::/64  *[Direct/0] 3d 22:00:34
                    >    via fe-1/2/0.2
2001:db8:0:3::/64  *[Direct/0] 3d 22:00:34
                    >    via fe-1/2/1.5
2001:db8:0:4::/64  *[RIPng/100] 3d 21:47:00, metric 2, tag 0
                    > to fe80::2a0:a514:0:64c via fe-1/2/1.5
2002:db7::2/128    *[Direct/0] 00:29:05
                    >    via lo0.2
2002:db8::2/128    *[Direct/0] 00:31:49
                    >    via lo0.2
2002:db9::2/128    *[Direct/0] 00:29:05
                    >    via lo0.2
```

Meaning Device R2 is sending the extra loopback interface /128 routes to Device R1.

Looking at the Routes That Device R1 Is Receiving from Device R2

Purpose Verify that Device R1 is receiving the expected routes.

Action From operational mode, enter the **show route receive-protocol ripng** command.

```
user@R1> show route receive-protocol ripng fe80::2a0:a514:0:24c

inet6.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::2/128    *[RIPng/100] 3d 21:55:49, metric 2, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
2001:db8::3/128    *[RIPng/100] 3d 21:55:46, metric 3, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
```

Meaning The output shows that the extra loopback interface addresses are excluded.

Checking the Routing Table

Purpose Verify that the routing table is populated with the expected routes.

Copyright © 2017, Juniper Networks, Inc.

Action From operational mode, enter the **show route protocol ripng** command.

```
user@R1> show route protocol ripng
inet6.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::2/128    *[RIPng/100] 3d 22:01:40, metric 2, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
2001:db8::3/128    *[RIPng/100] 3d 22:01:37, metric 3, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
ff02::9/128        *[RIPng/100] 00:00:08, metric 1
                          MultiRecv
```

Meaning The output shows that the routes have been learned from Device R2 and Device R3.

If you delete or deactivate the import policy, the routing table contains the extra loopback interface routes.

Related Documentation

- Understanding RIPng Import Policies to Filter Routes on page 123

Example: Testing a Routing Policy with Complex Regular Expressions

This example shows how to test a routing policy using the **test policy** command to ensure that the policy produces the results that you expect before you apply it in a production environment. Regular expressions, especially complex ones, can be tricky to get right. This example shows how to use the **test policy** command to make sure that your regular expressions have the intended effect.

- Requirements on page 129
- Overview on page 129
- Configuration on page 131
- Verification on page 135

Requirements

No special configuration beyond device initialization is required before you configure this example.

Overview

This example shows two routing devices with an external BGP (EBGP) connection between them. Device R2 uses the BGP session to send customer routes to Device R1. These static routes have multiple community values attached.

```
user@R2> show route match-prefix 172.16.* detail

inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
172.16.1.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Reject
```

Copyright © 2017, Juniper Networks, Inc.

```
                Address: 0x8fd0dc4
                Next-hop reference count: 8
                State: <Active Int Ext>
                Local AS: 64511
                Age: 21:32:13
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
                Communities: 64510:1 64510:10 64510:11 64510:100 64510:111

172.16.2.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Reject
                Address: 0x8fd0dc4
                Next-hop reference count: 8
                State: <Active Int Ext>
                Local AS: 64511
                Age: 21:32:13
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
                Communities: 64510:2 64510:20 64510:22 64510:200 64510:222

172.16.3.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Reject
                Address: 0x8fd0dc4
                Next-hop reference count: 8
                State: <Active Int Ext>
                Local AS: 64511
                Age: 21:32:13
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
                Communities: 64510:3 64510:30 64510:33 64510:300 64510:333

172.16.4.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Reject
                Address: 0x8fd0dc4
                Next-hop reference count: 8
                State: <Active Int Ext>
                Local AS: 64511
                Age: 21:32:13
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
                Communities: 64510:4 64510:40 64510:44 64510:400 64510:444
```

To test a complex regular expression, Device R2 has a policy called **test-regex** that locates routes. The policy is configured like this:

```
policy-statement test-regex {
  term find-routes {
    from community complex-regex;
    then accept;
```


Copyright © 2017, Juniper Networks, Inc.

```
    }
    term reject-the-rest {
      then reject;
    }
  }
  community complex-regex members "^64510:[13].*$";
```

This regular expression matches community values beginning with either 1 or 3.

Topology

Figure 20 on page 131 shows the sample network.

Figure 20: Routing Policy Test for Complex Regular Expressions

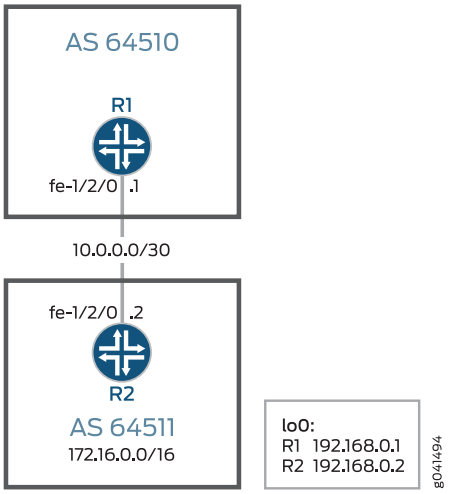

"CLI Quick Configuration" on page 131 shows the configuration for all of the devices in Figure 20 on page 131.

The section "Step-by-Step Procedure" on page 132 describes the steps on Device R2.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.1/30
set interfaces lo0 unit 0 family inet address 192.168.0.1/32
set protocols bgp group ext type external
set protocols bgp group ext peer-as 64511
set protocols bgp group ext neighbor 10.0.0.2
set routing-options router-id 192.168.0.1
set routing-options autonomous-system 64510
```

Device R2

```
set interfaces fe-1/2/0 unit 0 family inet address 10.0.0.2/30
set interfaces lo0 unit 0 family inet address 192.168.0.2/32
set protocols bgp group ext type external
```

Copyright © 2017, Juniper Networks, Inc.

```
set protocols bgp group ext peer-as 64510
set protocols bgp group ext neighbor 10.0.0.1
set policy-options policy-statement send-static term 1 from protocol static
set policy-options policy-statement send-static term 1 then accept
set policy-options policy-statement send-static term 2 then reject
set policy-options policy-statement test-regex term find-routes from community
  complex-regex
set policy-options policy-statement test-regex term find-routes then accept
set policy-options policy-statement test-regex term reject-the-rest then reject
set policy-options community complex-regex members "^64510:[13].*$"
set routing-options static route 172.16.1.0/24 reject
set routing-options static route 172.16.1.0/24 community 64510:1
set routing-options static route 172.16.1.0/24 community 64510:10
set routing-options static route 172.16.1.0/24 community 64510:11
set routing-options static route 172.16.1.0/24 community 64510:100
set routing-options static route 172.16.1.0/24 community 64510:111
set routing-options static route 172.16.2.0/24 reject
set routing-options static route 172.16.2.0/24 community 64510:2
set routing-options static route 172.16.2.0/24 community 64510:20
set routing-options static route 172.16.2.0/24 community 64510:22
set routing-options static route 172.16.2.0/24 community 64510:200
set routing-options static route 172.16.2.0/24 community 64510:222
set routing-options static route 172.16.3.0/24 reject
set routing-options static route 172.16.3.0/24 community 64510:3
set routing-options static route 172.16.3.0/24 community 64510:30
set routing-options static route 172.16.3.0/24 community 64510:33
set routing-options static route 172.16.3.0/24 community 64510:300
set routing-options static route 172.16.3.0/24 community 64510:333
set routing-options static route 172.16.4.0/24 reject
set routing-options static route 172.16.4.0/24 community 64510:4
set routing-options static route 172.16.4.0/24 community 64510:40
set routing-options static route 172.16.4.0/24 community 64510:44
set routing-options static route 172.16.4.0/24 community 64510:400
set routing-options static route 172.16.4.0/24 community 64510:444
set routing-options router-id 192.168.0.2
set routing-options autonomous-system 64511
```

Step-by-Step Procedure

The following example requires that you navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure Device R2:

1. Configure the interfaces.

   ```
   [edit interfaces]
   user@R2# set fe-1/2/0 unit 0 family inet address 10.0.0.2/30

   user@R2# set lo0 unit 0 family inet address 192.168.0.2/32
   ```

2. Configure BGP.

   Apply the import policy to the BGP peering session with Device R2.

   ```
   [edit protocols bgp group ext]
   ```


Copyright © 2017, Juniper Networks, Inc.

```
user@R2# set type external
user@R2# set peer-as 64510
user@R2# set neighbor 10.0.0.1
```

3. Configure the routing policy that sends static routes.

```
[edit policy-options policy-statement send-static]
user@R2# set term 1 from protocol static
user@R2# set term 1 then accept
user@R2# set term 2 then reject
```

4. Configure the routing policy that tests a regular expression.

```
[edit policy-options policy-statement test-regex]
user@R2# set term find-routes from community complex-regex
user@R2# set term find-routes then accept
user@R2# set term reject-the-rest then reject

[edit policy-options community]
user@R2# set complex-regex members "^64510:[13].*$"
```

5. Configure the static routes and attaches community values.

```
[edit routing-options static route 172.16.1.0/24]
user@R2# set reject
user@R2# set community [ 64510:1 64510:10 64510:11 64510:100 64510:111 ]

[edit routing-options static route 172.16.2.0/24]
user@R2# set reject
user@R2# set community [ 64510:2 64510:20 64510:22 64510:200 64510:222 ]

[edit routing-options static route 172.16.3.0/24]
user@R2# set reject
user@R2# set community [ 64510:3 64510:30 64510:33 64510:300 64510:333 ]

[edit routing-options static route 172.16.4.0/24]
user@R2# set reject
user@R2# set community [ 64510:4 64510:40 64510:44 64510:400 64510:444 ]
```

6. Configure the autonomous system (AS) number and the router ID.

   This affects Device R2's routing table, and as no impact on Device R1 and Device R3.

```
[edit routing-options ]
user@R2# set router-id 192.168.0.2
user@R2# set autonomous-system 64511
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, **show policy-options**, and **show routing-options** commands. If the output

Copyright © 2017, Juniper Networks, Inc.

does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@R2# show interfaces
fe-1/2/0 {
    unit 0 {
        family inet {
            address 10.0.0.2/30;
        }
    }
}
lo0 {
    unit 0 {
        family inet {
            address 192.168.0.2/32;
        }
    }
}
```

```
user@R2# show protocols
bgp {
    group ext {
        type external;
        peer-as 64510;
        neighbor 10.0.0.1;
    }
}
```

```
user@R2# show policy-options
policy-statement send-static {
    term 1 {
        from protocol static;
        then accept;
    }
    term 2 {
        then reject;
    }
}
policy-statement test-regex {
    term find-routes {
        from community complex-regex;
        then accept;
    }
    term reject-the-rest {
        then reject;
    }
}
community complex-regex members "^64510:[13].*$";
```

```
user@R2# show routing-options
static {
    route 172.16.1.0/24 {
        reject;
        community [ 64510:1 64510:10 64510:11 64510:100 64510:111 ];
    }
    route 172.16.2.0/24 {
        reject;
```


Copyright © 2017, Juniper Networks, Inc.

```
    community [ 64510:2 64510:20 64510:22 64510:200 64510:222 ];
  }
  route 172.16.3.0/24 {
    reject;
    community [ 64510:3 64510:30 64510:33 64510:300 64510:333 ];
  }
  route 172.16.4.0/24 {
    reject;
    community [ 64510:4 64510:40 64510:44 64510:400 64510:444 ];
  }
}
router-id 192.168.0.2;
autonomous-system 64511;
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Test to See Which Communities Match the Regular Expression

Purpose You can test the regular expression and its policy by using the test policy ***policy-name*** command.

Action

1. On Device R2, run the **test policy test-regex 0/0** command.

```
user@R2> test policy test-regex 0/0


inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.1.0/24      *[Static/5] 1d 00:32:50
                      Reject
172.16.3.0/24      *[Static/5] 1d 00:32:50
                      Reject

Policy test-regex: 2 prefix accepted, 5 prefix rejected
```

2. On Device R2, change the regular expression to match a community value containing any number of instances of the digit 2.

```
[edit policy-options community complex-regex]
user@R2# delete members "^64510:[13].*$"
user@R2# set members "^65020:2+$"
user@R2# commit
```

3. On Device R2, rerun the **test policy test-regex 0/0** command.

```
user@R2> test policy test-regex 0/0

inet.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.2.0/24      *[Static/5] 1d 00:31:36
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Reject

Policy test-regex: 1 prefix accepted, 6 prefix rejected
```

Meaning The 172.16.1.0 /24 and 172.16.3.0/24 routes both have communities attached that match the ^64510:[13].*$ expression. The 172.16.2.0/24 route has communities that match the ^65020:2+$ expression.

Related Documentation

- *Understanding Routing Policy Tests*
- *Understanding How to Define BGP Communities and Extended Communities*
- *Understanding AS Path Regular Expressions for Use as Routing Policy Match Conditions*

 Copyright © 2017, Juniper Networks, Inc.

CHAPTER 15

Configuring Traffic Control with Metrics in a RIPng Network

- Understanding RIPng Traffic Control with Metrics for Optimizing the Path Cost on page 137
- Example: Configuring the Metric Value Added to Imported RIPng Routes to Control the Route Selection Process on page 138

Understanding RIPng Traffic Control with Metrics for Optimizing the Path Cost

To tune a RIPng network and to control traffic flowing through the network, you increase or decrease the cost of the paths through the network. RIPng provides two ways to modify the path cost: an incoming metric and an outgoing metric, which are each set to 1 by default. In other words, by default, the metric of routes that RIPng imports from a neighbor or exports to a neighbor is incremented by 1. These routes include those learned from RIPng as well as those learned from other protocols. The metrics are attributes that 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 of 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.

If an exported route was learned from a member of the same RIPng group, the metric associated with that route is the normal RIPng metric. For example, a RIPng 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 RIPng 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.

Copyright © 2017, Juniper Networks, Inc.

You might want to increase the metric of routes to decrease the likelihood that a particular route is selected and installed in the routing table. This process is sometimes referred to as *route poisoning*. Some reasons that you might want to poison a route are that the route is relatively expensive to use, or it has relatively low bandwidth.

A route with a higher metric than another route becomes the active route only when the lower-metric route becomes unavailable. In this way, the higher-metric route serves as a backup path.

One way to increase the metric of imported routes is to configure an import policy. Another way is to include the **metric-in** statement in the RIPng neighbor configuration. One way to increase the metric of export routes is to configure an export policy. Another way is to include the **metric-out** statement in the RIPng neighbor configuration.

Related Documentation

- Example: Configuring the Metric Value Added to Imported RIPng Routes to Control the Route Selection Process on page 138

Example: Configuring the Metric Value Added to Imported RIPng Routes to Control the Route Selection Process

This example shows how to change the default metric to be added to incoming routes to control the route selection process.

- Requirements on page 138
- Overview on page 138
- Configuration on page 139
- Verification on page 142

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

Normally, when multiple routes are available, RIPng selects the route with the lowest hop count. Changing the default metric enables you to control the route selection process such that a route with a higher hop count can be preferred over of a route with a lower hop count.

Figure 21 on page 139 shows the topology used in this example.

Copyright © 2017, Juniper Networks, Inc.

Figure 21: RIPng Incoming Metrics Network Topology

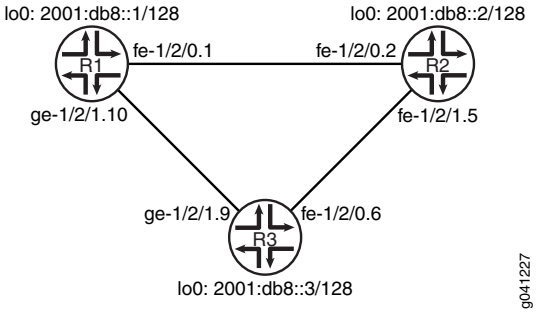

Device R1 has two potential paths to reach 2001:db8::2/128. The default behavior is to send traffic out the 2001:db8:0:1::/64 interface facing Device R2. Suppose, though, that the path through Device R3 is less expensive to use or has higher bandwidth links. This example shows how to use the **metric-in** statement to ensure that Device R1 uses the path through Device R3 to reach 2001:db8::2/128. "CLI Quick Configuration" on page 139 shows the configuration for all of the devices in Figure 21 on page 139. The section "Step-by-Step Procedure" on page 140 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 description to-R2
set interfaces fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces ge-1/2/1 unit 10 description to-R3
set interfaces ge-1/2/1 unit 10 family inet6 address 2001:db8:0:5::/64 eui-64
set interfaces lo0 unit 1 family inet6 address 2001:db8::1/128
set protocols ripng group primary export advertise-routes-through-ripng
set protocols ripng group primary neighbor ge-1/2/1.10
set protocols ripng group secondary export advertise-routes-through-ripng
set protocols ripng group secondary neighbor fe-1/2/0.1 metric-in 4
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
    ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 family inet6 address 2001:db8:0:2::/64 eui-64
set interfaces fe-1/2/1 unit 5 description to-R3
set interfaces fe-1/2/1 unit 5 family inet6 address 2001:db8:0:3::/64 eui-64
set interfaces lo0 unit 2 family inet6 address 2001:db8::2/128
set protocols ripng group ripng-group export advertise-routes-through-ripng
set protocols ripng group ripng-group neighbor fe-1/2/0.2
set protocols ripng group ripng-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
    direct
```

Copyright © 2017, Juniper Networks, Inc.

```
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 family inet6 address 2001:db8:0:4::/64 eui-64
set interfaces ge-1/2/1 unit 9 description to-R1
set interfaces ge-1/2/1 unit 9 family inet address 10.0.0.9/30
set interfaces ge-1/2/1 unit 9 family inet6 address 2001:db8:0:6::/64 eui-64
set interfaces lo0 unit 3 family inet6 address 2001:db8::3/128
set protocols ripng group ripng-group export advertise-routes-through-ripng
set protocols ripng group ripng-group neighbor fe-1/2/0.6
set protocols ripng group ripng-group neighbor ge-1/2/1.9
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure a RIPng metrics:

1. Configure the network interfaces.

```
[edit interfaces]
user@R1# set fe-1/2/0 unit 1 description to-R2
user@R1# set fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64

user@R1# set ge-1/2/1 unit 10 description to-R3
user@R1# set ge-1/2/1 unit 10 family inet6 address 2001:db8:0:5::/64 eui-64

user@R1# set lo0 unit 1 family inet6 address 2001:db8::1/128
```

2. Create the RIPng groups and add the interfaces.

   To configure RIPng in Junos OS, you must configure one or more groups that contain the interfaces on which RIPng is enabled. You do not need to enable RIPng on the loopback interface.

   For the interface that is facing Device R2, the **metric-in 4** setting causes this route to be less likely to be chosen as the active route.

```
[edit protocols ripng]
user@R1# set group primary neighbor ge-1/2/1.10
user@R1# set group secondary neighbor fe-1/2/0.1 metric-in 4
```

3. Create the routing policy to advertise both direct and RIPng-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-ripng term 1]
user@R1# set from protocol direct
user@R1# set from protocol ripng
user@R1# set then accept
```


Copyright © 2017, Juniper Networks, Inc.

4. Apply the routing policy.

   In Junos OS, you can only apply RIPng export policies at the group level.

   ```
   [edit protocols ripng]
   user@R1# set group primary export advertise-routes-through-ripng
   user@R1# set group secondary export advertise-routes-through-ripng
   ```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
    unit 1 {
        description to-R2;
        family inet6 {
            address 2001:db8:0:1::/64 {
                eui-64;
            }
        }
    }
}
ge-1/2/1 {
    unit 10 {
        description to-R3;
        family inet6 {
            address 2001:db8:0:5::/64 {
                eui-64;
            }
        }
    }
}
lo0 {
    unit 1 {
        family inet6 {
            address 2001:db8::1/128;
        }
    }
}
```

```
user@R1# show protocols
ripng {
    group primary {
        export advertise-routes-through-ripng;
        neighbor ge-1/2/0.10;
    }
    group secondary {
        export advertise-routes-through-ripng;
        neighbor fe-1/2/0.1 {
            metric-in 4;
        }
    }
}
```

Copyright © 2017, Juniper Networks, Inc.

```
user@R1# show policy-options
policy-statement advertise-routes-through-ripng {
  term 1 {
    from protocol [ direct ripng ];
    then accept;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- Verifying That the Expected Route Is Active on page 142
- Removing the metric-in Statement on page 142

Verifying That the Expected Route Is Active

Purpose Make sure that Device R1 uses the path through Device R3 to reach 2001:db8:0:2:2a0:a514:0:24c/128.

Action From operational mode, enter the **show route 2001:db8:0:2:2a0:a514:0:24c** command.

```
user@R1> show route 2001:db8:0:2:2a0:a514:0:24c
inet6.0: 16 destinations, 17 routes (16 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8:0:2::/64  *[RIPng/100] 01:54:35, metric 3, tag 0
                    > to fe80::2a0:a514:0:94c via ge-1/2/1.10
```

Meaning The **to fe80::2a0:a514:0:94c via ge-1/2/1.10** output shows that Device R1 uses the path through Device R3 to reach 2001:db8:0:2:2a0:a514:0:24c/128. The metric for this route is 3.

Removing the metric-in Statement

Purpose Delete or deactivate the metric-in statement to see what happens to the 2001:db8:0:2:2a0:a514:0:24c/128 route.

Action

1. From configuration mode, deactivate the **metric-in** statement.

   ```
   [edit protocols ripng group secondary neighbor fe-1/2/0.1]
   user@R1# deactivate metric-in
   user@R1# commit
   ```

2. From operational mode, enter the **show route 2001:db8:0:2:2a0:a514:0:24c** command.

   ```
   user@R1> show route 2001:db8:0:2:2a0:a514:0:24c
   inet6.0: 16 destinations, 17 routes (16 active, 0 holddown, 0 hidden)
   + = Active Route, - = Last Active, * = Both
   ```


Copyright © 2017, Juniper Networks, Inc.

```
2001:db8:0:2::/64  *[RIPng/100] 00:00:02, metric 2, tag 0
                    > to fe80::2a0:a514:0:24c via fe-1/2/0.1
```

Meaning The **to fe80::2a0:a514:0:24c via fe-1/2/0.1** output shows that Device R1 uses the path through Device R2 to reach 2001:db8:0:2:2a0:a514:0:24c/128. The metric for this route is 2.

Related Documentation

- Understanding RIPng Traffic Control with Metrics for Optimizing the Path Cost on page 137

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 16

Configuring RIPng Timers

- Understanding RIP Timers on page 145
- Example: Configuring RIPng Update Interval on page 146

Understanding RIP Timers

RIP uses several timers to regulate its operation.

The update interval is the interval at which routes that are learned by RIP are advertised to neighbors. This timer controls the interval between routing updates. The update interval is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can occur if all routing devices update their neighbors simultaneously.

To configure the update time interval, include the **update-interval** statement:

```
update-interval seconds;
```

seconds can be a value from 10 through 60.

You can set a route timeout interval. If a route is not refreshed after being installed in the routing table by the specified time interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.

To configure the route timeout for RIP, include the **route-timeout** statement:

```
route-timeout seconds;
```

seconds can be a value from 30 through 360. The default value is 180 seconds.

RIP routes expire when either a route timeout limit is met or a route metric reaches infinity, and the route is no longer valid. However, the expired route is retained in the routing table for a specified period so that neighbors can be notified that the route has been dropped. This time period is set by configuring the hold-down timer. Upon expiration of the hold-down timer, the route is removed from the routing table.

To configure the hold-down timer for RIP, include the **holddown** statement:

```
holddown seconds;
```

seconds can be a value from 10 through 180. The default value is 120 seconds.

Copyright © 2017, Juniper Networks, Inc.

NOTE: In Junos OS Release 11.1 and later, a retransmission timer is available for RIP demand circuits.

Generally, we recommend against changing the RIP timers, unless the effects of a change are well understood. The route timeout should be at least three times the update interval. Normally, the default values are best left in effect for standard operations.

Release History Table

Release	Description
11.1	In Junos OS Release 11.1 and later, a retransmission timer is available for RIP demand circuits.

Related Documentation

- Example: Configuring RIP Timers on page 38
- *Example: Configuring RIP Demand Circuits*

Example: Configuring RIPng Update Interval

This example shows how to configure the RIPng update interval and how to monitor the impact of the change.

- Requirements on page 146
- Overview on page 146
- Configuration on page 147
- Verification on page 150

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, Device R2 has an update interval of 60 seconds for its neighbor Device R1, and an update interval of 10 seconds for its neighbor Device R3.

This example is not necessarily practical, but it is shown for demonstration purposes. Generally, we recommend against changing the RIPng timers, unless the effects of a change are well understood. Normally, the default values are best left in effect for standard operations.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIPng.

Figure 22 on page 147 shows the topology used in this example.

Copyright © 2017, Juniper Networks, Inc.

Figure 22: RIPng Timers Network Topology

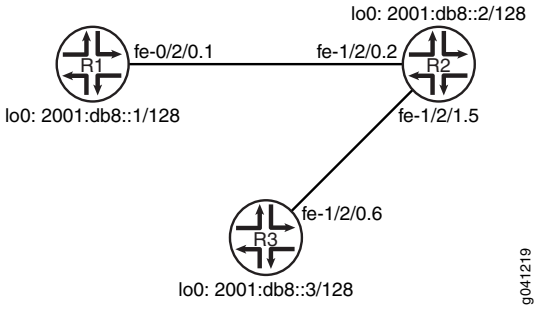

"CLI Quick Configuration" on page 147 shows the configuration for all of the devices in Figure 22 on page 147. The section "Step-by-Step Procedure" on page 148 describes the steps on Device R2.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 description to-R2
set interfaces fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces lo0 unit 1 family inet6 address 2001:db8::1/128
set protocols ripng group ripng-group export advertise-routes-through-ripng
set protocols ripng group ripng-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Device R2

```
set interfaces fe-1/2/0 unit 2 description to-R1
set interfaces fe-1/2/0 unit 2 family inet6 address 2001:db8:0:2::/64 eui-64
set interfaces fe-1/2/1 unit 5 description to-R3
set interfaces fe-1/2/1 unit 5 family inet6 address 2001:db8:0:3::/64 eui-64
set interfaces lo0 unit 2 family inet6 address 2001:db8::2/128
set protocols ripng group ripng-group export advertise-routes-through-ripng
set protocols ripng group ripng-group neighbor fe-1/2/0.2 update-interval 60
set protocols ripng group ripng-group neighbor fe-1/2/1.5 update-interval 10
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 description to-R2
set interfaces fe-1/2/0 unit 6 family inet6 address 2001:db8:0:4::/64 eui-64
set interfaces lo0 unit 3 family inet6 address 2001:db8::3/128
set protocols ripng group ripng-group export advertise-routes-through-ripng
```

Copyright © 2017, Juniper Networks, Inc.

```
set protocols ripng group ripng-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
    ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the RIPng update interval:

1. Configure the network interfaces.

   This example shows multiple loopback interface addresses to simulate attached networks.

   ```
   [edit interfaces]
   user@R2# set fe-1/2/0 unit 2 description to-R1
   user@R2# set fe-1/2/0 unit 2 family inet6 address 2001:db8:0:2::/64 eui-64

   user@R2# set fe-1/2/1 unit 5 description to-R3
   user@R2# set fe-1/2/1 unit 5 family inet6 address 2001:db8:0:3::/64 eui-64

   user@R2# set lo0 unit 2 family inet6 address 2001:db8::2/128
   ```

2. Configure different update intervals for the two RIPng neighbors.

   To configure RIPng in Junos OS, you must configure a group that contains the interfaces on which RIPng is enabled. You do not need to enable RIPng on the loopback interface.

   ```
   [edit protocols ripng group ripng-group]
   user@R2# set neighbor fe-1/2/0.2 update-interval 60
   user@R2# set neighbor fe-1/2/1.5 update-interval 10
   ```

3. Create the routing policy to advertise both direct and RIPng-learned routes.

   ```
   [edit policy-options policy-statement advertise-routes-through-ripng term 1]
   user@R2# set from protocol direct
   user@R2# set from protocol ripng
   user@R2# set then accept
   ```

4. Apply the routing policy.

   In Junos OS, you can only apply RIPng export policies at the group level.

   ```
   [edit protocols ripng group ripng-group]
   user@R2# set export advertise-routes-through-ripng
   ```


Copyright © 2017, Juniper Networks, Inc.

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R2# show interfaces
fe-1/2/0 {
    unit 2 {
        description to-R1;
        family inet6 {
            address 2001:db8:0:2::/64 {
                eui-64;
            }
        }
    }
}
fe-1/2/1 {
    unit 5 {
        description to-R3;
        family inet6 {
            address 2001:db8:0:3::/64 {
                eui-64;
            }
        }
    }
}
lo0 {
    unit 2 {
        family inet6 {
            address 2001:db8::2/128;
        }
    }
}
```

```
user@R2# show protocols
ripng {
    group ripng-group {
        export advertise-routes-through-ripng;
        neighbor fe-1/2/0.2 {
            update-interval 60;
        }
        neighbor fe-1/2/1.5 {
            update-interval 10;
        }
    }
}
```

```
user@R2# show policy-options
policy-statement advertise-routes-through-ripng {
    term 1 {
        from protocol [ direct ripng ];
        then accept;
    }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Copyright © 2017, Juniper Networks, Inc.

Verification

Confirm that the configuration is working properly.

- Checking the RIPng Updates Sent by Device R2 on page 150
- Checking the RIPng Updates Received by Device R2 on page 151
- Checking the RIPng Updates Received by Device R3 on page 151

Checking the RIPng Updates Sent by Device R2

Purpose Make sure that the RIPng update packets are sent at the expected interval.

Action From operational mode, enter the **show ripng statistics** command.

```
user@R2> show ripng statistics
RIPng info: port 521; holddown 120s.
    rts learned  rts held down  rqsts dropped  resps dropped
              4              0              0              0

fe-1/2/0.2:  2 routes learned; 5 routes advertised; timeout 180s; update interval
 60s
Counter                       Total   Last 5 min  Last minute
-------                   -----------  -----------  -----------
Updates Sent                        1            1            1
Triggered Updates Sent              0            0            0
Responses Sent                      0            0            0
Bad Messages                        0            0            0
Updates Received                    1            0            0
Bad Route Entries                   0            0            0
Updates Ignored                     0            0            0
RIPng Requests Received             0            0            0
RIPng Requests Ignored              0            0            0

fe-1/2/1.5:  2 routes learned; 5 routes advertised; timeout 180s; update interval
 10s
Counter                       Total   Last 5 min  Last minute
-------                   -----------  -----------  -----------
Updates Sent                        6            2            2
Triggered Updates Sent              0            0            0
Responses Sent                      0            0            0
Bad Messages                        0            0            0
Updates Received                    2            0            0
Bad Route Entries                   0            0            0
Updates Ignored                     0            0            0
RIPng Requests Received             0            0            0
RIPng Requests Ignored              0            0            0
```

Meaning The **update interval** field shows that the interval is 60 seconds for its neighbor Device R1 and 10 seconds for its neighbor Device R3. The **Updates Sent** field shows that Device R2 is sending updates to Device R1 at roughly 1/6 of the rate that it is sending updates to Device R3.

Copyright © 2017, Juniper Networks, Inc.

Checking the RIPng Updates Received by Device R2

Purpose Make sure that the RIPng update packets are sent at the expected interval.

Action From operational mode, enter the **show ripng statistics** command.

```
user@R1> show ripng statistics
RIPng info: port 521; holddown 120s.
    rts learned  rts held down  rqsts dropped  resps dropped
              5              8              0              0

fe-1/2/0.1:  5 routes learned; 2 routes advertised; timeout 180s; update interval
 30s
Counter                         Total   Last 5 min  Last minute
-------                     ----------  ----------  ----------
Updates Sent                         6           5           2
Triggered Updates Sent               0           0           0
Responses Sent                       0           0           0
Bad Messages                         0           0           0
Updates Received                     3           3           1
Bad Route Entries                    0           0           0
Updates Ignored                      0           0           0
RIPng Requests Received              0           0           0
RIPng Requests Ignored               0           0           0
```

Meaning The **Updates Received** field shows the number of updates received from Device R2.

Checking the RIPng Updates Received by Device R3

Purpose Make sure that the RIPng update packets are sent at the expected interval.

Action From operational mode, enter the **show ripng statistics** command.

```
user@R3> show ripng statistics
RIPng info: port 521; holddown 120s.
    rts learned  rts held down  rqsts dropped  resps dropped
              5              0              0              0


fe-1/2/0.6:  5 routes learned; 2 routes advertised; timeout 180s; update interval
 30s
Counter                         Total   Last 5 min  Last minute
-------                     ----------  ----------  ----------
Updates Sent                         5           5           2
Triggered Updates Sent               0           0           0
Responses Sent                       0           0           0
Bad Messages                         0           0           0
Updates Received                    16          15           6
Bad Route Entries                    0           0           0
Updates Ignored                      0           0           0
RIPng Requests Received              0           0           0
RIPng Requests Ignored               0           0           0
```

Copyright © 2017, Juniper Networks, Inc.

Meaning The **Updates Received** field shows the number of updates received from Device R2.

Related Documentation

- Understanding RIP Timers on page 37

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 17

Tracing RIPng Protocol Traffic

- Understanding RIPng Protocol Traffic Trace Operations on page 153
- Example: Tracing Global Routing Protocol Operations on page 154
- Example: Tracing RIPng Protocol Traffic on page 158

Understanding RIPng Protocol Traffic Trace Operations

You can trace various RIPng protocol traffic to help debug RIP protocol issues.

To trace RIP protocol traffic, include the **traceoptions** statement at the **[edit protocols ripng]** hierarchy level:

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

You can specify the following RIPng protocol-specific trace options using the **flag** statement:

- **error**—RIPng error packets
- **expiration**—RIPng route expiration processing
- **holddown**—RIPng hold-down processing
- **nsr-synchronization**—Nonstop routing synchronization events
- **packets**—All RIPng packets
- **request**—RIPng information packets
- **trigger**—RIPng triggered updates
- **update**—RIPng update packets

You can optionally specify one or more of the following flag modifiers:

- **detail**—Detailed trace information
- **receive**—Packets being received
- **send**—Packets being transmitted

Copyright © 2017, Juniper Networks, Inc.

NOTE: Use the **detail** flag modifier with caution as this might cause the CPU to become very busy.

Global tracing options are inherited from the configuration set by the **traceoptions** statement at the **[edit routing-options]** hierarchy level. You can override the following global trace options for the RIPng protocol using the **traceoptions flag** statement included at the **[edit protocols ripng]** hierarchy level:

- **all**—All tracing operations
- **general**—All normal operations and routing table changes (a combination of the normal and route trace operations)
- **normal**—Normal events
- **policy**—Policy processing
- **route**—Routing information
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

NOTE: Use the **trace** flag **all** with caution as this might cause the CPU to become very busy.

Related Documentation

- Example: Tracing RIPng Protocol Traffic on page 158
- Example: Tracing Global Routing Protocol Operations on page 154.

Example: Tracing Global Routing Protocol Operations

This example shows how to list and view files that are created when you enable global routing trace operations.

- Requirements on page 154
- Overview on page 155
- Configuration on page 155
- Verification on page 158

Requirements

You must have the **view** privilege.

Copyright © 2017, Juniper Networks, Inc.

Overview

To configure global routing protocol tracing, include the **traceoptions** statement at the **[edit routing-options]** hierarchy level:

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

The flags in a **traceoptions flag** statement are identifiers. When you use the **set** command to configure a flag, any flags that might already be set are not modified. In the following example, setting the **timer** tracing flag has no effect on the already configured **task** flag. Use the **delete** command to delete a particular flag.

```
[edit routing-options traceoptions]
user@host# show
flag task;
user@host# set traceoptions flag timer
user@host# show
flag task;
flag timer;
user@host# delete traceoptions flag task
user@host# show
flag timer;
```

This example shows how to configure and view a trace file that tracks changes in the routing table. The steps can be adapted to apply to trace operations for any Junos OS hierarchy level that supports trace operations.

TIP: To view a list of hierarchy levels that support tracing operations, enter the help apropos traceoptions command in configuration mode.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
set routing-options traceoptions file routing-table-changes
set routing-options traceoptions file size 10m
set routing-options traceoptions file files 10
set routing-options traceoptions flag route
set routing-options static route 1.1.1.2/32 next-hop 10.0.45.6
```

Copyright © 2017, Juniper Networks, Inc.

Configuring Trace Operations

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the trace operations:

1. Configure trace operations.

```
[edit routing-options traceoptions]
user@host# set file routing-table-changes
user@host# set file size 10m
user@host# set file files 10
user@host# set flag route
```

2. Configure a static route to cause a change in the routing table.

```
[edit routing-options static]
user@host# set route 1.1.1.2/32 next-hop 10.0.45.6
```

3. If you are done configuring the device, commit the configuration.

```
[edit]
user@host# commit
```

Viewing the Trace File

Step-by-Step Procedure

To view the trace file:

1. In operational mode, list the log files on the system.

```
user@host> file list /var/log
/var/log:
...
routing-table-changes
...
```

2. View the contents of the **routing-table-changes** file.

```
user@host> file show /var/log/routing-table-changes
Dec 15 11:09:29 trace_on: Tracing to "/var/log/routing-table-changes" started
Dec 15 11:09:29.496507
Dec 15 11:09:29.496507 Tracing flags enabled: route
Dec 15 11:09:29.496507
Dec 15 11:09:29.533203 inet_routerid_notify: Router ID: 192.168.4.1
Dec 15 11:09:29.533334 inet_routerid_notify: No Router ID assigned
Dec 15 11:09:29.533381 inet_routerid_notify: No Router ID assigned
Dec 15 11:09:29.533420 inet_routerid_notify: No Router ID assigned
Dec 15 11:09:29.534915 inet_routerid_notify: Router ID: 192.168.4.1
Dec 15 11:09:29.542934 inet_routerid_notify: No Router ID assigned
Dec 15 11:09:29.549253 inet_routerid_notify: No Router ID assigned
Dec 15 11:09:29.556878 inet_routerid_notify: No Router ID assigned
```


Copyright © 2017, Juniper Networks, Inc.

```
Dec 15 11:09:29.582990 rt_static_reinit: examined 3 static nexthops, 0
unreferenced
Dec 15 11:09:29.589920
Dec 15 11:09:29.589920 task_reconfigure reinitializing done
...
```

3. Filter the output of the log file.

```
user@host> file show /var/log/routing-table-changes | match 1.1.1.2
Dec 15 11:15:30.780314 ADD      1.1.1.2/32          nhid 0 gw 10.0.45.6
   Static   pref 5/0 metric  at-0/2/0.0 <ctive Int Ext>
Dec 15 11:15:30.782276 KRT Request: send len 216 v104 seq 0 ADD route/user
af 2 table 0 infot 0 addr 1.1.1.2 nhop-type unicast nhindex 663
```

4. View the tracing operations in real time by running the **monitor start** command with an optional **match** condition.

```
user@host> monitor start routing-table-changes | match 1.1.1.2
Aug 10 19:21:40.773467 BGP RECV         0.0.0.0/0
Aug 10 19:21:40.773685 bgp_rcv_nlri: 0.0.0.0/0
Aug 10 19:21:40.773778 bgp_rcv_nlri: 0.0.0.0/0 belongs to meshgroup
Aug 10 19:21:40.773832 bgp_rcv_nlri: 0.0.0.0/0 qualified bnp->ribact 0x0
l2afcb 0x0
```

5. Deactivate the static route.

```
user@host# deactivate routing-options static route 1.1.1.2/32
user@host# commit
```

```
*** routing-table-changes ***
Dec 15 11:42:59.355557 CHANGE   1.1.1.2/32          nhid 663 gw 10.0.45.6
     Static   pref 5/0 metric  at-0/2/0.0 <Delete Int Ext>
Dec 15 11:42:59.426887 KRT Request: send len 216 v104 seq 0 DELETE route/user
 af 2 table 0 infot 0 addr 1.1.1.2 nhop-type discard filtidx 0
Dec 15 11:42:59.427366 RELEASE  1.1.1.2/32          nhid 663 gw 10.0.45.6
     Static   pref 5/0 metric  at-0/2/0.0 <Release Delete Int Ext>
```

6. Halt the **monitor** command by pressing Enter and typing **monitor stop**.

```
[Enter]
user@host> monitor stop
```

7. When you are finished troubleshooting, consider deactivating trace logging to avoid any unnecessary impact to system resources.

   When configuration is deactivated, it appears in the configuration with the **inactive** tag.

```
[edit routing-options]
user@host# deactivate traceoptions
user@host# commit
```

```
[edit routing-options]
user@host# show

inactive: traceoptions {
    file routing-table-changes size 10m files 10;
    flag route;
```

Copyright © 2017, Juniper Networks, Inc.

```
}
static {
    inactive: route 1.1.1.2/32 next-hop 10.0.45.6;
}
```

8. To reactivate trace operations, use the **activate** configuration-mode statement.

```
[edit routing-options]
user@host# activate traceoptions
user@host# commit
```

Results

From configuration mode, confirm your configuration by entering the **show routing-options** command. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@host# show routing-options
traceoptions {
   file routing-table-changes size 10m files 10;
   flag route;
}
static {
   route 1.1.1.2/32 next-hop 10.0.45.6;
}
```

Verification

Confirm that the configuration is working properly.

Verifying That the Trace Log File Is Operating

Purpose Make sure that events are being written to the log file.

Action

```
user@host> show log routing-table-changes
Dec 15 11:09:29 trace_on: Tracing to "/var/log/routing-table-changes" started
```

Related Documentation

- *Understanding Global Routing Protocol Tracing Operations*
- CLI Explorer

Example: Tracing RIPng Protocol Traffic

This example shows how to trace RIPng protocol operations.

- Requirements on page 159
- Overview on page 159

Copyright © 2017, Juniper Networks, Inc.

- Configuration on page 159
- Verification on page 162

Requirements

No special configuration beyond device initialization is required before configuring this example.

Overview

In this example, Device R1 is set to trace routing information updates.

An export policy is also shown because an export policy is required as part of the minimum configuration for RIPng.

Figure 23 on page 159 shows the topology used in this example.

Figure 23: RIPng Trace Operations Network Topology

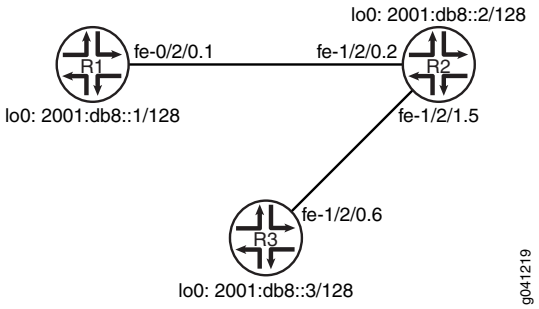

"CLI Quick Configuration" on page 159 shows the configuration for all of the devices in Figure 23 on page 159. The section "Step-by-Step Procedure" on page 160 describes the steps on Device R1.

Configuration

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

Device R1

```
set interfaces fe-1/2/0 unit 1 description to-R2
set interfaces fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64
set interfaces lo0 unit 1 family inet6 address 2001:db8::1/128
set protocols ripng traceoptions file ripng-trace-file
set protocols ripng traceoptions flag route
set protocols ripng group ripng-group export advertise-routes-through-ripng
set protocols ripng group ripng-group neighbor fe-1/2/0.1
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
    direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
    ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Copyright © 2017, Juniper Networks, Inc.

Device R2

```
set interfaces fe-1/2/0 unit 2 description to-R1
set interfaces fe-1/2/0 unit 2 family inet6 address 2001:db8:0:2::/64 eui-64
set interfaces fe-1/2/1 unit 5 description to-R3
set interfaces fe-1/2/1 unit 5 family inet6 address 2001:db8:0:3::/64 eui-64
set interfaces lo0 unit 2 family inet6 address 2001:db8::2/128
set protocols ripng group ripng-group export advertise-routes-through-ripng
set protocols ripng group ripng-group neighbor fe-1/2/0.2
set protocols ripng group ripng-group neighbor fe-1/2/1.5
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Device R3

```
set interfaces fe-1/2/0 unit 6 description to-R2
set interfaces fe-1/2/0 unit 6 family inet6 address 2001:db8:0:4::/64 eui-64
set interfaces lo0 unit 3 family inet6 address 2001:db8::3/128
set protocols ripng group ripng-group export advertise-routes-through-ripng
set protocols ripng group ripng-group neighbor fe-1/2/0.6
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  direct
set policy-options policy-statement advertise-routes-through-ripng term 1 from protocol
  ripng
set policy-options policy-statement advertise-routes-through-ripng term 1 then accept
```

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see *Using the CLI Editor in Configuration Mode* in the *CLI User Guide*.

To configure the RIPng update interval:

1. Configure the network interfaces.

   This example shows multiple loopback interface addresses to simulate attached networks.

   ```
   [edit interfaces]
   user@R1# set fe-1/2/0 unit 1 description to-R2
   user@R1# set fe-1/2/0 unit 1 family inet6 address 2001:db8:0:1::/64 eui-64


   user@R1# set lo0 unit 1 family inet6 address 2001:db8::1/128
   ```

2. Configure the RIPng group, and add the interface to the group.

   To configure RIPng in Junos OS, you must configure a group that contains the interfaces on which RIPng is enabled. You do not need to enable RIPng on the loopback interface.

   ```
   [edit protocols ripng group ripng-group]
   user@R1# set neighbor fe-1/2/0.1
   ```

3. Configure RIPng tracing operations.

   ```
   [edit protocols ripng traceoptions]
   ```


Copyright © 2017, Juniper Networks, Inc.

```
user@R1# set file ripng-trace-file
user@R1# set flag route
```

4. Create the routing policy to advertise both direct and RIPng-learned routes.

```
[edit policy-options policy-statement advertise-routes-through-ripng term 1]
user@R1# set from protocol direct
user@R1# set from protocol ripng
user@R1# set then accept
```

5. Apply the routing policy.

   In Junos OS, you can only apply RIPng export policies at the group level.

```
[edit protocols ripng group ripng-group]
user@R1# set export advertise-routes-through-ripng
```

Results From configuration mode, confirm your configuration by entering the **show interfaces**, **show protocols**, and **show policy-options** commands. If the output does not display the intended configuration, repeat the configuration instructions in this example to correct it.

```
user@R1# show interfaces
fe-1/2/0 {
  unit 1 {
    description to-R2;
    family inet6 {
      address 2001:db8:0:1::/64 {
        eui-64;
      }
    }
  }
}
lo0 {
  unit 1 {
    family inet6 {
      address 2001:db8::1/128;
    }
  }
}
```

```
user@R1# show protocols
ripng {
  traceoptions {
    file ripng-trace-file;
    flag route;
  }
  group ripng-group {
    export advertise-routes-through-ripng;
    neighbor fe-1/2/0.1;
  }
}
```

```
user@R1# show policy-options
```

Copyright © 2017, Juniper Networks, Inc.

```
policy-statement advertise-routes-through-ripng {
  term 1 {
    from protocol [ direct ripng ];
    then accept;
  }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

Checking the Log File

Purpose Make sure that the RIPng route updates are logged in the configured log file.

Action

1. Deactivate the extra loopback interface address on Device R3.

   ```
   [edit interfaces lo0 unit 3 family inet6]
   user@R3# deactivate address 2001:db8::3/128
   user@R3# commit
   ```

2. From operational mode, enter the **show log ripng-trace-file** command with the **| match 2001:db8::3** option.

   ```
   user@R1> show log ripng-trace-file | match 2001:db8::3

   Mar  6 14:57:03.516867  2001:db8::3/128: metric-in: 3, change: 3 -> 3; # gw:
   1, pkt_upd_src fe80::2a0:a514:0:24c, inx: 0, rte_upd_src fe80::2a0:a514:0:24c
   Mar  6 14:57:32.786286  2001:db8::3/128: metric-in: 3, change: 3 -> 3; # gw:
   1, pkt_upd_src fe80::2a0:a514:0:24c, inx: 0, rte_upd_src fe80::2a0:a514:0:24c
   Mar  6 14:58:02.584669  2001:db8::3/128: metric-in: 3, change: 3 -> 3; # gw:
   1, pkt_upd_src fe80::2a0:a514:0:24c, inx: 0, rte_upd_src fe80::2a0:a514:0:24c
   Mar  6 14:58:30.213894  2001:db8::3/128: metric-in: 3, change: 3 -> 3; # gw:
   1, pkt_upd_src fe80::2a0:a514:0:24c, inx: 0, rte_upd_src fe80::2a0:a514:0:24c
   Mar  6 14:59:00.115110  2001:db8::3/128: metric-in: 3, change: 3 -> 3; # gw:
   1, pkt_upd_src fe80::2a0:a514:0:24c, inx: 0, rte_upd_src fe80::2a0:a514:0:24c
   Mar  6 14:59:05.826644 Setting RIPng rtbit on route 2001:db8::3/128, tsi =
   0xbb69880
   Mar  6 14:59:13.014652  2001:db8::3/128: metric-in: 16, change: 3 -> 16; # gw:
    1, pkt_upd_src fe80::2a0:a514:0:24c, inx: 0, rte_upd_src fe80::2a0:a514:0:24c
   Mar  6 14:59:13.015132 CHANGE   2001:db8::3/128     nhid 566 gw
   fe80::2a0:a514:0:24c RIPng    pref 100/0 metric 3/0 fe-1/2/0.1 **Delete Int>
   Mar  6 14:59:13.015197 Best route to 2001:db8::3/128 got deleted. Doing route
   calculation on the stored rte-info
   ```

Meaning The output shows that the route to 2001:db8::3/128 was deleted.

Related Documentation

- Understanding RIPng Protocol Traffic Trace Operations on page 153

Copyright © 2017, Juniper Networks, Inc.

PART 4

Troubleshooting

- Troubleshooting Network Issues on page 165
- Monitoring and Troubleshooting a RIP Network on page 173

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 18

Troubleshooting Network Issues

- Working with Problems on Your Network on page 165
- Isolating a Broken Network Connection on page 166
- Identifying the Symptoms of a Broken Network Connection on page 167
- Isolating the Causes of a Network Problem on page 168
- Taking Appropriate Action for Resolving the Network Problem on page 169
- Evaluating the Solution to Check Whether the Network Problem Is Resolved on page 170

Working with Problems on Your Network

Problem **Description:** This checklist provides links to troubleshooting basics, an example network, and includes a summary of the commands you might use to diagnose problems with the router and network.

Solution

Table 8: Checklist for Working with Problems on Your Network

Tasks	Command or Action
"Isolating a Broken Network Connection" on page 166	
1. Identifying the Symptoms of a Broken Network Connection on page 167	ping (*ip-address* | *hostname*) show route (*ip-address* | *hostname*) traceroute (*ip-address* | *hostname*)
2. Isolating the Causes of a Network Problem on page 168	show < configuration | interfaces | protocols | route >
3. Taking Appropriate Action for Resolving the Network Problem on page 169	[edit] delete routing options static route *destination-prefix* commit and-quit show route *destination-prefix*
4. Evaluating the Solution to Check Whether the Network Problem Is Resolved on page 170	show route (*ip-address* | *hostname*) ping (*ip-address* | *hostname*) count 3 traceroute (*ip-address* | *hostname*)

Copyright © 2017, Juniper Networks, Inc.

Isolating a Broken Network Connection

By applying the standard four-step process illustrated in Figure 24 on page 166, you can isolate a failed node in the network. Note that the functionality described in this section is not supported in versions 15.1X49, 15.1X49-D30, or 15.1X49-D40.

Figure 24: Process for Diagnosing Problems in Your Network

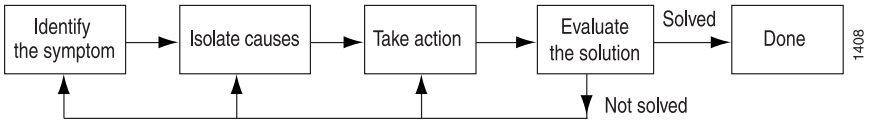

Before you embark on the four-step process, however, it is important that you are prepared for the inevitable problems that occur on all networks. While you might find a solution to a problem by simply trying a variety of actions, you can reach an appropriate solution more quickly if you are systematic in your approach to the maintenance and monitoring of your network. To prepare for problems on your network, understand how the network functions under normal conditions, have records of baseline network activity, and carefully observe the behavior of your network during a problem situation.

Figure 25 on page 166 shows the network topology used in this topic to illustrate the process of diagnosing problems in a network.

Figure 25: Network with a Problem

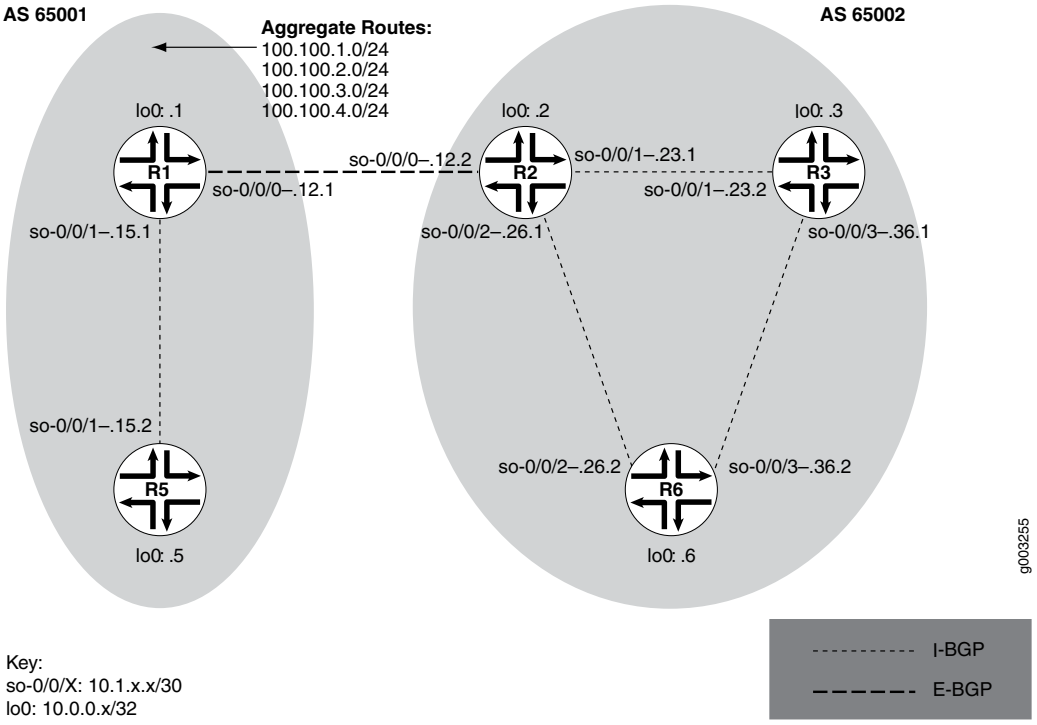

Key:
so-0/0/X: 10.1.x.x/30
lo0: 10.0.0.x/32

The network in Figure 25 on page 166 consists of two autonomous systems (ASs). AS 65001 includes two routers, and AS 65002 includes three routers. The border router (**R1**) in AS 65001 announces aggregated prefixes **100.100/24** to the AS 65002 network. The

Copyright © 2017, Juniper Networks, Inc.

problem in this network is that **R6** does not have access to **R5** because of a loop between **R2** and **R6**.

To isolate a failed connection in your network, follow the steps in these topics:

- Identifying the Symptoms of a Broken Network Connection on page 167
- Isolating the Causes of a Network Problem on page 168
- Taking Appropriate Action for Resolving the Network Problem on page 169
- Taking Appropriate Action for Resolving the Network Problem on page 169
- Evaluating the Solution to Check Whether the Network Problem Is Resolved on page 170

Identifying the Symptoms of a Broken Network Connection

Problem **Description:** The symptoms of a problem in your network are usually quite obvious, such as the failure to reach a remote host.

Solution To identify the symptoms of a problem on your network, start at one end of your network and follow the routes to the other end, entering all or one of the following Junos OS command-line interfaces (CLI) operational mode commands:

```
user@host> ping (ip-address | host-name)
user@host> show route (ip-address | host-name)
user@host> traceroute (ip-address | host-name)
```

Sample Output

```
user@R6> ping 10.0.0.5
PING 10.0.0.5 (10.0.0.5): 56 data bytes
36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2db   0 0000  01  01 a8c6 10.1.26.2  10.0.0.5

36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2de   0 0000  01  01 a8c3 10.1.26.2  10.0.0.5

36 bytes from 10.1.26.1: Time to live exceeded
Vr HL TOS  Len   ID Flg  off TTL Pro  cks      Src      Dst
 4  5  00 0054 e2e2   0 0000  01  01 a8bf 10.1.26.2  10.0.0.5

^C
--- 10.0.0.5 ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss

user@R6> show route 10.0.0.5

inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32        *[IS-IS/165] 00:02:39, metric 10
                     > to 10.1.26.1  via so-0/0/2.0

user@R6> traceroute 10.0.0.5
```

```
traceroute to 10.0.0.5 (10.0.0.5), 30 hops max, 40 byte packets
 1  10.1.26.1 (10.1.26.1)  0.649 ms  0.521 ms  0.490 ms
 2  10.1.26.2 (10.1.26.2)  0.521 ms  0.537 ms  0.507 ms
 3  10.1.26.1 (10.1.26.1)  0.523 ms  0.536 ms  0.514 ms
 4  10.1.26.2 (10.1.26.2)  0.528 ms  0.551 ms  0.523 ms
 5  10.1.26.1 (10.1.26.1)  0.531 ms  0.550 ms  0.524 ms
```

Meaning

The sample output shows an unsuccessful **ping** command in which the packets are being rejected because the time to live is exceeded. The output for the **show route** command shows the interface (**10.1.26.1**) that you can examine further for possible problems. The **traceroute** command shows the loop between **10.1.26.1** (**R2**) and **10.1.26.2** (**R6**), as indicated by the continuous repetition of the two interface addresses.

Isolating the Causes of a Network Problem

Problem **Description:** A particular symptom can be the result of one or more causes. Narrow down the focus of your search to find each individual cause of the unwanted behavior.

Solution To isolate the cause of a particular problem, enter one or all of the following Junos OS CLI operational mode command:

```
user@host> show < configuration | bgp | interfaces | isis | ospf | route >
```

Your particular problem may require the use of more than just the commands listed above. See the appropriate command reference for a more exhaustive list of commonly used operational mode commands.

Sample Output

```
user@R6> show interfaces terse
Interface               Admin Link Proto Local                 Remote
so-0/0/0                up    up
so-0/0/0.0              up    up   inet  10.1.56.2/30
                                   iso
so-0/0/2                up    up
so-0/0/2.0              up    up   inet  10.1.26.2/30
                                   iso
so-0/0/3                up    up
so-0/0/3.0              up    up   inet  10.1.36.2/30
                                   iso
[...Output truncated...]
```

The following sample output is from **R2**:

```
user@R2> show route 10.0.0.5

inet.0: 22 destinations, 25 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32         *[Static/5] 00:16:21
                    > to 10.1.26.2 via so-0/0/2.0
                    [BGP/170] 3d 20:23:35, MED 5, localpref 100
```


Copyright © 2017, Juniper Networks, Inc.

```
      AS path: 65001 I
    > to 10.1.12.1 via so-0/0/0.0
```

Meaning

The sample output shows that all interfaces on **R6** are up. The output from **R2** shows that a static route [**Static/5**] configured on **R2** points to **R6** (**10.1.26.2**) and is the preferred route to **R5** because of its low preference value. However, the route is looping from **R2** to **R6**, as indicated by the missing reference to **R5** (**10.1.15.2**).

Taking Appropriate Action for Resolving the Network Problem

Problem **Description:** The appropriate action depends on the type of problem you have isolated. In this example, a static route configured on **R2** is deleted from the [**routing-options**] hierarchy level. Other appropriate actions might include the following:

Solution

- Check the local router's configuration and edit it if appropriate.
- Troubleshoot the intermediate router.
- Check the remote host configuration and edit it if appropriate.
- Troubleshoot routing protocols.
- Identify additional possible causes.

To resolve the problem in this example, enter the following Junos OS CLI commands:

```
[edit]
user@R2# delete routing-options static route destination-prefix
user@R2# commit and-quit
user@R2# show route destination-prefix
```

Sample Output

```
[edit]
user@R2#  delete routing-options static route 10.0.0.5/32

[edit]
user@R2#  commit and-quit
commit complete
Exiting configuration mode

user@R2>  show route 10.0.0.5

inet.0: 22 destinations, 24 routes (22 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32         *[BGP/170] 3d 20:26:17, MED 5, localpref 100
                       AS path: 65001 I
                     > to 10.1.12.1 via so-0/0/0.0
```

Copyright © 2017, Juniper Networks, Inc.

Meaning

The sample output shows the static route deleted from the [**routing-options**] hierarchy and the new configuration committed. The output for the **show route** command now shows the BGP route as the preferred route, as indicated by the asterisk (*).

Evaluating the Solution to Check Whether the Network Problem Is Resolved

Problem **Description:** If the problem is solved, you are finished. If the problem remains or a new problem is identified, start the process over again.

You can address possible causes in any order. In relation to the network in "Isolating a Broken Network Connection" on page 166, we chose to work from the local router toward the remote router, but you might start at a different point, particularly if you have reason to believe that the problem is related to a known issue, such as a recent change in configuration.

Solution To evaluate the solution, enter the following Junos OS CLI commands:

```
user@host> show route (ip-address |host-name)
user@host> ping (ip-address | host-name)
user@host> traceroute (ip-address | host-name)
```

Sample Output

```
user@R6> show route 10.0.0.5

inet.0: 20 destinations, 20 routes (20 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32         *[BGP/170]  00:01:35, MED 5, localpref 100, from 10.0.0.2
                      AS path: 65001 I
                    > to 10.1.26.1 via so-0/0/2.0

user@R6> ping 10.0.0.5
PING 10.0.0.5 (10.0.0.5): 56 data bytes
64 bytes from 10.0.0.5: icmp_seq=0 ttl=253 time=0.866 ms
64 bytes from 10.0.0.5: icmp_seq=1 ttl=253 time=0.837 ms
64 bytes from 10.0.0.5: icmp_seq=2 ttl=253 time=0.796 ms
^C
--- 10.0.0.5 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max/stddev = 0.796/0.833/0.866/0.029 ms

user@R6> traceroute 10.0.0.5
traceroute to 10.0.0.5 (10.0.0.5), 30 hops max, 40 byte packets
 1  10.1.26.1 (10.1.26.1)  0.629 ms  0.538 ms  0.497 ms
 2  10.1.12.1 (10.1.12.1)  0.534 ms  0.538 ms  0.510 ms
 3  10.0.0.5 (10.0.0.5)  0.776 ms  0.705 ms  0.672 ms
```

Meaning

The sample output shows that there is now a connection between **R6** and **R5.** The **show route** command shows that the BGP route to **R5** is preferred, as indicated by the asterisk

 Copyright © 2017, Juniper Networks, Inc.

(*). The **ping** command is successful and the **traceroute** command shows that the path from **R6** to **R5** is through **R2** (**10.1.26.1**), and then through **R1** (**10.1.12.1**).

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 19

Monitoring and Troubleshooting a RIP Network

- Monitoring RIP Routing Information on page 173
- Verifying a RIP Configuration on page 174

Monitoring RIP Routing Information

Purpose

NOTE: This topic applies only to the J-Web Application package.

Use the monitoring functionality to monitor RIP routing on routing devices.

Action

To view RIP routing information in the J-Web interface, select **Monitor** > **Routing** > **RIP Information**.

To view RIP routing information in the CLI, enter the following CLI commands:

- **show rip statistics**
- **show rip neighbor**

Meaning

Table 9 on page 173 summarizes key output fields in the RIP routing display in the J-Web interface.

Table 9: Summary of Key RIP Routing Output Fields

Field	Values	Additional Information
RIP Statistics		
Protocol Name	The RIP protocol name.	
Port number	The port on which RIP is enabled.	
Hold down time	The interval during which routes are neither advertised nor updated.	

Copyright © 2017, Juniper Networks, Inc.

Table 9: Summary of Key RIP Routing Output Fields *(continued)*

Field	Values	Additional Information
Global routes learned	Number of RIP routes learned on the logical interface.	
Global routes held down	Number of RIP routes that are not advertised or updated during the hold-down interval.	
Global request dropped	Number of requests dropped.	
Global responses dropped	Number of responses dropped.	
RIP Neighbors		
Neighbor	Name of the RIP neighbor.	This value is the name of the interface on which RIP is enabled. Click the name to see the details for this neighbor.
State	State of the RIP connection: **Up** or **Dn** (Down).	
Source Address	Local source address.	This value is the configured address of the interface on which RIP is enabled.
Destination Address	Destination address.	This value is the configured address of the immediate RIP adjacency.
Send Mode	The mode of sending RIP messages.	
Receive Mode	The mode in which messages are received.	
In Metric	Value of the incoming metric configured for the RIP neighbor.	

Verifying a RIP Configuration

To verify a RIP configuration, perform the following tasks:

- Verifying the RIP-Enabled Interfaces on page 174
- Verifying Reachability of All Hosts in the RIP Network on page 175

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.

Copyright © 2017, Juniper Networks, Inc.

Sample Output

```
user@host> show rip neighbor
Source           Destination     Send   Receive   In
Neighbor           State  Address         Address          Mode   Mode     Met
--------           -----  -------         -----------      ----   -------  ---
ge-0/0/0.0            Dn (null)          (null)           mcast  both       1
ge-0/0/1.0            Up 192.168.220.5   224.0.0.9        mcast  both       1
```

Meaning The output shows a list of the RIP neighbors that are configured on the device. Verify the following information:

- Each configured interface is present. Interfaces are listed in alphabetical order.
- Each configured interface is up. The state of the interface is listed in the **Destination State** column. A state of **Up** indicates that the link is passing RIP traffic. A state of **Dn** indicates that the link is not passing RIP traffic. In a point-to-point link, this state generally means that either the end point is not configured for RIP or the link is unavailable.

See Also
- Verifying a RIP Configuration on page 174

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 Juniper Networks device.

Action For each device in the RIP network:

1. In the J-Web interface, select **Troubleshoot>Traceroute**.
2. In the Remote Host box, type the name of a host for which you want to verify reachability from the device.
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
```

Meaning Each numbered row in the output indicates a routing hop in the path to the host. The three-time increments indicate the round-trip time (RTT) between the device and the hop for each traceroute packet.

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

Copyright © 2017, Juniper Networks, Inc.

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

See Also

- RIP Overview on page 3

Related Documentation

- RIP Configuration Overview on page 17
- **show rip statistics on page 257** in the CLI Explorer
- **show rip neighbor on page 255** in the CLI Explorer
- *traceroute* in the CLI Explorer
- RIP Overview on page 3

Copyright © 2017, Juniper Networks, Inc.

PART 5

Configuration Statements and Operational Commands

- RIP Configuration Statements on page 179
- RIPng Configuration Statements on page 215
- RIP Operational Commands on page 233
- RIPng Operational Commands on page 453

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 20

RIP Configuration Statements

- any-sender on page 180
- authentication-key (Protocols RIP) on page 181
- authentication-type (Protocols RIP) on page 182
- bfd-liveness-detection (Protocols RIP) on page 183
- check-zero on page 186
- demand-circuit (Protocols RIP) on page 187
- dynamic-peers on page 188
- export on page 189
- graceful-restart (Protocols RIP) on page 190
- group (Protocols RIP) on page 191
- holddown (Protocols RIP) on page 193
- import (Protocols RIP) on page 194
- interface-type (Protocols RIP) on page 195
- max-retrans-time on page 196
- message-size on page 197
- metric-in (Protocols RIP) on page 198
- metric-out on page 199
- neighbor on page 200
- peer (Protocols RIP) on page 201
- preference (Protocols RIP) on page 202
- receive (Protocols RIP) on page 203
- rib-group (Protocols RIP) on page 204
- rip on page 205
- route-timeout (Protocols RIP) on page 206
- routing-instances (Multiple Routing Entities) on page 207
- send (Protocols RIP) on page 209
- traceoptions (Protocols RIP) on page 210
- update-interval (Protocols RIP) on page 213

Copyright © 2017, Juniper Networks, Inc.

any-sender

Syntax any-sender;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced in Junos OS Release 8.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Disable strict sender address checks.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

authentication-key (Protocols RIP)

Syntax authentication-key *password*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Require authentication for RIP route queries received on an interface.

Options
password—Authentication password. If the password does not match, the packet is rejected. The password can be from 1 through 16 contiguous characters long and can include any ASCII strings.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Route Authentication for RIP on page 31

Copyright © 2017, Juniper Networks, Inc.

authentication-type (Protocols RIP)

Syntax authentication-type *type*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the type of authentication for RIP route queries received on an interface.

Default If you do not include this statement and the **authentication-key** statement, RIP authentication is disabled.

Options *type*—Authentication type:

- **md5**—Use the MD5 algorithm to create an encoded checksum of the packet. The encoded checksum is included in the transmitted packet. The receiving routing device uses the authentication key to verify the packet, discarding it if the digest does not match. This algorithm provides a more secure authentication scheme.
- **none**—Disable authentication. If **none** is configured, the configured authentication key is ignored.
- **simple**—Use a simple password. The password is included in the transmitted packet, which makes this method of authentication relatively insecure. The password can be from 1 through 16 contiguous letters or digits long.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Route Authentication for RIP on page 31
- authentication-key on page 181

Copyright © 2017, Juniper Networks, Inc.

bfd-liveness-detection (Protocols RIP)

Syntax

```
bfd-liveness-detection {
  authentication {
    algorithm algorithm-name;
    key-chain key-chain-name;
    loose-check;
  }
  detection-time {
    threshold milliseconds;
  }
  minimum-interval milliseconds;
  minimum-receive-interval milliseconds;
  multiplier number;
  no-adaptation;
  transmit-interval {
    minimum-interval milliseconds;
    threshold milliseconds;
  }
  version (1 | automatic);
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rip group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip group *group-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information

Statement introduced in Junos OS Release 8.0.
Options **detection-time threshold** and **transmit-interval threshold** introduced in Junos OS Release 8.2.
Support for logical systems introduced in Junos OS Release 8.3.
Option **no-adaptation** introduced in Junos OS Release 9.0.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Options **authentication algorithm**, **authentication key-chain**, and **authentication loose-check** introduced in Junos OS Release 9.6.
Options **authentication algorithm**, **authentication key-chain**, and **authentication loose-check** introduced in Junos OS Release 9.6 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure bidirectional failure detection timers and authentication.

The remaining statements are explained separately. See CLI Explorer.

Options

authentication algorithm *algorithm-name* —Configure the algorithm used to authenticate the specified BFD session: **simple-password**, **keyed-md5**, **keyed-sha-1**, **meticulous-keyed-md5**, or **meticulous-keyed-sha-1**.

Copyright © 2017, Juniper Networks, Inc.

authentication key-chain ***key-chain-name***—Associate a security key with the specified BFD session using the name of the security keychain. The name you specify must match one of the keychains configured in the **authentication-key-chains key-chain** statement at the **[edit security]** hierarchy level.

authentication loose-check—(Optional) Configure loose authentication checking on the BFD session. Use only for transitional periods when authentication is not configured at both ends of the BFD session.

detection-time threshold ***milliseconds***—Configure a threshold for the adaptation of the BFD session detection time. When the detection time adapts to a value equal to or greater than the threshold, a single trap and a single system log message are sent.

minimum-interval ***milliseconds***—Configure the minimum interval after which the local routing device transmits a hello packet and then expects to receive a reply from the neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can specify the minimum transmit and receive intervals separately using the **transmit-interval minimum-interval** and **minimum-receive-interval** statements.

Range: 1 through 255,000 milliseconds

minimum-receive-interval ***milliseconds***—Configure the minimum interval after which the local routing device expects to receive a reply from a neighbor with which it has established a BFD session. Optionally, instead of using this statement, you can configure the minimum receive interval using the **minimum-interval** statement.

Range: 1 through 255,000 milliseconds

multiplier ***number***—Configure the number of hello packets not received by a neighbor that causes the originating interface to be declared down.

Range: 1 through 255

Default: 3

no-adaptation—Configure BFD sessions not to adapt to changing network conditions. We recommend that you not disable BFD adaptation unless it is preferable not to have BFD adaptation enabled in your network.

transmit-interval threshold ***milliseconds***—Configure the threshold for the adaptation of the BFD session transmit interval. When the transmit interval adapts to a value greater than the threshold, a single trap and a single system message are sent. The interval threshold must be greater than the minimum transmit interval.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

transmit-interval minimum-interval ***milliseconds***—Configure a minimum interval after which the local routing device transmits hello packets to a neighbor. Optionally, instead of using this statement, you can configure the minimum transmit interval using the **minimum-interval** statement.

Range: 1 through 255,000

Copyright © 2017, Juniper Networks, Inc.

version—Configure the BFD version to detect: **1** (BFD version 1) or **automatic** (autodetect the BFD version).

Default: automatic

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring BFD for RIP on page 52
- Example: Configuring BFD Authentication for RIP on page 59

Copyright © 2017, Juniper Networks, Inc.

check-zero

Syntax (check-zero | no-check-zero);

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description Check whether the reserved fields in a RIP packet are zero:

- **check-zero**—Discard version 1 packets that have nonzero values in the reserved fields and version 2 packets that have nonzero values in the fields that must be zero. This default behavior implements the RIP version 1 and version 2 specifications.
- **no-check-zero**—Receive RIP version 1 packets with nonzero values in the reserved fields or RIP version 2 packets with nonzero values in the fields that must be zero. This is in spite of the fact that they are being sent in violation of the specifications in RFC 1058 and RFC 2453.

Default check-zero

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

Copyright © 2017, Juniper Networks, Inc.

demand-circuit (Protocols RIP)

Syntax demand-circuit;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip group *group-name*],
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip group *group-name*],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information Statement introduced in Release 11.1 of Junos OS.

Description Configure a neighboring interface to act as a RIP demand circuit. To complete the demand circuit, you must configure both ends of the pair as demand circuits. When configured, the device sends RIP information only when changes occur in the routing database.

Default Disabled. You must explicitly configure two neighboring interfaces to act as a RIP demand circuit.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring RIP Demand Circuits on page 47
- RIP Demand Circuits Overview on page 44
- max-retrans-time on page 196

Copyright © 2017, Juniper Networks, Inc.

dynamic-peers

Syntax dynamic-peers;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip group *group-name* neighbor *neighbor-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure an interface to have dynamic peers in a point-to-multipoint RIP network.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Point-to-Multipoint RIP Networks on page 78

Copyright © 2017, Juniper Networks, Inc.

export

Syntax export [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols rip group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name*],
[edit protocols rip group *group-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Apply a policy to routes being exported to the neighbors.

By default, RIP does not export routes it has learned to its neighbors. To enable RIP to export routes, apply one or more export policies.

If no routes match the policies, the local routing device does not export any routes to its neighbors. Export policies override any metric values determined through calculations involving the values configured with the metric-in and metric-out statements.

NOTE: **The export policy on RIP does not support manipulating routing information of the next hop.**

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- import on page 194

Copyright © 2017, Juniper Networks, Inc.

graceful-restart (Protocols RIP)

Syntax

```
graceful-restart {
  disable;
  restart-time seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rip],
[edit protocols rip]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description

Configure graceful restart for RIP.

Options

disable—Disables graceful restart for RIP.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Junos OS High Availability Library for Routing Devices*

Copyright © 2017, Juniper Networks, Inc.

group (Protocols RIP)

Syntax

```
group group-name {
  bfd-liveness-detection {
    authentication {
      algorithm algorithm-name;
      key-chain key-chain-name;
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    transmit-interval {
      threshold milliseconds;
      minimum-interval milliseconds;
    }
    multiplier number;
    version (0 | 1 | automatic);
  }
  demand-circuit;
  export policy;
  max-retrans-time  seconds;
  metric-out metric;
  preference number;
  route-timeout seconds;
  update-interval seconds;
  neighbor neighbor-name {
    authentication-key password;
    authentication-type type;
    bfd-liveness-detection {
      authentication {
        algorithm algorithm-name;
        key-chain key-chain-name;
        loose-check;
      }
      detection-time {
        threshold milliseconds;
      }
      minimum-interval milliseconds;
      minimum-receive-interval milliseconds;
      transmit-interval {
        threshold milliseconds;
        minimum-interval milliseconds;
      }
      multiplier number;
      version (0 | 1 | automatic);
    }
    (check-zero | no-check-zero);
    demand-circuit;
    import policy-name;
    max-retrans-time seconds;
    message-size number;
```

Copyright © 2017, Juniper Networks, Inc.

```
        metric-in metric;
        metric-out metric;
        receive receive-options;
        route-timeout seconds;
        send send-options;
        update-interval seconds;
    }
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit protocols rip],
[edit routing-instances *routing-instance-name* protocols rip]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Configure a set of RIP neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group. Each group must contain at least one neighbor. You should create a group for every export policy.

Options

group-name—Name of a group, up to 16 characters long.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring RIP*

Copyright © 2017, Juniper Networks, Inc.

holddown (Protocols RIP)

Syntax holddown *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit protocols rip],
[edit routing-instances *routing-instance-name* protocols rip]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.

Description
Configure how long the expired route is retained in the routing table before being removed.

When the hold-down timer runs on RIP demand circuits, routes are advertised as unreachable on other interfaces. When the hold-down timer expires, the route is removed from the routing table if all destinations detect that the route is unreachable or the remaining destinations are down.

Options
seconds—Estimated time to wait before making updates to the routing table.

Range: 10 through 180 seconds

Default: 180 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring RIP Timers on page 38
- RIP Demand Circuits Overview on page 44

Copyright © 2017, Juniper Networks, Inc.

import (Protocols RIP)

Syntax import [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Apply one or more policies to routes being imported by the local routing device from neighbors.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Applying Policies to RIP Routes Imported from Neighbors on page 85
- Junos OS Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices
- export on page 189

Copyright © 2017, Juniper Networks, Inc.

interface-type (Protocols RIP)

Syntax interface-type p2mp;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip group *group-name* neighbor *neighbor-name*]

Release Information Statement introduced in Junos OS Release 11.4.

Description Configure the type of interface in a RIP network.

This statement enables a RIP device to have single or multiple peers through an interface.

Options **p2mp**—Configure an interface in a RIP network as a point-to-multipoint interface.

Required Privilege Level interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring Point-to-Multipoint RIP Networks on page 78

Copyright © 2017, Juniper Networks, Inc.

max-retrans-time

Syntax max-retrans-time *seconds*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols rip group *group-name*],
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip group *group-name*],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information Statement introduced in Junos OS Release 11.1.

Description RIP demand circuits send update messages every 5 seconds to an unresponsive peer. Configure the retransmission timer to limit the number of times the demand circuit resends update messages to an unresponsive peer. If the configured retransmission threshold is reached, routes from the next hop router are marked as unreachable and the hold-down timer starts. You must configure a pair of RIP demand circuits for this timer to take effect.

To determine the number of times to resend the update message, use the following calculation:

5 seconds x number of retransmissions = retransmission seconds

Options ***seconds***—The total amount of time the demand circuit resends update messages to an unresponsive peer. The seconds range corresponds to sending an update message a minimum of 1 time (5 seconds) and a maximum of 36 times (180 seconds).

Range: 5 through 180 seconds

Default: 5 seconds

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring RIP Demand Circuits on page 47
- RIP Demand Circuits Overview on page 44
- demand-circuit on page 187

Copyright © 2017, Juniper Networks, Inc.

message-size

Syntax message-size *number*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Specify the number of route entries to be included in every RIP update message. To ensure interoperability with other vendors' equipment, use the standard of 25 route entries per message.

Options ***number***—Number of route entries per update message.

Range: 25 through 255 entries

Default: 25 entries

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

metric-in (Protocols RIP)

Syntax metric-in *metric*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Specify the metric to add to incoming routes when the routing device advertises into RIP routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIP routes learned through a specific neighbor.

Options *metric*—Metric value.

Range: 1 through 16

Default: 1

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Configuring the Metric Value Added to Imported RIP Routes on page 72

Copyright © 2017, Juniper Networks, Inc.

metric-out

Syntax metric-out *metric*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIP routes sent from this neighbor.

Options *metric*—Metric value.

Range: 1 through 16

Default: 1

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

neighbor

Syntax

```
neighbor neighbor-name {
  authentication-key password;
  authentication-type type;
  bfd-liveness-detection {
    authentication {
      algorithm algorithm-name;
      key-chain key-chain-name;
      loose-check;
    }
    detection-time {
      threshold milliseconds;
    }
    minimum-interval milliseconds;
    minimum-receive-interval milliseconds;
    transmit-interval {
      threshold milliseconds;
      minimum-interval milliseconds;
    }
    multiplier number;
    version (0 | 1 | automatic);
  }
  (check-zero | no-check-zero);
  demand-circuit;
  import policy-name;
  max-retrans-time seconds;
  message-size number;
  metric-in metric;
  metric-out metric;
  receive receive-options;
  route-timeout seconds;
  send send-options;
  update-interval seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rip group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name*],
[edit protocols rip group *group-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description

Configure neighbor-specific RIP parameters, thereby overriding the defaults set for the routing device.

Options

neighbor-name—Name of an interface over which a routing device communicates to its neighbors.

Copyright © 2017, Juniper Networks, Inc.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

peer (Protocols RIP)

Syntax
peer *IP address*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced in Junos OS Release 11.4.

Description
Configure a static peer for an interface in a point-to-multipoint RIP network.

Options
address—IP address of the static peer to be configured.

Required Privilege Level
interface—To view this statement in the configuration.
interface-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring Point-to-Multipoint RIP Networks on page 78

Copyright © 2017, Juniper Networks, Inc.

preference (Protocols RIP)

Syntax preference *preference*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name*],
[edit protocols rip group *group-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.

Description
Specify the preference of external routes learned by RIP as compared to those learned from other routing protocols.

By default, Junos OS assigns a preference of 100 to routes that originate from RIP. When Junos OS determines a route's preference to become the active route, the software selects the route with the lowest preference and installs this route into the forwarding table.

Options
preference—Preference value. A lower value indicates a more preferred route.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: 100

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Route Preferences Overview*

Copyright © 2017, Juniper Networks, Inc.

receive (Protocols RIP)

Syntax receive *receive-options*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure RIP receive options.

Options ***receive-options***—One of the following:

- **both**—Accept both RIP version 1 and version 2 packets.
- **none**—Do not receive RIP packets.
- **version-1**—Accept only RIP version 1 packets.
- **version-2**—Accept only RIP version 2 packets.

Default: both

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 93
- send on page 209

Copyright © 2017, Juniper Networks, Inc.

rib-group (Protocols RIP)

Syntax rib-group *group-name*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit protocols rip],
[edit routing-instances *routing-instance-name* protocols rip]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Install RIP routes into multiple routing tables by configuring a routing table group.

Options *group-name*—Name of the routing table group.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Redistributing Routes Between Two RIP Instances on page 100

Copyright © 2017, Juniper Networks, Inc.

rip

Syntax rip {...}

Hierarchy Level
[edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Enable RIP routing on the routing device.

Default RIP is disabled on the routing device.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring RIP*

Copyright © 2017, Juniper Networks, Inc.

route-timeout (Protocols RIP)

Syntax route-timeout *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit logical-systems *logical-system-name* protocols rip group *group-name*],
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip],
[edit protocols rip group *group-name*],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip],
[edit routing-instances *routing-instance-name* protocols rip group *group-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the route timeout interval for RIP. If a route is not refreshed after being installed in the routing table by the specified timeout interval, the route is marked as invalid and is removed from the routing table after the hold-down period expires.

Options ***seconds***—Estimated time to wait before making updates to the routing table.

Range: 30 through 360 seconds

Default: 180 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring RIP Timers on page 38
- RIP Demand Circuits Overview on page 44

Copyright © 2017, Juniper Networks, Inc.

routing-instances (Multiple Routing Entities)

Syntax routing-instances *routing-instance-name* { ... }

Hierarchy Level [edit],
[edit logical-systems *logical-system-name*]

Release Information Statement introduced before Junos OS Release 7.4.
remote-vtep-v6-list statement introduced in Junos OS Release 17.3 for MX Series routers with MPC and MIC interfaces.

Description Configure an additional routing entity for a router. You can create multiple instances of BGP, IS-IS, OSPF, OSPFv3, and RIP for a router. You can also create multiple routing instances for separating routing tables, routing policies, and interfaces for individual wholesale subscribers (retailers) in a Layer 3 wholesale network.

Each routing instance consist of the following:

- A set of routing tables
- A set of interfaces that belong to these routing tables
- A set of routing option configurations

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table is my-instance.inet.0. All routes for **my-instance** are installed into my-instance.inet.0.

Routes are installed into the default routing instance inet.0 by default, unless a routing instance is specified.

In Junos OS Release 9.0 and later, you can no longer specify a routing-instance name of *master*, *default*, or *bgp* or include special characters within the name of a routing instance.

In Junos OS Release 9.6 and later, you can include a slash (/) in a routing-instance name only if a logical system is not configured. That is, you cannot include the slash character in a routing-instance name if a logical system other than the default is explicitly configured. Routing-instance names, further, are restricted from having the form __.*__ (beginning and ending with underscores). The colon **:** character cannot be used when multitopology routing (MTR) is enabled.

Default Routing instances are disabled for the router.

Options ***routing-instance-name***——Name of the routing instance. This must be a non-reserved string of not more than 128 characters.

remote-vtep-list—Configure static remote VXLAN tunnel endpoints.

remote-vtep-v6-list—Configure static IPv6 remote VXLAN tunnel endpoints.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Example: Configuring Interprovider Layer 3 VPN Option A*
- *Example: Configuring Interprovider Layer 3 VPN Option B*
- *Example: Configuring Interprovider Layer 3 VPN Option C*

Copyright © 2017, Juniper Networks, Inc.

send (Protocols RIP)

Syntax send *send-options*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure RIP send options.

Options *send-options*—One of the following:

- **broadcast**—Broadcast RIP version 2 packets (RIP version 1 compatible).
- **multicast**—Multicast RIP version 2 packets. This is the default.
- **none**—Do not send RIP updates.
- **version-1**—Broadcast RIP version 1 packets.

Default: multicast

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring the Sending and Receiving of RIPv1 and RIPv2 Packets on page 93
- receive on page 203

Copyright © 2017, Juniper Networks, Inc.

traceoptions (Protocols RIP)

Syntax

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

Hierarchy Level

[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit protocols rip],
[edit routing-instances *routing-instance-name* protocols rip]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Set RIP protocol-level tracing options.

NOTE: The **traceoptions** statement is not supported on QFabric systems.

Default

The default RIP protocol-level trace options are inherited from the global **traceoptions** statement.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIP tracing output in the file **/var/log/rip-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

RIP Tracing Options

Copyright © 2017, Juniper Networks, Inc.

- **auth**—RIP authentication
- **error**—RIP error packets
- **expiration**—RIP route expiration processing
- **holddown**—RIP hold-down processing
- **nsr-synchronization**—Nonstop routing synchronization events
- **packets**—All RIP packets
- **request**—RIP information packets such as request, poll, and poll entry packets
- **trigger**—RIP triggered updates
- **update**—RIP update packets

Global Tracing Options

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **receive-detail**—Provide detailed trace information for packets being received.
- **send**—Trace the packets being transmitted.
- **send-detail**—Provide detailed trace information for packets being transmitted.

no-world-readable—(Optional) Prevent any user from reading the log file.

Copyright © 2017, Juniper Networks, Inc.

size *size*—(Optional) Maximum size of each trace file, in kilobytes (KB) or megabytes (MB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- Example: Tracing RIP Protocol Traffic on page 108

Copyright © 2017, Juniper Networks, Inc.

update-interval (Protocols RIP)

Syntax update-interval *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols rip],
[edit logical-systems *logical-system-name* protocols rip group *group-name*],
[edit logical-systems *logical-system-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*],
[edit protocols rip],
[edit protocols rip group *group-name*],
[edit protocols rip group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols rip],
[edit routing-instances *routing-instance-name* protocols rip group *group-name*],
[edit routing-instances *routing-instance-name* protocols rip group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 12.1 for the QFX Series.
Statement introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Configure the interval at which routes learned by RIP are sent to neighbors. This timer controls the interval between routing updates. This timer is set to 30 seconds, by default, with a small random amount of time added when the timer is reset. This added time prevents congestion that can happen if all routing devices update their neighbors simultaneously.

Options ***seconds***—Estimated time to wait before making updates to the routing table.

Range: 10 through 60 seconds

Default: 30 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring RIP Timers on page 38

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 21

RIPng Configuration Statements

- export (Protocols RIPng) on page 216
- graceful-restart (Protocols RIPng) on page 217
- group (Protocols RIPng) on page 218
- holddown (Protocols RIPng) on page 219
- import (Protocols RIPng) on page 220
- metric-in (Protocols RIPng) on page 221
- metric-out (Protocols RIPng) on page 222
- neighbor (Protocols RIPng) on page 223
- preference (Protocols RIPng) on page 224
- receive (Protocols RIPng) on page 225
- ripng on page 226
- route-timeout (Protocols RIPng) on page 227
- routing-instances (Multiple Routing Entities) on page 228
- send (Protocols RIPng) on page 229
- traceoptions (Protocols RIPng) on page 230
- update-interval (Protocols RIPng) on page 232

Copyright © 2017, Juniper Networks, Inc.

export (Protocols RIPng)

Syntax export [*policy-names*];

Hierarchy Level [edit logical-systems *logical-system-name* protocols ripng group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name*],
[edit protocols ripng group *group-name*],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description Apply a policy or list of policies to routes being exported to the neighbors.

By default, RIPng does not export routes it has learned to its neighbors. To have RIPng export routes, apply one or more export policies. To apply export policies and to filter routes being exported from the local routing device to its neighbors, include the **export** statement and list the name of the policy to be evaluated.

You can define one or more export policies. If no routes match the policies, the local routing device does not export any routes to its neighbors. Export policies override any metric values determined through calculations involving the values configured with the **metric-in** and **metric-out** statements.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- import on page 220

 Copyright © 2017, Juniper Networks, Inc.

graceful-restart (Protocols RIPng)

Syntax

```
graceful-restart {
  disable;
  restart-time seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description

Configure graceful restart for RIPng.

Options

disable—Disables graceful restart for RIPng.

The remaining statement is explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation

- *Junos OS High Availability Library for Routing Devices*

Copyright © 2017, Juniper Networks, Inc.

group (Protocols RIPng)

Syntax

```
group group-name {
  export [ policy-names ];
  metric-out metric;
  neighbor neighbor-name {
    import policy-name;
    metric-in metric;
    receive <none>;
    route-timeout seconds;
    send <none>;
    update-interval seconds;
  }
  preference number;
  route-timeout seconds;
  update-interval seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description

Configure a set of RIPng neighbors that share an export policy and metric. The export policy and metric govern what routes to advertise to neighbors in a given group.

Each group must contain at least one neighbor. You should create a group for each export policy that you have.

Options

group-name—Name of a group, up to 16 characters long.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

holddown (Protocols RIPng)

Syntax holddown *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description Configure how long the expired route is retained in the routing table before being removed.

Options
seconds—Estimated time to wait before removing expired routes from the routing table.
Default: 180 seconds
Range: 10 through 180 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring RIPng Update Interval on page 146

Copyright © 2017, Juniper Networks, Inc.

import (Protocols RIPng)

Syntax import [*policy-names*];

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit protocols ripng],
[edit protocols ripng group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description Apply one or more policies to routes being imported into the local routing device from its neighbors.

Options ***policy-names***—Name of one or more policies.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Applying Policies to RIPng Routes Imported from Neighbors on page 123
- export on page 216

Copyright © 2017, Juniper Networks, Inc.

metric-in (Protocols RIPng)

Syntax metric-in *metric*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit protocols ripng],
[edit protocols ripng group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description Specify the metric to add to incoming routes when advertising into RIPng routes that were learned from other protocols. Use this statement to configure the routing device to prefer RIPng routes learned through a specific neighbor.

Options
metric—Metric value.
Range: 1 through 16
Default: 1

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring the Metric Value Added to Imported RIPng Routes to Control the Route Selection Process on page 138

Copyright © 2017, Juniper Networks, Inc.

metric-out (Protocols RIPng)

Syntax metric-out *metric*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit protocols ripng group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description
Specify the metric value to add to routes transmitted to the neighbor. Use this statement to control how other routing devices prefer RIPng routes sent from this neighbor.

When an export policy is configured, RIPng exports all learned routes to neighbors configured with the **neighbor** statement.

If a route being exported was learned from a member of the same RIPng group, the metric associated with that route (unless modified by an export policy) is the normal RIPng metric. For example, a RIPng route with a metric of 5 learned from a neighbor configured with a **metric-in** value of 2 is advertised with a combined metric of 7 when advertised to RIPng neighbors in the same group. However, if this route was learned from a RIPng neighbor in a different group or from a different protocol, the route is advertised with the metric value configured for that group with the **metric-out** statement. The default value for **metric-out** is 1.

To modify the metric for routes advertised outside a group, include the **metric-out** statement.

Options
metric—Metric value.
Range: 1 through 16
Default: 1

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring the Metric Value Added to Imported RIPng Routes to Control the Route Selection Process on page 138
- Understanding RIPng Traffic Control with Metrics for Optimizing the Path Cost on page 137

Copyright © 2017, Juniper Networks, Inc.

neighbor (Protocols RIPng)

Syntax

```
neighbor neighbor-name {
  import [ policy-names ];
  metric-in metric;
  receive <none>;
  route-timeout seconds;
  send <none>;
  update-interval seconds;
}
```

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ripng group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name*],
[edit protocols ripng group *group-name*],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name*]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description

Configure neighbor-specific RIPng parameters, thereby overriding the defaults set for the routing device.

Options

neighbor-name—Name of an interface over which a routing device communicates to its neighbors.

The remaining statements are explained separately. See CLI Explorer.

Required Privilege Level

routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

preference (Protocols RIPng)

Syntax preference *preference*;

Hierarchy Level [edit logical-systems *logical-system-name* protocols ripng group *group-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name*],
[edit protocols ripng group *group-name*],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name*]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description Specify the preference of external routes learned by RIPng as compared to those learned from other routing protocols.

By default, Junos OS assigns a preference of 100 to routes that originate from RIPng. When Junos OS determines that a route is to become the active route, the software selects the route with the lowest preference and installs this route into the forwarding table.

To modify the default RIPng preference value, include the **preference** statement.

Options ***preference***—Preference value. A lower value indicates a more preferred route.

Range: 0 through 4,294,967,295 ($2^{32} - 1$)

Default: 100

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

receive (Protocols RIPng)

Syntax receive <none>;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit protocols ripng],
[edit protocols ripng group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description Enable or disable receiving of update messages.

Options **none**—(Optional) Disable receiving update messages.
Default: Enabled

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- send on page 229

ripng

Syntax ripng {...}

Hierarchy Level [edit logical-systems *logical-system-name* protocols],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols],
[edit protocols],
[edit routing-instances *routing-instance-name* protocols]

Release Information Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description Enable RIPng routing on the routing device.

Default RIPng is disabled on the routing device.

Required Privilege Level routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Copyright © 2017, Juniper Networks, Inc.

route-timeout (Protocols RIPng)

Syntax route-timeout *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng]

Release Information
Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description Configure the route timeout interval for RIPng.

Options
seconds—Estimated time to wait before making updates to the routing table.
Range: 30 through 360 seconds
Default: 180 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring RIPng Update Interval on page 146

Copyright © 2017, Juniper Networks, Inc.

routing-instances (Multiple Routing Entities)

Syntax routing-instances *routing-instance-name* { ... }

Hierarchy Level [edit],
[edit logical-systems *logical-system-name*]

Release Information Statement introduced before Junos OS Release 7.4.
remote-vtep-v6-list statement introduced in Junos OS Release 17.3 for MX Series routers with MPC and MIC interfaces.

Description Configure an additional routing entity for a router. You can create multiple instances of BGP, IS-IS, OSPF, OSPFv3, and RIP for a router. You can also create multiple routing instances for separating routing tables, routing policies, and interfaces for individual wholesale subscribers (retailers) in a Layer 3 wholesale network.

Each routing instance consist of the following:

- A set of routing tables
- A set of interfaces that belong to these routing tables
- A set of routing option configurations

Each routing instance has a unique name and a corresponding IP unicast table. For example, if you configure a routing instance with the name **my-instance**, its corresponding IP unicast table is my-instance.inet.0. All routes for **my-instance** are installed into my-instance.inet.0.

Routes are installed into the default routing instance inet.0 by default, unless a routing instance is specified.

In Junos OS Release 9.0 and later, you can no longer specify a routing-instance name of *master*, *default*, or *bgp* or include special characters within the name of a routing instance.

In Junos OS Release 9.6 and later, you can include a slash (/) in a routing-instance name only if a logical system is not configured. That is, you cannot include the slash character in a routing-instance name if a logical system other than the default is explicitly configured. Routing-instance names, further, are restricted from having the form __.*__ (beginning and ending with underscores). The colon **:** character cannot be used when multitopology routing (MTR) is enabled.

Default Routing instances are disabled for the router.

Options ***routing-instance-name***——Name of the routing instance. This must be a non-reserved string of not more than 128 characters.

remote-vtep-list—Configure static remote VXLAN tunnel endpoints.

remote-vtep-v6-list—Configure static IPv6 remote VXLAN tunnel endpoints.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- *Example: Configuring Interprovider Layer 3 VPN Option A*
- *Example: Configuring Interprovider Layer 3 VPN Option B*
- *Example: Configuring Interprovider Layer 3 VPN Option C*

send (Protocols RIPng)

Syntax
send <none>;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit logical-systems *logical-system-name* routing-instances *routing-instances-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*],
[edit protocols ripng],
[edit protocols ripng group *group-name* neighbor *neighbor-name*],
[edit routing-instances *routing-instance-name* protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng group *group-name* neighbor *neighbor-name*]

Release Information
Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description
Enable or disable sending of update messages.

Options
none—(Optional) Disable sending of update messages.
Default: Enabled

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- receive on page 225

Copyright © 2017, Juniper Networks, Inc.

traceoptions (Protocols RIPng)

Syntax

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

Hierarchy Level

[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng]

Release Information

Statement introduced before Junos OS Release 7.4.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description

Set RIPng protocol-level tracing options.

Default

The default RIPng protocol-level trace options are inherited from the global **traceoptions** statement.

Options

disable—(Optional) Disable the tracing operation. One use of this option is to disable a single operation when you have defined a broad group of tracing operations, such as **all**.

file *filename*—Name of the file to receive the output of the tracing operation. Enclose the name in quotation marks. We recommend that you place RIPng tracing output in the file **/var/log/ripng-log**.

files *number*—(Optional) Maximum number of trace files. When a trace file named ***trace-file*** reaches its maximum size, it is renamed ***trace-file*.0**, then ***trace-file*.1**, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum number of files, you must also specify a maximum file size with the **size** option.

Range: 2 through 1000 files

Default: 10 files

flag *flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple **flag** statements.

RIPng Tracing Options

- **error**—RIPng error packets
- **expiration**—RIPng route expiration processing
- **holddown**—RIPng hold-down processing
- **nsr-synchronization**—Nonstop routing synchronization events

Copyright © 2017, Juniper Networks, Inc.

- **packets**—All RIPng packets
- **request**—RIPng information packets such as request, poll, and poll entry packets
- **trigger**—RIPng triggered updates
- **update**—RIPng update packets

Global Tracing Options

- **all**—All tracing operations
- **general**—A combination of the **normal** and **route** trace operations
- **normal**—All normal operations

Default: If you do not specify this option, only unusual or abnormal operations are traced.

- **policy**—Policy operations and actions
- **route**—Routing table changes
- **state**—State transitions
- **task**—Routing protocol task processing
- **timer**—Routing protocol timer processing

flag-modifier—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

- **detail**—Provide detailed trace information.
- **receive**—Trace the packets being received.
- **receive-detail**—Provide detailed trace information for packets being received.
- **send**—Trace the packets being transmitted.
- **send-detail**—Provide detailed trace information for packets being transmitted.

no-world-readable—(Optional) Do not allow any user to read the log file.

size ***size***—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named ***trace-file*** reaches this size, it is renamed ***trace-file*.0**. When the ***trace-file*** again reaches its maximum size, ***trace-file*.0** is renamed ***trace-file*.1** and ***trace-file*** is renamed ***trace-file*.0**. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten. If you specify a maximum file size, you must also specify a maximum number of trace files with the **files** option.

Syntax: ***x*k** to specify KB, ***x*m** to specify MB, or ***x*g** to specify GB

Range: 10 KB through the maximum file size supported on your system

Default: 128 KB

world-readable—(Optional) Allow any user to read the log file.

Copyright © 2017, Juniper Networks, Inc.

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Tracing RIPng Protocol Traffic on page 158

update-interval (Protocols RIPng)

Syntax
update-interval *seconds*;

Hierarchy Level
[edit logical-systems *logical-system-name* protocols ripng],
[edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols ripng],
[edit protocols ripng],
[edit routing-instances *routing-instance-name* protocols ripng]

Release Information
Statement introduced in Junos OS Release 7.6.
Statement introduced in Junos OS Release 9.0 for EX Series switches.
Support for routing instances introduced in Junos OS Release 9.0.

Description
Configure the interval at which routes learned by RIPng are sent to neighbors.

Options
seconds—Estimated time to wait before making updates to the routing table.

Range: 10 through 60 seconds

Default: 30 seconds

Required Privilege Level
routing—To view this statement in the configuration.
routing-control—To add this statement to the configuration.

Related Documentation
- Example: Configuring RIP Timers on page 38

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 22

RIP Operational Commands

- clear rip general-statistics
- clear rip statistics
- restart
- show policy
- show policy conditions
- show rip general-statistics
- show rip neighbor
- show rip statistics
- show route
- show route active-path
- show route advertising-protocol
- show route all
- show route best
- show route brief
- show route detail
- show route exact
- show route export
- show route extensive
- show route forwarding-table
- show route hidden
- show route inactive-path
- show route instance
- show route next-hop
- show route output
- show route protocol
- show route receive-protocol
- show route table

Copyright © 2017, Juniper Networks, Inc.

- show route terse
- test policy

Copyright © 2017, Juniper Networks, Inc.

clear rip general-statistics

List of Syntax
Syntax on page 235
Syntax (EX Series Switches and QFX Series) on page 235

Syntax
clear rip general-statistics
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)
clear rip general-statistics

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Clear RIP general statistics.

Options
none—Clear RIP general statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- show rip general-statistics on page 253

List of Sample Output
clear rip general-statistics on page 235

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear rip general-statistics

```
user@host> clear rip general-statistics
```

Copyright © 2017, Juniper Networks, Inc.

clear rip statistics

List of Syntax
Syntax on page 236
Syntax (EX Series Switches and QFX Series) on page 236

Syntax
clear rip statistics
<instance (all | *instance-name*)>
<logical-system (all | *logical-system-name*)>
<*neighbor*>
<peer (all | *address*)>

Syntax (EX Series Switches and QFX Series)
clear rip statistics
<instance (all | *instance-name*)>
<*neighbor*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Clear RIP statistics.

Options
none—Reset RIP counters for all neighbors for all routing instances.

instance (all | *instance-name*)—(Optional) Clear RIP statistics for all instances or for the specified routing instance only.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor—(Optional) Clear RIP statistics for the specified neighbor only.

peer (all | *address*)—(Optional) Clear RIP statistics for a single peer or all peers.

Required Privilege Level
clear

Related Documentation
- show rip statistics on page 257

List of Sample Output
clear rip statistics on page 236

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear rip statistics

```
user@host> clear rip statistics
```


Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

restart

List of Syntax

Syntax on page 238
Syntax (ACX Series Routers) on page 238
Syntax (EX Series Switches) on page 238
Syntax (MX Series Routers) on page 239
Syntax (QFX Series) on page 239
Syntax (Routing Matrix) on page 239
Syntax (TX Matrix Routers) on page 239
Syntax (TX Matrix Plus Routers) on page 240
Syntax (QFX Series) on page 240

Syntax

```
restart
<adaptive-services |ancpd-service | application-identification |audit-process |
  auto-configuration |captive-portal-content-delivery |ce-l2tp-service |chassis-control |
  class-of-service |clksyncd-service |database-replication|datapath-trace-service
  |dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture |
  ecc-error-logging | ethernet-connectivity-fault-management
  |ethernet-link-fault-management |event-processing | firewall
  |general-authentication-service | gracefully | iccp-service |idp-policy | immediately
  |interface-control | ipsec-key-management | kernel-replication | l2-learning | l2cpd-service
  | l2tp-service | l2tp-universal-edge | lacp | license-service |link-management
  |local-policy-decision-function |mac-validation |mib-process | mountd-service
  |mpls-traceroute |mspd | multicast-snooping |named-service | nfsd-service |
  packet-triggered-subscribers |peer-selection-service |pgm | pic-services-logging | pki-service
  |ppp | ppp-service |pppoe | protected-system-domain-service |
  redundancy-interface-process | remote-operations | root-system-domain-service | routing
  <logical-system logical-system-name> | sampling | sbc-configuration-process | sdk-service
  |service-deployment | services | snmp |soft |static-subscribers |statistics-service|
  subscriber-management | subscriber-management-helper | tunnel-oamd |usb-control|
  vrrp |web-management>
<gracefully | immediately | soft>
```

Syntax (ACX Series Routers)

```
restart
<adaptive-services |audit-process | auto-configuration | autoinstallation |chassis-control |
  class-of-service |clksyncd-service |database-replication| dhcp-service | diameter-service
  | disk-monitoring | dynamic-flow-capture | ethernet-connectivity-fault-management
  |ethernet-link-fault-management |event-processing | firewall
  |general-authentication-service | gracefully | immediately |interface-control |
  ipsec-key-management | l2-learning | lacp |link-management |mib-process | mountd-service
  |mpls-traceroute |mspd | named-service | nfsd-service | pgm | pki-service |ppp | pppoe |
  redundancy-interface-process | remote-operations | routing | sampling | sdk-service
  |secure-neighbor-discovery | service-deployment | services | snmp |soft |statistics-service|
  subscriber-management | subscriber-management-helper | tunnel-oamd | vrrp>
```

Syntax (EX Series Switches)

```
restart
<autoinstallation | chassis-control | class-of-service | database-replication | dhcp |
  dhcp-service | diameter-service | dot1x-protocol | ethernet-link-fault-management |
  ethernet-switching | event-processing | firewall | general-authentication-service |
  interface-control | kernel-replication | l2-learning | lacp | license-service | link-management
  | lldpd-service | mib-process | mountd-service | multicast-snooping | pgm |
```


Copyright © 2017, Juniper Networks, Inc.

redundancy-interface-process | remote-operations | routing | secure-neighbor-discovery | service-deployment | sflow-service | snmp | vrrp | web-management>

Syntax (MX Series Routers)

```
restart
<adaptive-services | ancpd-service |application-identification|audit-process |
  auto-configuration |captive-portal-content-delivery |ce-l2tp-service |chassis-control |
  class-of-service |clksyncd-service |database-replication| datapath-trace-service
  |dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture |
  ecc-error-logging | ethernet-connectivity-fault-management
  |ethernet-link-fault-management |event-processing | firewall |
  general-authentication-service | gracefully |iccp-service |idp-policy |immediately
  |interface-control | ipsec-key-management | kernel-replication | l2-learning |l2cpd-service
  |l2tp-service | l2tp-universal-edge | lacp | license-service |link-management
  |local-policy-decision-function |mac-validation |mib-process | mountd-service
  |mpls-traceroute |mspd | multicast-snooping |named-service | nfsd-service |
  packet-triggered-subscribers |peer-selection-service | pgm | pic-services-logging |
  pki-service |ppp | ppp-service |pppoe | protected-system-domain-service |
  redundancy-interface-process | remote-operations |root-system-domain-service | routing
  |routing <logical-system logical-system-name>  | sampling | sbc-configuration-process |
  sdk-service |service-deployment |services | snmp |soft |static-subscribers |statistics-service|
  subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control|
  vrrp |web-management>
<all-members>
<gracefully | immediately | soft>
<local>
<member member-id>
```

Syntax (QFX Series)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services |
  diameter-service | dlsw | ethernet-connectivity | event-processing | fibre-channel | firewall
  | general-authentication-service | igmp-host-services | interface-control |
  ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
  named-service | network-access-service | nstrace-process | pgm | ppp | pppoe |
  redundancy-interface-process | remote-operations |logical-system-name>  | routing |
  sampling |secure-neighbor-discovery | service-deployment | snmp | usb-control |
  web-management>
<gracefully | immediately | soft>
```

Syntax (Routing Matrix)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | disk-monitoring |
  dynamic-flow-capture | ecc-error-logging | event-processing | firewall | interface-control
  | ipsec-key-management | kernel-replication | l2-learning | l2tp-service | lacp |
  link-management | mib-process | pgm | pic-services-logging | ppp | pppoe |
  redundancy-interface-process | remote-operations | routing <logical-system
  logical-system-name>  | sampling | service-deployment | snmp>
<all | all-lcc | lcc number>
<gracefully | immediately | soft>
```

Syntax (TX Matrix Routers)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp-service |
  diameter-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging |
  event-processing | firewall | interface-control | ipsec-key-management | kernel-replication
  | l2-learning | l2tp-service | lacp | link-management | mib-process |pgm | pic-services-logging
```

Copyright © 2017, Juniper Networks, Inc.

| ppp | pppoe | redundancy-interface-process | remote-operations | routing <logical-system *logical-system-name*> | sampling | service-deployment | snmp| statistics-service>
<all-chassis | all-lcc | lcc *number* | scc>
<gracefully | immediately | soft>

Syntax (TX Matrix Plus Routers)

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging | event-processing | firewall | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2tp-service | lacp | link-management | mib-process | pgm | pic-services-logging | ppp | pppoe | redundancy-interface-process | remote-operations | routing <logical-system *logical-system-name*> | sampling | service-deployment | snmp| statistics-service>
<all-chassis | all-lcc | all-sfc | lcc *number* | sfc *number*>
<gracefully | immediately | soft>

Syntax (QFX Series)

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services | diameter-service | dlsw | ethernet-connectivity | event-processing | fibre-channel | firewall | general-authentication-service | igmp-host-services | interface-control | ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process | named-service | network-access-service | nstrace-process | pgm | ppp | pppoe | redundancy-interface-process | remote-operations |*logical-system-name*> | routing | sampling |secure-neighbor-discovery | service-deployment | snmp | usb-control | web-management>
<gracefully | immediately | soft>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 12.2 for ACX Series routers.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Options added:

- **dynamic-flow-capture** in Junos OS Release 7.4.
- **dlsw** in Junos OS Release 7.5.
- **event-processing** in Junos OS Release 7.5.
- **ppp** in Junos OS Release 7.5.
- **l2ald** in Junos OS Release 8.0.
- **link-management** in Release 8.0.
- **pgcp-service** in Junos OS Release 8.4.
- **sbc-configuration-process** in Junos OS Release 9.5.
- **services pgcp gateway** in Junos OS Release 9.6.
- **sfc** and **all-sfc** for the TX Matrix Router in Junos OS Release 9.6.

Description

Restart a Junos OS process.

Copyright © 2017, Juniper Networks, Inc.

CAUTION: Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.

Options **none**—Same as **gracefully**.

adaptive-services—(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all T1600 routers connected to the TX Matrix Plus router.

all-members—(MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.

all-sfc—(TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).

ancpd-service—(Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.

application-identification—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.

audit-process—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.

auto-configuration—(Optional) Restart the Interface Auto-Configuration process.

autoinstallation—(EX Series switches only) (Optional) Restart the autoinstallation process.

captive-portal-content-delivery—(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.

ce-l2tp-service—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.

Copyright © 2017, Juniper Networks, Inc.

chassis-control—(Optional) Restart the chassis management process.

class-of-service—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.

clksyncd-service—(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).

database-replication—(EX Series switches and MX Series routers only) (Optional) Restart the database replication process.

datapath-trace-service—(Optional) Restart the packet path tracing process.

dhcp—(EX Series switches only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.

dhcp-service—(Optional) Restart the Dynamic Host Configuration Protocol process.

dialer-services—(EX Series switches only) (Optional) Restart the ISDN dial-out process.

diameter-service—(Optional) Restart the diameter process.

disk-monitoring—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.

dlsw—(QFX Series only) (Optional) Restart the data link switching (DLSw) service.

dot1x-protocol—(EX Series switches only) (Optional) Restart the port-based network access control process.

dynamic-flow-capture—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

ecc-error-logging—(Optional) Restart the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.

ethernet-connectivity-fault-management—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

ethernet-link-fault-management—(EX Series switches and MX Series routers only) (Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.

ethernet-switching—(EX Series switches only) (Optional) Restart the Ethernet switching process.

event-processing—(Optional) Restart the event process (eventd).

fibre-channel—(QFX Series only) (Optional) Restart the Fibre Channel process.

Copyright © 2017, Juniper Networks, Inc.

firewall—(Optional) Restart the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

general-authentication-service—(EX Series switches and MX Series routers only) (Optional) Restart the general authentication process.

gracefully—(Optional) Restart the software process.

iccp-service—(Optional) Restart the Inter-Chassis Communication Protocol (ICCP) process.

idp-policy—(Optional) Restart the intrusion detection and prevention (IDP) protocol process.

immediately—(Optional) Immediately restart the software process.

interface-control—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

ipsec-key-management—(Optional) Restart the IPsec key management process.

isdn-signaling—(QFX Series only) (Optional) Restart the ISDN signaling process, which initiates ISDN connections.

kernel-replication—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

l2-learning—(Optional) Restart the Layer 2 address flooding and learning process.

l2cpd-service—(Optional) Restart the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

l2tp-service— (M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Layer 2 Tunneling Protocol (L2TP) process, which sets up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented.

l2tp-universal-edge—(MX Series routers only) (Optional) Restart the L2TP process, which establishes L2TP tunnels and PPP sessions through L2TP tunnels.

lacp—(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific router that is connected to the TX Matrix Plus router.

Copyright © 2017, Juniper Networks, Inc.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

license-service—(EX Series switches only) (Optional) Restart the feature license management process.

link-management— (TX Matrix and TX Matrix Plus routers and EX Series switches only) (Optional) Restart the Link Management Protocol (LMP) process, which establishes and maintains LMP control channels.

lldpd-service—(EX Series switches only) (Optional) Restart the Link Layer Discovery Protocol (LLDP) process.

local—(MX Series routers only) (Optional) Restart the software process for the local Virtual Chassis member.

local-policy-decision-function— (Optional) Restart the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

mac-validation— (Optional) Restart the Media Access Control (MAC) validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

member *member-id*—(MX Series routers only) (Optional) Restart the software process for a specific member of the Virtual Chassis configuration. Replace ***member-id*** with a value of **0** or **1**.

mib-process—(Optional) Restart the Management Information Base (MIB) version II process, which provides the router's MIB II agent.

mobile-ip—(Optional) Restart the Mobile IP process, which configures Junos OS Mobile IP features.

mountd-service—(EX Series switches and MX Series routers only) (Optional) Restart the service for NFS mount requests.

mpls-traceroute—(Optional) Restart the MPLS Periodic Traceroute process.

mspd—(Optional) Restart the Multiservice process.

Copyright © 2017, Juniper Networks, Inc.

multicast-snooping—(EX Series switches and MX Series routers only) (Optional) Restart the multicast snooping process, which makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

named-service—(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

network-access-service—(QFX Series only) (Optional) Restart the network access process, which provides the router's Challenge Handshake Authentication Protocol (CHAP) authentication service.

nfsd-service—(Optional) Restart the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

packet-triggered-subscribers—(Optional) Restart the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

peer-selection-service—(Optional) Restart the Peer Selection Service process.

pgcp-service—(Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the **services pgcp gateway** option.

pgm—(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.

pic-services-logging—(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

pki-service—(Optional) Restart the PKI Service process.

ppp—(Optional) Restart the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

ppp-service—(Optional) Restart the Universal edge PPP process, which is the encapsulation protocol process for transporting IP traffic across universal edge routers.

pppoe—(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

protected-system-domain-service—(Optional) Restart the Protected System Domain (PSD) process.

redundancy-interface-process—(Optional) Restart the ASP redundancy process.

Copyright © 2017, Juniper Networks, Inc.

remote-operations—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.

root-system-domain-service—(Optional) Restart the Root System Domain (RSD) service.

routing—(ACX Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the routing protocol process.

routing <logical-system *logical-system-name*>—(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables. Optionally, restart the routing protocol process for the specified logical system only.

sampling—(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

sbc-configuration-process—(Optional) Restart the session border controller (SBC) process of the border signaling gateway (BSG).

scc—(TX Matrix routers only) (Optional) Restart the software process on the TX Matrix router (or switch-card chassis).

sdk-service—(Optional) Restart the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

secure-neighbor-discovery—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

sfc *number*—(TX Matrix Plus routers only) (Optional) Restart the software process on the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with **0**.

service-deployment—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

services—(Optional) Restart a service.

services pgcp gateway *gateway-name*—(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the **pgcp-service** option.

sflow-service—(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.

snmp—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

Copyright © 2017, Juniper Networks, Inc.

soft—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

static-subscribers—(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

statistics-service—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.

subscriber-management—(Optional) Restart the Subscriber Management process.

subscriber-management-helper—(Optional) Restart the Subscriber Management Helper process.

tunnel-oamd—(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

usb-control—(MX Series routers) (Optional) Restart the USB control process.

vrrp—(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

web-management—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

Required Privilege Level reset

Related Documentation
- *Overview of Junos OS CLI Operational Mode Commands*

List of Sample Output restart interface-control gracefully on page 247

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

restart interface-control gracefully

```
user@host> restart interface-control gracefully
Interface control process started, pid 41129
```

Copyright © 2017, Juniper Networks, Inc.

show policy

List of Syntax
Syntax on page 248
Syntax (EX Series Switches) on page 248

Syntax
show policy
<logical-system (all | *logical-system-name*)>
<*policy-name*>
<*statistics* >

Syntax (EX Series Switches)
show policy
<*policy-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
statistics option introduced in Junos OS Release 16.1 for MX Series routers.

Description
Display information about configured routing policies.

Options
none—List the names of all configured routing policies.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

policy-name—(Optional) Show the contents of the specified policy.

statistics—(Optional) Use in conjunction with the **test policy** command to show the length of time (in microseconds) required to evaluate a given policy and the number of times it has been executed. This information can be used, for example, to help structure a policy so it is evaluated efficiently. Timers shown are per route; times are not cumulative. Statistics are incremented even when the router is learning (and thus evaluating) routes from peering routers.

Required Privilege Level
view

Related Documentation
- *show policy damping*
- test policy on page 450

List of Sample Output
show policy on page 249
show policy policy-name on page 249
show policy statistics policy-name on page 249
show policy (Multicast Scoping) on page 250
show policy (Route Filter and source Address Filter Lists) on page 250

Copyright © 2017, Juniper Networks, Inc.

Output Fields Table 10 on page 249 lists the output fields for the **show policy** command. Output fields are listed in the approximate order in which they appear.

Table 10: show policy Output Fields

Field Name	Field Description
policy-name	Name of the policy listed.
term	Name of the user-defined policy term. The term name **unnamed** is used for policy elements that occur outside of user defined terms
from	Match condition for the policy.
then	Action for the policy.

Sample Output

show policy

```
user@host> show policy
Configured policies:
__vrf-export-red-internal__
__vrf-import-red-internal__
red-export
rf-test-policy
multicast-scoping
```

show policy policy-name

```
user@host> show policy vrf-import-red-internal
Policy vrf-import-red-internal:
    from
         203.0.113.0/28  accept
         203.0.113.32/28  accept
    then reject
```

show policy statistics policy-name

```
user@host> show policy statistics iBGP-v4-RR-Import
Policy iBGP-v4-RR-Import:
    [1243328] Term Lab-Infra:
        from [1243328 0]  proto BGP
             [28 0] route filter:
                 10.11.0.0/8 orlonger
                 10.13.0.0/8 orlonger
        then [28 0] accept
    [1243300] Term External:
        from [1243300 1]  proto BGP
         [1243296 0]  community Ext-Com1 [64496:1515 ]
         [1243296 0]  prefix-list-filter Customer-Routes
         [1243296 0]  aspath AS6221
             [1243296 1] route filter:
                 172.16.49.0/12 orlonger
                 172.16.50.0/12 orlonger
                 172.16.51.0/12 orlonger
```

Copyright © 2017, Juniper Networks, Inc.

```
                    172.16.52.0/12 orlonger
                    172.16.56.0/12 orlonger
                    172.16.60.0/12 orlonger
                then [1243296 2] community + Ext-Com2 [64496:2000 ] [1243296 0] accept
            [4] Term Final:
                then [4 0] reject
```

show policy (Multicast Scoping)

```
user@host> show policy multicast-scoping
Policy multicast-scoping:
    from
         multicast-scope == 8
  then
    accept
```

show policy (Route Filter and source Address Filter Lists)

```
user@host> show policy rf-test-policy
Policy rf-test-policy:
    Term term1:
        from source-address-filter-list saf-list-1
             source-address filter:
                 192.0.2.0/29 longer
                 192.0.2.64/28 exact
                 192.0.2.128/28 exact
                 192.0.2.160/28 orlonger
    Term term2:
        from route-filter-list rf-list-1
             route filter:
                 198.51.100.0/29 upto 198.51.100.0/30
                 198.51.100.8/29 upto 198.51.100.8/30 accept
    Term unnamed:
        then reject
```


Copyright © 2017, Juniper Networks, Inc.

show policy conditions

Syntax

show policy conditions
<*condition-name*>
<detail>
<dynamic>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show policy conditions
<*condition-name*>
<detail>
<dynamic>

Release Information

Command introduced in Junos OS Release 9.0.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display all the configured conditions as well as the routing tables with which the configuration manager is interacting. If the **detail** keyword is included, the output also displays dependent routes for each condition.

Options

none—Display all configured conditions and associated routing tables.

condition-name—(Optional) Display information about the specified condition only.

detail—(Optional) Display the specified level of output.

dynamic—(Optional) Display information about the conditions in the dynamic database.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show policy conditions detail on page 252

Output Fields

Table 11 on page 251 lists the output fields for the **show policy conditions** command. Output fields are listed in the approximate order in which they appear.

Table 11: show policy conditions Output Fields

Field Name	Field Description	Level of Output
Condition	Name of configured condition.	All levels
event	Condition type. If the **if-route-exists** option is configured, the event type is: **Existence of a route in a specific routing table**.	All levels
Dependent routes	List of routes dependent on the condition, along with the latest generation number.	**detail**

Copyright © 2017, Juniper Networks, Inc.

Table 11: show policy conditions Output Fields *(continued)*

Field Name	Field Description	Level of Output
Condition tables	List of routing tables associated with the condition, along with the latest generation number and number of dependencies.	All levels
If-route-exists conditions	List of conditions configured to look for a route in the specified table.	All levels

Sample Output

show policy conditions detail

```
user@host> show policy conditions detail
Configured conditions:
Condition cond1, event: Existence of a route in a specific routing table
Dependent routes:
 172.16.4.4/32, generation 3
 6.6.6.6/32, generation 3
 10.10.10.10/32, generation 3

Condition cond2, event: Existence of a route in a specific routing table
Dependent routes:
None

Condition tables:
Table inet.0, generation 4, dependencies 3, If-route-exists conditions: cond1
(static) cond2 (static)
```


Copyright © 2017, Juniper Networks, Inc.

show rip general-statistics

List of Syntax
Syntax on page 253
Syntax (EX Series Switches and QFX Series) on page 253

Syntax
show rip general-statistics
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches and QFX Series)
show rip general-statistics

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Display brief RIP statistics.

Options
none—Display brief RIP statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- clear rip general-statistics on page 235

List of Sample Output
show rip general-statistics on page 254

Output Fields
Table 12 on page 253 lists the output fields for the **show rip general-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 12: show rip general-statistics Output Fields

Field Name	Field Description
bad msgs	Number of invalid messages received.
no recv intf	Number of packets received with no matching interface.
curr memory	Amount of memory currently used by RIP.
max memory	Most memory used by RIP.

Copyright © 2017, Juniper Networks, Inc.

Sample Output

show rip general-statistics

```
user@host> show rip general-statistics
RIPv2 I/O info:
     bad msgs       :         0
     no recv intf   :         0
     curr memory    :         0
     max memory     :         0
```


Copyright © 2017, Juniper Networks, Inc.

show rip neighbor

List of Syntax
Syntax on page 255
Syntax (EX Series Switches and QFX Series) on page 255

Syntax
show rip neighbor
<instance (all | *instance-name*)>
<logical-system (all | *logical-system-name*)>
<*name*>

Syntax (EX Series Switches and QFX Series)
show rip neighbor
<instance (all | *instance-name*)>
<*name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Display information about RIP neighbors.

Options
none—Display information about all RIP neighbors for all instances.

instance (all | *instance-name*)—(Optional) Display RIP neighbor information for all instances or for only the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display detailed information about only the specified RIP neighbor.

Required Privilege Level
view

List of Sample Output
show rip neighbor on page 256
show rip neighbor (With Demand Circuits Configured) on page 256

Output Fields
Table 13 on page 256 lists the output fields for the **show rip neighbor** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 13: show rip neighbor Output Fields

Field Name	Field Description
Neighbor	Name of the RIP neighbor. NOTE: Beginning with Junos OS Release 11.1, when you configure demand circuits, the output displays a demand circuit (DC) flag next to neighbor interfaces configured for demand circuits. If you configure demand circuits at the **[edit protocols rip group *group-name* neighbor *neighbor-name*]** hierarchy level, the output shows only the neighboring interface that you specifically configured as a demand circuit. If you configure demand circuits at the **[edit protocols rip group *group-name*]** hierarchy level, all of the interfaces in the group are configured as demand circuits. Therefore, the output shows all of the interfaces in that group as demand circuits.
State	State of the connection: **Up** or **Dn** (Down).
Source Address	Address of the port on the local router.
Destination Address	Address of the port on the remote router.
Send Mode	Send options: **broadcast**, **multicast**, **none**, or **version 1**.
Receive Mode	Type of packets to accept: **both**, **none**, **version 1**, or **version 2**.
In Met	Metric added to incoming routes when advertising into RIP routes that were learned from other protocols.

Sample Output

show rip neighbor

```
user@host> show rip neighbor
                    Local  Source          Destination     Send    Receive  In
Neighbor            State  Address         Address         Mode    Mode     Met
--------            -----  -------         -----------     ----    -------  ---
ge-2/3/0.0             Up  192.168.9.105   192.168.9.107   bcast   both       1
at-5/1/1.42            Dn  (null)          (null)          mcast   v2 only    3
at-5/1/0.42            Dn  (null)          (null)          mcast   both       3
at-5/1/0.0             Up  198.51.100.0    224.0.0.9       mcast   both       3
so-0/0/0.0             Up  192.168.9.97    224.0.0.9       mcast   both       3
```

show rip neighbor (With Demand Circuits Configured)

```
user@host> show rip neighbor
                    Local  Source          Destination     Send   Receive  In
Neighbor            State  Address         Address         Mode   Mode     Met
--------            -----  -------         -----------     ----   -------  ---
so-0/1/0.0(DC)         Up  10.10.10.2      224.0.0.9       mcast  both       1
so-0/2/0.0(DC)         Up  192.0.2.2       224.0.0.9       mcast  both       1
```


Copyright © 2017, Juniper Networks, Inc.

show rip statistics

List of Syntax
Syntax on page 257
Syntax (EX Series Switches and QFX Series) on page 257

Syntax
show rip statistics
<instance (all | *instance-name*)>
<logical-system (all | *logical-system-name*)>
<*name*>
<peer (all | *address*)>

Syntax (EX Series Switches and QFX Series)
show rip statistics
<instance (all | *instance-name*)>
<*name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 12.1 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Display RIP statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.

Options
none—Display RIP statistics for all routing instances.

instance (all | *instance-name*)—(Optional) Display RIP statistics for all instances or for only the specified routing instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display detailed information about only the specified RIP neighbor.

peer (all | *address*)—(Optional) Display RIP statistics for a single peer or all peers.

Required Privilege Level
view

Related Documentation
- clear rip statistics on page 236

List of Sample Output
show rip statistics on page 259

Output Fields
Table 14 on page 258 lists the output fields for the **show rip statistics** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 14: show rip statistics Output Fields

Field Name	Field Description
RIP info	Information about RIP on the specified interface: • **port**—UDP port number used for RIP. • **update interval**—Interval between routing table updates, in seconds. • **holddown**—Hold-down interval, in seconds. • **timeout**—Timeout interval, in seconds. • **restart in progress**—Graceful restart status. Displayed when RIP is or has been in the process of graceful restart. • **restart time**—Estimated time for the graceful restart to finish, in seconds. • **restart will complete in**—Remaining time for the graceful restart to finish, in seconds. • **rts learned**—Number of routes learned through RIP. • **rts held down**—Number of routes held down by RIP. • **rqsts dropped**—Number of received request packets that were dropped. • **resps dropped**—Number of received response packets that were dropped.
logical-interface	Name of the logical interface and its statistics: • **routes learned**—Number of routes learned on the logical interface. • **routes advertised**—Number of routes advertised by the logical interface.
Counter	List of counter types: • **Updates Sent**—Number of update messages sent. • **Triggered Updates Sent**—Number of triggered update messages sent. • **Responses Sent**—Number of response messages sent. • **Bad Messages**—Number of invalid messages received. • **RIPv1 Updates Received**—Number of RIPv1 update messages received. • **RIPv1 Bad Route Entries**—Number of RIPv1 invalid route entry messages received. • **RIPv1 Updates Ignored**—Number of RIPv1 update messages ignored. • **RIPv2 Updates Received**—Number of RIPv2 update messages received. • **RIPv2 Bad Route Entries**—Number of RIPv2 invalid route entry messages received. • **RIPv2 Updates Ignored**—Number of RIPv2 update messages that were ignored. • **Authentication Failures**—Number of received update messages that failed authentication. • **RIP Requests Received**—Number of RIP request messages received. • **RIP Requests Ignored**—Number of RIP request messages ignored.
Total	Total number of packets for the selected counter.
Last 5 min	Number of packets for the selected counter in the most recent 5-minute period.
Last minute	Number of packets for the selected counter in the most recent 1-minute period.

Copyright © 2017, Juniper Networks, Inc.

Sample Output

show rip statistics

```
user@host> show rip statistics so-0/0/0.0
RIP info: port 520; update interval: 30s; holddown 180s; timeout 120s
restart in progress: restart time 60s; restart will complete in 55s
    rts learned  rts held down  rqsts dropped  resps dropped
              0              0              0              0
so-0/0/0.0:  0 routes learned; 501 routes advertised
Counter                         Total   Last 5 min  Last minute
-------                   -----------  -----------  -----------
Updates Sent                        0            0            0
Triggered Updates Sent              0            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              0            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
```

Copyright © 2017, Juniper Networks, Inc.

show route

List of Syntax

Syntax on page 260
Syntax (EX Series Switches) on page 260

Syntax

show route
<all>
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>
<private>
<te-ipv4-prefix-ip *te-ipv4-prefix-ip*>
<te-ipv4-prefix-node-ip *te-ipv4-prefix-node-ip*>
<te-ipv4-prefix-node-iso *te-ipv4-prefix-node-iso*>

Syntax (EX Series Switches)

show route
<all>
<*destination-prefix*>
<private>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Option **private** introduced in Junos OS Release 9.5.
Option **private** introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.
Option **display-client-data** introduced in Junos OS Release 16.2R1 on MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series routers.
Options **te-ipv4-prefix-ip**, **te-ipv4-prefix-node-ip**, and **te-ipv4-prefix-node-iso** introduced in Junos OS Release 17.2R1 on MX Series and PTX Series.

Description

Display the active entries in the routing tables.

Options

none—Display brief information about all active entries in the routing tables.

all—(Optional) Display information about all routing tables, including private, or internal, routing tables.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

private—(Optional) Display information only about all private, or internal, routing tables.

display-client-data —(Optional) Display client id and cookie information for routes installed by the routing protocol process client applications.

te-ipv4-prefix-ip ***te-ipv4-prefix-ip***—(Optional) Display IPv4 address of the traffic-engineering prefix, without the mask length if present in the routing table.

Copyright © 2017, Juniper Networks, Inc.

te-ipv4-prefix-node-ip ***te-ipv4-prefix-node-ip***—(Optional) Display all prefixes that have originated from the traffic-engineering node. You can filter IPv4 node addresses from the traffic-engineered routes in the **lsdist.0** table.

te-ipv4-prefix-node-iso ***te-ipv4-prefix-node-iso***—(Optional) Display all prefixes that have originated from the traffic-engineering node. You can filter IPv4 routes with the specified ISO circuit ID from the **lsdist.0** table.

Required Privilege Level view

Related Documentation

- *Understanding IS-IS Configuration*
- *Example: Configuring IS-IS*
- *Examples: Configuring Internal BGP Peering*
- *Examples: Configuring External BGP Peering*
- *Examples: Configuring OSPF Routing Policy*
- *Verifying and Managing Junos OS Enhanced Subscriber Management*

List of Sample Output

show route on page 264
show route (VPN) on page 265
show route (with Destination Prefix) on page 265
show route destination-prefix detail on page 265
show route extensive on page 265
show route extensive (ECMP) on page 266
show route extensive (Multipath Resolution) on page 266
show route (Enhanced Subscriber Management) on page 270
show route (IPv6 Flow Specification) on page 271
show route display-client-data detail on page 271
show route te-ipv4-prefix-ip on page 272
show route te-ipv4-prefix-ip extensive on page 272
show route te-ipv4-prefix-node-iso on page 275
show route te-ipv4-prefix-node-iso extensive on page 275
show route te-ipv4-prefix-node-iso detail on page 278

Output Fields Table 15 on page 261 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 15: show route Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.

Copyright © 2017, Juniper Networks, Inc.

Table 15: show route Output Fields *(continued)*

Field Name	Field Description
***number* routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly. However, if you have configured advertisement of multiple routes (with the **add-path** or **advertise-inactive** statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong. • **hidden** (routes that are not used because of a routing policy).
destination-prefix	Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as: • ***MPLS-label*** (for example, 80001). • ***interface-name*** (for example, ge-1/0/2). • ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id* :*source*** (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96): ◦ ***neighbor-address***—Address of the neighbor. ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**. ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. ◦ ***vc-id***—Virtual circuit identifier. ◦ ***source***—Source of the advertisement: **Local** or **Remote**.
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
weeks*:*days* *hours*:*minutes*:*seconds	How long the route been known (for example, **2w4d 13:11:14**, or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).
metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.

Copyright © 2017, Juniper Networks, Inc.

Table 15: show route Output Fields *(continued)*

Field Name	Field Description
localpref	Local preference value included in the route.
from	Interface from which the route was received.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
encapsulated	Extended next-hop encoding capability enabled for the specified BGP community for routing IPv4 traffic over IPv6 tunnels. When BGP receives routes without the tunnel community, IPv4-0ver IPv6 tunnels are not created and BGP routes are resolved without encapsulation.
Route Labels	Stack of labels carried in the BGP route update.
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the destination is **Discard**, traffic is dropped.

Copyright © 2017, Juniper Networks, Inc.

Table 15: show route Output Fields *(continued)*

Field Name	Field Description
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • **Weight**—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • **Balance**—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. • ***lsp-path-name***—Name of the LSP used to reach the next hop. • ***label-action***—MPLS label and operation occurring at the next hop. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label). For VPNs, expect to see multiple **push** operations, corresponding to the inner and outer labels required for VPN routes (in the case of a direct PE-to-PE connection, the VPN route would have the inner label push only).
Private unicast	(Enhanced subscriber management for MX Series routers) Indicates that an access-internal route is managed by enhanced subscriber management. By contrast, access-internal routes *not* managed by enhanced subscriber management are displayed with associated next-hop and media access control (MAC) address information.
balance	Distribution of the load based on the underlying operational interface bandwidth for equal-cost multipaths (ECMP) across the nexthop gateways in percentages.

Sample Output

show route

```
user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
                   *[MVPN/70] 19:53:41, metric2 1
                      Indirect
1:65500:1:10.0.0.40/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
                    [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
                    [BGP/170] 19:53:25, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
```


Copyright © 2017, Juniper Networks, Inc.

show route (VPN)

The following sample output shows a VPN route with composite next hops enabled. The first **Push** operation corresponds to the outer label. The second **Push** operation corresponds to the inner label.

```
user@host> show route 192.0.2.0

13979:665001.inet.0: 871 destinations, 3556 routes (871 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.0/24        [BGP/170] 00:28:32, localpref 100, from 10.9.9.160
                      AS path: 13980 ?, validation-state: unverified
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                    [BGP/170] 00:28:28, localpref 100, from 10.9.9.169
                      AS path: 13980 ?, validation-state: unverified
                    > to 10.100.0.42 via ae2.0, Push 126016, Push 300368(top)
                   #[Multipath/255] 00:28:28, metric2 102
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                      to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
```

show route (with Destination Prefix)

```
user@host> show route 172.16.0.0/12

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12      *[Static/5] 2w4d 12:54:27
                    > to 192.168.167.254 via fxp0.0
```

show route destination-prefix detail

```
user@host> show route 198.51.100.0 detail

inet.0: 15 destinations, 20 routes (15 active, 0 holddown, 0 hidden)
198.51.100.0/24 (2 entries, 2 announced)
        *BGP    Preference: 170/-101
         ...
         BGP-Static Preference: 4294967292
               Next hop type: Discard
                 Address: 0x9041ae4
                 Next-hop reference count: 2
                 State: <NoReadvrt Int Ext AlwaysFlash>
        Inactive reason: Route Preference
        Local AS:   200
        Age: 4d 1:40:40
        Validation State: unverified
        Task: RT
        Announcement bits (1): 2-BGP_RT_Background
        AS path: 4 5 6 I
```

show route extensive

```
user@host> show route extensive
v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
    *BGP   Preference: 170/-101
```

Copyright © 2017, Juniper Networks, Inc.

```
                PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group
203.0.113.1
                Next hop type: Indirect
                Address: 0x92455b8
                Next-hop reference count: 2
                Source: 10.0.0.30
                Protocol next hop: 10.0.0.40
                Indirect next hop: 2 no-forward
                State: <Active Int Ext>
                          Local AS: 64510 Peer AS: 64511
                Age: 3  Metric2: 1
               Validation State: unverified
                Task: BGP_64510.10.0.0.30+179
                Announcement bits (2): 0-PIM.v1 1-mvpn global task
                AS path: I (Originator) Cluster list:  10.0.0.30
                AS path:  Originator ID: 10.0.0.40
                Communities: target:64502:100 encapsulation:0L:14           Import
Accepted
                Localpref: 100
                Router ID: 10.0.0.30
                Primary Routing Table bgp.mvpn.0
                Indirect next hops: 1
                     Protocol next hop: 10.0.0.40 Metric: 1
                     Indirect next hop: 2 no-forward
                     Indirect path forwarding next hops: 1
                           Next hop type: Router
                           Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
                     10.0.0.40/32 Originating RIB: inet.3
                           Metric: 1                       Node path count: 1
                           Forwarding nexthops: 1
                              Nexthop: 10.0.24.4 via lt-0/3/0.24
```

show route extensive (ECMP)

```
user@host> show route extensive
*IS-IS Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 1048577
                Address: 0xXXXXXXXXXX
                Next-hop reference count: YY
                Next hop: 172.16.50.2 via ae1.0 balance 43%, selected
                Session Id: 0x141
                Next hop: 192.0.2.2 via ae0.0 balance 57%
```

show route extensive (Multipath Resolution)

```
user@host> show route extensive
inet.0: 37 destinations, 37 routes (36 active, 0 holddown, 1 hidden)
10.1.1.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.1.1.2/32 -> {indirect(1048574)}
        *Static Preference: 5
                Next hop type: Indirect, Next hop index: 0
                Address: 0xb39d1b0
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 581
                Next hop: 12.1.1.2 via ge-2/0/1.0, selected
                Session Id: 0x144
                Next hop: 13.1.1.2 via ge-2/0/2.0, selected
                Session Id: 0x145
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Protocol next hop: 10.1.1.1
                    Indirect next hop: 0xb2b20f0 1048574 INH Session ID: 0x143
                    State: <Active Int Ext>
                    Age: 2:53  Metric2: 0
                    Validation State: unverified
                    Task: RT
                    Announcement bits (2): 0-KRT 2-Resolve tree 1
                    AS path: I
                    Indirect next hops: 1
                            Protocol next hop: 10.1.1.1
                           Indirect next hop: 0xb2b20f0 1048574 INH Session ID: 0x143

                            Indirect path forwarding next hops: 2
                                    Next hop type: Router
                                    Next hop: 12.1.1.2 via ge-2/0/1.0
                                    Session Id: 0x144
                                    Next hop: 13.1.1.2 via ge-2/0/2.0
                                    Session Id: 0x145
                         10.1.1.1/32 Originating RIB: inet.0
                            Node path count: 1
                              Node flags: 1
                            Forwarding nexthops: 2 (Merged)
                            Nexthop: 12.1.1.2 via ge-2/0/1.0
                             Nexthop: 13.1.1.2 via ge-2/0/2.0
user@host> show route active-path extensive
user@host> show route 141.1.1.1 active-path extensive

inet.0: 1000061 destinations, 1000082 routes (1000061 active, 0 holddown, 0 hidden)
141.1.1.1/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 141.1.1.1/32 -> {indirect(1051215)}
unicast reverse-path: 0
[ae0.0 ae1.0]
Page 0 idx 0, (group Internet-IPv4 type External) Type 1 val 0xbb2e53d8 (adv_entry)
Advertised metrics:
Nexthop: Self
AS path: [500] 410 I
Communities:
Path 141.1.1.1 from 50.0.0.11 Vector len 4. Val: 0
*BGP Preference: 170/-101
Next hop type: Indirect, Next hop index: 0
Address: 0x2e9aacdc
Next-hop reference count: 500000
Source: 50.0.0.11
Next hop type: Router, Next hop index: 0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 25, Push 20(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 20: None;
Label element ptr: 0xb5dc1780
Label parent element ptr: 0x18d48080
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 25, Push 22(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 22: None;
Label element ptr: 0xb5dc1700
Label parent element ptr: 0x18d41000
```

Copyright © 2017, Juniper Networks, Inc.

```
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 24, Push 48(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 48: None;
Label element ptr: 0x18d40800
Label parent element ptr: 0x18d49780
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 24, Push 49(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 49: None;
Label element ptr: 0xb5dc1680
Label parent element ptr: 0x18d48f00
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Label operation: Push 3851, Push 25, Push 21(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 21: None;
Label element ptr: 0xb5dc1600
Label parent element ptr: 0x18d44d80
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Label operation: Push 3851, Push 25, Push 25(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 25: None;
Label element ptr: 0xb5dc1580
Label parent element ptr: 0x18d3da80
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1, selected
Label operation: Push 3851, Push 24, Push 68(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 68: None;
Label element ptr: 0x18d41500
Label parent element ptr: 0x18d49000
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Label operation: Push 3851, Push 24, Push 69(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 69: None;
Label element ptr: 0xb5dc1500
Label parent element ptr: 0x18d48300
Label element references: 2
```


Copyright © 2017, Juniper Networks, Inc.

```
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Protocol next hop: 50.0.0.11
Label operation: Push 3851
Label TTL action: prop-ttl
Load balance label: Label 3851: None;
Indirect next hop: 0x1883e200 1051215 INH Session ID: 0xb0d
State:
Local AS: 500 Peer AS: 500
Age: 1:40:03 Metric2: 2
Validation State: unverified
Task: BGP_500.50.0.0.11
Announcement bits (5): 0-KRT 8-KRT 9-BGP_RT_Background 10-Resolve tree 5 11-Resolve
 tree 8
AS path: 410 I
Accepted
Route Label: 3851
Localpref: 100
Router ID: 50.0.0.11
Indirect next hops: 1
Protocol next hop: 50.0.0.11 Metric: 2
Label operation: Push 3851
Label TTL action: prop-ttl
Load balance label: Label 3851: None;
Indirect next hop: 0x1883e200 1051215 INH Session ID: 0xb0d
Indirect path forwarding next hops (Merged): 8
Next hop type: Router
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
50.0.0.11/32 Originating RIB: inet.3
Metric: 1 Node path count: 4
Node flags: 1
Indirect nexthops: 4
Protocol Nexthop: 50.0.0.4 Metric: 1 Push 24
Indirect nexthop: 0x1880f200 1048597 INH Session ID: 0xb0c
Path forwarding nexthops link: 0x36120400
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
50.0.0.4/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
```

Copyright © 2017, Juniper Networks, Inc.

```
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
Protocol Nexthop: 50.0.0.5 Metric: 1 Push 24
Indirect nexthop: 0x18810000 1048596 INH Session ID: 0xb0b
Path forwarding nexthops link: 0x1545be00
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
50.0.0.5/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
Protocol Nexthop: 50.0.0.6 Metric: 1 Push 25
Indirect nexthop: 0x1880e600 1048588 INH Session ID: 0xb0a
Path forwarding nexthops link: 0x3611f440
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
50.0.0.6/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
Protocol Nexthop: 50.0.0.7 Metric: 1 Push 25
Indirect nexthop: 0x1880dc00 1048586 INH Session ID: 0xb09
Path forwarding nexthops link: 0x15466d80
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
50.0.0.7/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
```

show route (Enhanced Subscriber Management)

```
user@host> show route
inet.0: 41 destinations, 41 routes (40 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

198.51.100.11/24    *[Access-internal/12] 00:00:08
                     > to #0 10.0.0.1.93.65 via demux0.1073741824
```


Copyright © 2017, Juniper Networks, Inc.

```
198.51.100.12/24    *[Access-internal/12] 00:00:08
                    Private unicast
```

show route (IPv6 Flow Specification)

```
user@host> show route
inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::10:255:185:19/128
                   *[Direct/0] 05:11:27
                    > via lo0.0
2001:db8::11:11:11:0/120
                   *[BGP/170] 00:28:58, localpref 100
                      AS path: 2000 I, validation-state: unverified
                    > to 2001:db8::13:14:2:2 via ge-1/1/4.0
2001:db8::13:14:2:0/120*[Direct/0] 00:45:07
                    > via ge-1/1/4.0
2001:db8::13:14:2:1/128*[Local/0] 00:45:18
                      Local via ge-1/1/4.0
fe80::2a0:a50f:fc71:71d5/128
                   *[Direct/0] 05:11:27
                    > via lo0.0
fe80::5e5e:abff:feb0:933e/128
                   *[Local/0] 00:45:18
                      Local via ge-1/1/4.0

inet6flow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::11:11:11:10/128,*,proto=6,dstport=80,srcport=65535/term:1
                   *[BGP/170] 00:28:58, localpref 100, from 2001:db8::13:14:2:2
                      AS path: 2000 I, validation-state: unverified
                      Fictitious
2001:db8::11:11:11:30/128,*,icmp6-type=128,len=100,dscp=10/term:2
                   *[BGP/170] 00:20:54, localpref 100, from 2001:db8::13:14:2:2
                      AS path: 2000 I, validation-state: unverified
                      Fictitious
```

show route display-client-data detail

```
user@host> show route 198.51.100.0/24 display-client-data detail
inet.0: 59 destinations, 70 routes (59 active, 0 holddown, 0 hidden)
198.51.100.0/24 (1 entry, 1 announced)
        State: <FlashAll>
        *BGP-Static Preference: 5/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xa5c2af8
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 1641
                Next hop: 192.0.2.1 via ge-2/1/1.0, selected
                Session Id: 0x160
                Protocol next hop: 192.0.2.1
                Indirect next hop: 0xa732cb0 1048621 INH Session ID: 0x17e
                State: <Active Int Ext AlwaysFlash NSR-incapable Programmed>
                Age: 3:13       Metric2: 0
                Validation State: unverified
                Announcement bits (3): 0-KRT 5-LDP 6-Resolve tree 3
```

Copyright © 2017, Juniper Networks, Inc.

```
                    AS path: I
                    Client id: 1, Cookie: 1
```

show route te-ipv4-prefix-ip

```
user@host> show route te-ipv4-prefix-ip 10.10.10.10
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152
                    *[IS-IS/15] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0101.0101.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0202.0202.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0303.0303.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0404.0404.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0505.0505.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0606.0606.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0707.0707.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
```

show route te-ipv4-prefix-ip extensive

```
user@host>show route te-ipv4-prefix-ip 10.10.10.10 extensive
lsdist.0: 298 destinations, 298 routes (298 active, 0 holddown, 0 hidden)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State:<Active NotInstall>
                Local AS:   100
                Age: 7:58
```


Copyright © 2017, Juniper Networks, Inc.

```
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0101.0101.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
       State: <Active NotInstall
       Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0>

PREFIX { Node { AS:100 ISO:0100.0202.0202.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
       Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0303.0303.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0404.0404.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
```

Copyright © 2017, Juniper Networks, Inc.

```
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0505.0505.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0606.0606.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0707.0707.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Address: 0xa1a2ac4
                    Next-hop reference count: 298
                    Next hop:
                    State: <Active NotInstall>
                    Local AS:   100
                    Age: 7:58
                    Validation State: unverified
                    Task: IS-IS
                    AS path: I
                    Prefix SID: 1000, Flags: 0x40, Algo: 0
```

show route te-ipv4-prefix-node-iso

```
user@host> show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152
                   *[IS-IS/15] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:1.1.1.1/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:2.2.2.2/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:3.3.3.3/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:4.4.4.4/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:5.5.5.5/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:6.6.6.6/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:7.7.7.7/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
```

show route te-ipv4-prefix-node-iso extensive

```
user@host> show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00 extensive
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
```

Copyright © 2017, Juniper Networks, Inc.

```
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:1.1.1.1/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1001, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:2.2.2.2/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1002, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:3.3.3.3/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
```


Copyright © 2017, Juniper Networks, Inc.

```
                Prefix SID: 1003, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:4.4.4.4/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1004, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:5.5.5.5/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1005, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:6.6.6.6/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1006, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:7.7.7.7/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
```

Copyright © 2017, Juniper Networks, Inc.

```
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1007, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0
```

show route te-ipv4-prefix-node-iso detail

```
user@host> show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00 detail
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:1.1.1.1/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1001, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:2.2.2.2/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
```


Copyright © 2017, Juniper Networks, Inc.

```
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1002, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:3.3.3.3/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1003, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:4.4.4.4/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1004, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:5.5.5.5/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1005, Flags: 0xe0, Algo: 0
```

Copyright © 2017, Juniper Networks, Inc.

```
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:6.6.6.6/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1006, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:7.7.7.7/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1007, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0
```


Copyright © 2017, Juniper Networks, Inc.

show route active-path

List of Syntax
Syntax on page 281
Syntax (EX Series Switches) on page 281

Syntax
show route active-path
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route active-path
<brief | detail | extensive | terse>

Release Information
Command introduced in Junos OS Release 8.0.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.

Options
none—Display all active routes.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route active-path on page 281
show route active-path brief on page 282
show route active-path detail on page 282
show route active-path extensive on page 283
show route active-path terse on page 285

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route active-path

```
user@host> show route active-path

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    *[Direct/0] 21:33:52
```

```
                    > via lo0.0
10.255.71.50/32    *[IS-IS/15] 00:18:13, metric 10
                    > to 172.16.100.1 via so-2/1/3.0
172.16.100.1/24      *[Direct/0] 00:18:36
                    > via so-2/1/3.0
172.16.100.1/32      *[Local/0] 00:18:41
                      Local via so-2/1/3.0
192.168.64.0/21    *[Direct/0] 21:33:52
                    > via fxp0.0
192.168.70.19/32   *[Local/0] 21:33:52
                      Local via fxp0.0
```

show route active-path brief

The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see show route active-path on page 281.

show route active-path detail

```
user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via lo0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

10.255.71.50/32 (1 entry, 1 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 397
                Next-hop reference count: 4
                Next hop: 172.16.100.1 via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:31      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
                AS path: I

172.16.100.0/24 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:54
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
```

 Copyright © 2017, Juniper Networks, Inc.

```
                AS path: I

172.16.100.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: so-2/1/3.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 21:59
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.64.0/21 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via fxp0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.70.19/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: fxp0.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I
```

show route active-path extensive

```
user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via lo0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

10.255.71.50/32 (1 entry, 1 announced)
```

```
TSI:
KRT in-kernel 10.255.71.50/32 -> {172.16.100.1}
IS-IS level 2, LSP fragment 0
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 397
                Next-hop reference count: 4
                Next hop: 172.16.100.1 via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 24:08      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
                AS path: I

172.16.100.1/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 24:31
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

172.16.100.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: so-2/1/3.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 24:36
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.64.0/21 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via fxp0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.70.19/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: fxp0.0
```


Copyright © 2017, Juniper Networks, Inc.

```
State: ‹Active NoReadvrt Int›
Local AS:   200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I
```

show route active-path terse

```
user@host> show route active-path terse

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.19/32    D   0                       >lo0.0
* 10.255.71.50/32    I  15         10            >172.16.100.1.
* 172.16.100.0/24       D   0                       >so-2/1/3.0
* 172.16.100.2/32       L   0                        Local
* 192.168.64.0/21    D   0                       >fxp0.0
* 192.168.70.19/32   L   0                        Local
```

Copyright © 2017, Juniper Networks, Inc.

show route advertising-protocol

Syntax show route advertising-protocol *protocol neighbor-address*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Release Information Command introduced before Junos OS Release 7.4.

Description Display the routing information as it has been prepared for advertisement to a particular neighbor of a particular dynamic routing protocol.

Options **brief | detail | extensive | terse**—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor-address—Address of the neighboring router to which the route entry is being transmitted.

protocol—Protocol transmitting the route:

- **bgp**—Border Gateway Protocol
- **dvmrp**—Distance Vector Multicast Routing Protocol
- **msdp**—Multicast Source Discovery Protocol
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation

Additional Information Routes displayed are routes that the routing table has exported into the routing protocol and that have been filtered by the associated protocol's **export** routing policy statements. Starting with Junos OS Release 13.3, you can display the routing instance table **foo** for any address family, on a VPN route reflector, or a VPN AS boundary router that is advertising local VPN routes. However, If you do not specify the **table** in the command, the output displays each VRF prefix twice.

Required Privilege Level view

Related Documentation
- *Example: Configuring the MED Attribute That Determines the Exit Point in an AS*

List of Sample Output
show route advertising-protocol bgp (Layer 3 VPN) on page 289
show route advertising-protocol bgp detail on page 289

Copyright © 2017, Juniper Networks, Inc.

show route advertising-protocol bgp detail (Aggregate Extended Community Bandwidth) on page 289
show route advertising-protocol bgp detail (Labeled Unicast) on page 290
show route advertising-protocol bgp detail (Layer 2 VPN) on page 290
show route advertising-protocol bgp detail (Layer 3 VPN) on page 290
show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address) on page 291

Output Fields Table 16 on page 287 lists the output fields for the **show route advertising-protocol** command. Output fields are listed in the approximate order in which they appear.

Table 16: show route advertising-protocol Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example, inet.0.	All levels
number **destinations**	Number of destinations for which there are routes in the routing table.	All levels
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)	All levels
Prefix	Destination prefix.	**brief** none
destination-prefix **(entry , announced)**	Destination prefix. The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination.	**detail extensive**
BGP group and type	BGP group name and type (**Internal** or **External**).	**detail extensive**
Route Distinguisher	Unique 64-bit prefix augmenting each IP subnet.	**detail extensive**
Advertised Label	Incoming label advertised by the Label Distribution Protocol (LDP). When an IP packet enters a label-switched path (LSP), the ingress router examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label.	**detail extensive**
Label-Base, range	First label in a block of labels and label block size. A remote PE router uses this first label when sending traffic toward the advertising PE router.	**detail extensive**
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routers by advertising VPN labels. VPN labels transit over either a Resource Reservation Protocol (RSVP) or a Label Distribution Protocol (LDP) label-switched path (LSP) tunnel.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 16: show route advertising-protocol Output Fields *(continued)*

Field Name	Field Description	Level of Output
Nexthop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the next-hop advertisement to the peer is **Self**, and the RIB-out next hop is a specific IP address, the RIB-out IP address is included in the extensive output. See show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address) on page 291.	All levels
MED	Multiple exit discriminator value included in the route.	**brief**
Lclpref or **Localpref**	Local preference value included in the route.	All levels
Queued	When BGP route prioritization is enabled and a route is present in a priority queue, this shows which priority queue the route is in.	All levels except **brief**
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if configured on the router, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.	All levels
Route Labels	Stack of labels carried in the BGP route update.	**detail extensive**
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	**detail extensive**
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	**detail extensive**
Communities	Community path attribute for the route. See the output field table for the show route detail command for all possible values for this field.	**detail extensive**
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 16: show route advertising-protocol Output Fields *(continued)*

Field Name	Field Description	Level of Output
Attrset AS	Number, local preference, and path of the autonomous system (AS) that originated the route. These values are stored in the **Attrset** attribute at the originating router.	**detail extensive**
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	**detail extensive**
control flags	Control flags: **none** or **Site Down**.	**detail extensive**
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	**detail extensive**

Sample Output

show route advertising-protocol bgp (Layer 3 VPN)

```
user@host> show route advertising-protocol bgp 10.255.14.171
      VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.172/32   Self                     1        100 I
      VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.181/32   Self                     2        100 I
```

show route advertising-protocol bgp detail

```
user@host> show route advertising-protocol bgp 111.222.1.3 detail
bgp20.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
111.222.1.11/32 (1 entry, 1 announced)
 BGP group pe-pe type Internal
     Route Distinguisher: 111.255.14.11:69
     Advertised Label: 100000
     next hop: Self
     Localpref: 100
     AS path: 2 I
     Communities: target:69:20
     AIGP 210
111.8.0.0/16 (1 entry, 1 announced)
 BGP group pe-pe type Internal
     Route Distinguisher: 111.255.14.11:69
     Advertised Label: 100000
     Next hop: Self
     Localpref: 100
     AS path: 2 I
     Communities: target:69:20
     AIGP 210
```

show route advertising-protocol bgp detail (Aggregate Extended Community Bandwidth)

```
user@host> show route advertising-protocol bgp 10.0.4.2 10.0.2.0/30 detail
inet.0: 20 destinations, 26 routes (20 active, 0 holddown, 0 hidden)
* 10.0.2.0/30 (2 entries, 1 announced)
 BGP group external2 type External
     Nexthop: Self
```

Copyright © 2017, Juniper Networks, Inc.

```
    AS path: [65000] 65001 I
    Communities: bandwidth:65000:80000000
```

show route advertising-protocol bgp detail (Labeled Unicast)

```
user@host>show route advertising bgp 1.1.1.3 detail
inet.0: 69 destinations, 70 routes (69 active, 0 holddown, 0 hidden)
* 1.1.1.8/32 (2 entries, 2 announced)
BGP group ibgp type Internal
Route Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: 1.1.1.4
MED: 7
Localpref: 100
AS path: [5] I
Cluster ID: 3.3.3.3
Originator ID: 1.1.1.1
Entropy label capable
inet6.0: 26 destinations, 28 routes (26 active, 0 holddown, 0 hidden)
* 100::1/128 (2 entries, 1 announced)
BGP group ibgp type Internal
Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: ::ffff:1.1.1.4
Localpref: 100
AS path: [5] I
Cluster ID: 3.3.3.3
Originator ID: 1.1.1.1
```

show route advertising-protocol bgp detail (Layer 2 VPN)

```
user@host> show route advertising-protocol bgp 192.168.24.1 detail
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
192.168.16.1:1:1:1/96 (1 entry, 1 announced)
 BGP group int type Internal
     Route Distinguisher: 192.168.16.1:1
     Label-base : 32768, range : 3
     Nexthop: Self
     Localpref: 100
     AS path: I
     Communities: target:65412:100
     AIGP 210
     Layer2-info: encaps:VLAN, control flags:, mtu:
```

show route advertising-protocol bgp detail (Layer 3 VPN)

```
user@host>  show route advertising-protocol bgp 10.255.14.176 detail
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 10.255.14.174:2
     VPN Label: 101264
     Nexthop: Self
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AIGP 210
     AttrSet AS: 100
         Localpref: 100
         AS path: I
...
```


Copyright © 2017, Juniper Networks, Inc.

show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address)

```
user@host> show route advertising-protocol bgp 200.0.0.2 170.0.1.0/24 extensive all
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 6 hidden)
  170.0.1.0/24 (2 entries, 1 announced)
 BGP group eBGP-INTEROP type External
     Nexthop: Self (rib-out 10.100.3.2)
     AS path: [4713] 200 I
...
```

Copyright © 2017, Juniper Networks, Inc.

show route all

List of Syntax
Syntax on page 292
Syntax (EX Series Switches) on page 292

Syntax
show route all
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route all

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display information about all routes in all routing tables, including private, or internal, tables.

Options
none—Display information about all routes in all routing tables, including private, or internal, tables.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show route brief on page 297
- show route detail on page 299

List of Sample Output
show route all on page 292

Output Fields
In Junos OS Release 9.5 and later, only the output fields for the **show route all** command display all routing tables, including private, or hidden, routing tables. The output field table of the show route command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later.

Sample Output

show route all

The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```
user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:24:39, metric 1
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Receive
1                  *[MPLS/0] 2d 02:24:39, metric 1
                    Receive
2                  *[MPLS/0] 2d 02:24:39, metric 1
                    Receive
800017             *[VPLS/7] 1d 14:00:16
                     > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 14:00:26
                     > via vt-3/2/0.32772, Pop

user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:19:12, metric 1
                    Receive
1                  *[MPLS/0] 2d 02:19:12, metric 1
                    Receive
2                  *[MPLS/0] 2d 02:19:12, metric 1
                    Receive
800017             *[VPLS/7] 1d 13:54:49
                     > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 13:54:59
                     > via vt-3/2/0.32772, Pop
vt-3/2/0.32769      [VPLS/7] 1d 13:54:49
                      Unusable
vt-3/2/0.32772      [VPLS/7] 1d 13:54:59
                      Unusable
```

Copyright © 2017, Juniper Networks, Inc.

show route best

List of Syntax
Syntax on page 294
Syntax (EX Series Switches) on page 294

Syntax
show route best *destination-prefix*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route best *destination-prefix*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

destination-prefix—Address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show route brief on page 297
- show route detail on page 299

List of Sample Output
show route best on page 294
show route best detail on page 295
show route best extensive on page 296
show route best terse on page 296

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route best

```
user@host> show route best 10.255.70.103
```


Copyright © 2017, Juniper Networks, Inc.

```
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[OSPF/10] 1d 13:19:20, metric 2
                    > to 10.31.1.6 via ge-3/1/0.0
                      via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[RSVP/7] 1d 13:20:13, metric 2
                    > via so-0/3/0.0, label-switched-path green-r1-r3

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.0/8         *[Direct/0] 2d 01:43:34
                    > via fxp2.0
                    [Direct/0] 2d 01:43:34
                    > via fxp1.0
```

show route best detail

```
user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                Next hop: via so-0/3/0.0
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:06        Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 5
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100016
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:59        Metric: 2
                Task: RSVP
                Announcement bits (1): 1-Resolve tree 2
                AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp2.0, selected
                State: <Active Int>
```

Copyright © 2017, Juniper Networks, Inc.

```
                Age: 2d 1:44:20
                Task: IF
                AS path: I
        Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp1.0, selected
                State: <NotBest Int>
                Inactive reason: No difference
                Age: 2d 1:44:20
                Task: IF
                AS path: I
```

show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see show route best detail on page 295.

show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   O  10          2            >10.31.1.6
                                                   so-0/3/0.0

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

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   R   7          2            >so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.0.0.0/8         D   0                        >fxp2.0
                     D   0                        >fxp1.0
```


Copyright © 2017, Juniper Networks, Inc.

show route brief

List of Syntax
Syntax on page 297
Syntax (EX Series Switches) on page 297

Syntax
show route brief
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route brief
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display brief information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show route all on page 292
- show route best on page 294

List of Sample Output
show route brief on page 297

Output Fields
For information about output fields, see the Output Field table of the show route command.

Sample Output

show route brief

```
user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                      Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                    > via lo0.0
```

Copyright © 2017, Juniper Networks, Inc.

```
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.0.0/18      *[Static/5] 1w5d 20:30:29
                    > to 192.168.167.254 via fxp0.0
192.168.40.0/22    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.64.0/18    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.164.0/22   *[Direct/0] 2w4d 13:11:14
                    > via fxp0.0
192.168.164.51/32  *[Local/0] 2w4d 13:11:14
                      Local via fxp0.0
207.17.136.192/32  *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16     *[Direct/0] 1w5d 20:30:28
                    > via fe-0/0/3.0
100.101.2.3/32     *[Local/0] 1w5d 20:30:28
                      Local via fe-0/0/3.0
172.16.233.5/32       *[OSPF/10] 1w5d 20:30:29, metric 1
                      MultiRecv
```


Copyright © 2017, Juniper Networks, Inc.

show route detail

List of Syntax
Syntax on page 299
Syntax (EX Series Switches) on page 299

Syntax
show route detail
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route detail
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Comand introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Display detailed information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table on all systems.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route detail on page 310
show route detail (with BGP Multipath) on page 316
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 317
show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 317

Output Fields
Table 17 on page 299 describes the output fields for the **show route detail** command. Output fields are listed in the approximate order in which they appear.

Table 17: show route detail Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.

Copyright © 2017, Juniper Networks, Inc.

Table 17: show route detail Output Fields *(continued)*

Field Name	Field Description
***number* routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)

***route-destination* (entry, announced)**	Route destination (for example:10.0.0.1/24). The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.

label stacking	**(**Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).

Copyright © 2017, Juniper Networks, Inc.

Table 17: show route detail Output Fields *(continued)*

Field Name	Field Description
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value. **Preference2** values are signed integers, that is, **Preference2** values can be either positive or negative values. However, Junos OS evaluates **Preference2** values as unsigned integers that are represented by positive values. Based on the **Preference2** values, Junos OS evaluates a preferred route differently in the following scenarios: • **Both Signed Preference2 values** ◦ Route A = -101 ◦ Route B = -156 Where both the **Preference2** values are signed, Junos OS evaluates only the unsigned value of **Preference2** and Route A, which has a lower **Preference2** value is preferred. • **Unsigned Preference2 values** Now consider both unsigned **Preference2** values: ◦ Route A = 4294967096 ◦ Route B = 200 Here, Junos OS considers the lesser **Preference2** value and Route B with a **Preference2** value of 200 is preferred because it is less than 4294967096. • **Combination of signed and unsigned Preference2 values** When **Preference2** values of two routes are compared, and for one route the **Preference2** is a signed value, and for the other route it is an unsigned value, Junos OS prefers the route with the positive **Preference2** value over the negative **Preference2** value. For example, consider the following signed and unsigned **Preference2** values: ◦ Route A = -200 ◦ Route B = 200 In this case, Route B with a **Preference2** value of 200 is preferred although this value is greater than -200, because Junos OS evaluates only the unsigned value of the **Preference2** value.
Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).

Copyright © 2017, Juniper Networks, Inc.

Table 17: show route detail Output Fields *(continued)*

Field Name	Field Description
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 18 on page 305.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 19 on page 307.
Local AS	AS number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.

Copyright © 2017, Juniper Networks, Inc.

Table 17: show route detail Output Fields *(continued)*

Field Name	Field Description
Metric*n*	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see show route table.
Task	Name of the protocol that has added the route.
Announcement bits	The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the KRT for installing the route into the Packet Forwarding Engine, to a resolve tree, a L2 VC, or even a VPN. For example, ***n*-Resolve inet** indicates that the specified route is used for route resolution for next hops found in the routing table. • ***n***—An index used by Juniper Networks customer support only.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.

Copyright © 2017, Juniper Networks, Inc.

Table 17: show route detail Output Fields *(continued)*

Field Name	Field Description
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
ORR Generation-ID	Displays the optimal route reflection (ORR) generation identifier. ISIS and OSPF interior gateway protocol (IGP) updates filed whenever any of the corresponding ORR route has its metric valued changed, or if the ORR route is added or deleted.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 20 on page 309 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or **Site Down**.
mtu	Maximum transmission unit (MTU) information.

Copyright © 2017, Juniper Networks, Inc.

Table 17: show route detail Output Fields *(continued)*

Field Name	Field Description
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 18 on page 305 describes all possible values for the Next-hop Types output field.

Table 18: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Dynamic List	Dynamic list next hop

Copyright © 2017, Juniper Networks, Inc.

Table 18: Next-hop Types Output Field Values *(continued)*

Next-Hop Type	Description
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.
Router	A specific node or set of nodes to which the routing device forwards packets that match the route prefix. To qualify as next-hop type router, the route must meet the following criteria: • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.

Copyright © 2017, Juniper Networks, Inc.

Table 18: Next-hop Types Output Field Values *(continued)*

Next-Hop Type	Description
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 19 on page 307 describes all possible values for the State output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).

Table 19: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.

Copyright © 2017, Juniper Networks, Inc.

Table 19: State Output Field Values *(continued)*

Value	Description
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Programmed	Route installed programatically by on-box or off-box applications using API.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available
Router ID	Path through a neighbor with lower ID is available.

Copyright © 2017, Juniper Networks, Inc.

Table 19: State Output Field Values *(continued)*

Value	Description
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.
ProtectionCand	Indicates paths requesting protection.
ProtectionPath	Indicates the route entry that can be used as a protection path.

Table 20 on page 309 describes the possible values for the Communities output field.

Table 20: Communities Output Field Values

Value	Description
area-number	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is **0**. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: ***local AS number:link-bandwidth-number***	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
link-bandwidth-number	Link-bandwidth number: from **0** through **4,294,967,295** (bytes per second).
local AS number	Local AS number: from **1** through **65,535**.
options	1 byte. Currently this is only used if the route type is **5** or **7**. Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	1 byte, encoded as **1** or **2** for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); **3** for summary routes; **5** for external routes (area number must be**0**); **7** for NSSA routes; or **129** for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x8000**. The format is ***area-number:ospf-route-type:options***.

Copyright © 2017, Juniper Networks, Inc.

Table 20: Communities Output Field Values *(continued)*

Value	Description
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x0306**. The format is ***area-number:ospf-route-type:options***.
target	Defines which VPN the route participates in; **target** has the format ***32-bit IP address:16-bit number***. For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between **0x1** and **0x7fff**. This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above **0x8000**. This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route detail

```
user@host> show route detail

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:31:43
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

10.31.1.0/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:30:17
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    69
                Age: 1:30:17    Metric: 1
                ORR Generation-ID: 1
      Area: 0.0.0.0
                Task: OSPF
                AS path: I
```


Copyright © 2017, Juniper Networks, Inc.

```
10.31.1.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:30:20
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:29:56    Metric: 2
                Area: 0.0.0.0
                ORR Generation-ID: 1
     Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

172.16.233.2/32 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

172.16.233.22/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:43
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
```

Copyright © 2017, Juniper Networks, Inc.

```
                State: <Active Int>
                Local AS:    69
                Age: 1:25:49    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

10.255.71.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    69
                Age: 1:25:49    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0   (1 entry, 1 announced)
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    69
                Age: 1:31:45    Metric: 1
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

  299840 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 299840 /52 -> {indirect(1048575)}
          *RSVP   Preference: 7/2
                  Next hop type: Flood
                  Address: 0x9174a30
                  Next-hop reference count: 4
                  Next hop type: Router, Next hop index: 798
                  Address: 0x9174c28
                  Next-hop reference count: 2
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Next hop: 172.16.0.2 via lt-1/2/0.9 weight 0x1
                    Label-switched-path R2-to-R4-2p2mp
                    Label operation: Pop
                    Next hop type: Router, Next hop index: 1048574
                    Address: 0x92544f0
                    Next-hop reference count: 2
                    Next hop: 172.16.0.2 via lt-1/2/0.7 weight 0x1
                    Label-switched-path R2-to-R200-p2mp
                    Label operation: Pop
                    Next hop: 172.16.0.2 via lt-1/2/0.5 weight 0x8001
                    Label operation: Pop
                    State: <Active Int>
                    Age: 1:29       Metric: 1
                    Task: RSVP
                    Announcement bits (1): 0-KRT
                    AS path: I...

800010 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:29:30
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 10.255.70.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:29:30    Metric2: 2
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
```

Copyright © 2017, Juniper Networks, Inc.

```
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:43
                Task: MLD
                Announcement bits (1): 0-KRT
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:31:44
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
```


Copyright © 2017, Juniper Networks, Inc.

```
                Age: 1:25:49    Metric2: 1
                AIGP 210
                Task: BGP_69.10.255.70.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 10.255.70.103
                Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:31:40    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:31:40    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    99
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512

inet.0: 45 destinations, 47 routes (44 active, 0 holddown, 1 hidden)
1.1.1.3/32 (1 entry, 1 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Router, Next hop index: 580
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Address: 0x9db6ed0
                    Next-hop reference count: 8
                    Next hop: 10.1.1.6 via lt-1/0/10.5, selected
                    Session Id: 0x18a
                    State: <Active Int>
                    Local AS:     2
                    Age: 1:32       Metric: 10
                    Validation State: unverified
                    ORR Generation-ID: 1
                    Task: IS-IS
                    Announcement bits (3): 0-KRT 5-Resolve tree 4 6-Resolve_IGP_FRR
task
                    AS path: I

inet.0: 61 destinations, 77 routes (61 active, 1 holddown, 0 hidden)
1.1.1.1/32 (2 entries, 1 announced)
        *OSPF   Preference: 10
                Next hop type: Router, Next hop index: 673
                Address: 0xc008830
                Next-hop reference count: 3
                Next hop: 10.1.1.1 via ge-0/0/2.0, selected
                Session Id: 0x1b7
                State: <Active Int>
                Local AS:     1
                Age: 3:06:59    Metric: 100
                Validation State: unverified
                ORR Generation-ID: 1
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 1-KRT 9-Resolve tree 2
                AS path: I
```

show route detail (with BGP Multipath)

```
user@host> show route detail

10.1.1.8/30 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 262142
                Address: 0x901a010
                Next-hop reference count: 2
                Source: 10.1.1.2
                Next hop: 10.1.1.2 via ge-0/3/0.1, selected
                Next hop: 10.1.1.6 via ge-0/3/0.5
                State: <Active Ext>
                Local AS:     1 Peer AS:     2
                Age: 5:04:43
                Validation State: unverified
                Task: BGP_2.10.1.1.2+59955
                Announcement bits (1): 0-KRT
                AS path: 2 I
                Accepted Multipath
                Localpref: 100
                Router ID: 172.16.1.2
         BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 678
                Address: 0x8f97520
                Next-hop reference count: 9
                Source: 10.1.1.6
                Next hop: 10.1.1.6 via ge-0/3/0.5, selected
                State: <NotBest Ext>
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Inactive reason: Not Best in its group - Active preferred
                    Local AS:     1 Peer AS:     2
                    Age: 5:04:43
                 Validation State: unverified
                    Task: BGP_2.10.1.1.6+58198
                    AS path: 2 I
                    Accepted MultipathContrib
                    Localpref: 100
                    Router ID: 172.16.1.3
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
                Next hop: via vt-0/1/0.1
                Next-hop index: 661
                Label operation: Pop
                Address: 0x9172130
                Next hop: via so-0/0/3.0
                Next-hop index: 654
                Label operation: Swap 299872
                State: **Active Int>
                Local AS:  1001
                Age: 8:20       Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
               FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
 src 192.168.142.2
```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
                Next-hop reference count: 3
                Next hop: 1.3.8.2 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 1.3.4.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    10
                Age: 54:05      Metric: 1
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Validation State: unverified
                    Task: LDP
                    Announcement bits (1): 0-KRT
                    AS path: I
                    FECs bound to route: P2MP root-addr 172.16.1.1, grp: 232.1.1.1,
src: 192.168.219.11
                    Primary Upstream : 172.16.1.3:0--172.16.1.2:0
                       RPF Nexthops  :
                            ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
                            ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
                    Backup Upstream : 172.16.1.3:0--172.16.1.6:0
                       RPF Nexthops :
                            ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xfffe
                            ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xfffe
```


Copyright © 2017, Juniper Networks, Inc.

show route exact

List of Syntax
Syntax on page 319
Syntax (EX Series Switches) on page 319

Syntax
show route exact *destination-prefix*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route exact *destination-prefix*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display only the routes that exactly match the specified address or range of addresses.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

destination-prefix—Address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route exact on page 319
show route exact detail on page 320
show route exact extensive on page 320
show route exact terse on page 320

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route exact

```
user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                     > to 192.168.71.254 via fxp0.0
```

Copyright © 2017, Juniper Networks, Inc.

show route exact detail

```
user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2d 3:30:26
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I
```

show route exact extensive

```
user@host> show route exact 207.17.136.0/24 extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:25:18
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I
```

show route exact terse

```
user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 207.17.136.0/24    S   5                      >192.168.71.254
```


Copyright © 2017, Juniper Networks, Inc.

show route export

List of Syntax
Syntax on page 321
Syntax (EX Series Switches) on page 321

Syntax
show route export
<brief | detail>
<instance <*instance-name*> | *routing-table-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route export
<brief | detail>
<instance <*instance-name*> | *routing-table-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.

Options
none—(Same as **brief**.) Display standard information about policy-based export for all instances and routing tables on all systems.

brief | detail—(Optional) Display the specified level of output.

instance <*instance-name*>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route export inet** command).

Required Privilege Level
view

List of Sample Output
show route export on page 322
show route export detail on page 322
show route export instance detail on page 323

Output Fields
Table 21 on page 322 lists the output fields for the **show route export** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 21: show route export Output Fields

Field Name	Field Description	Level of Output
Table or *table-name*	Name of the routing tables that either import or export routes.	All levels
Routes	Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one.	**brief** none
Export	Whether the table is currently exporting routes to other tables: **Y** or **N** (Yes or No).	**brief** none
Import	Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)	**detail**
Flags	(**instance** keyword only) Flags for this feature on this instance: • **config auto-policy**—The policy was deduced from the configured IGP export policies. • **cleanup**—Configuration information for this instance is no longer valid. • **config**—The instance was explicitly configured.	**detail**
Options	(**instance** keyword only) Configured option displays the type of routing tables the feature handles: • **unicast**—Indicates *instance*.inet.0. • **multicast**—Indicates *instance*.inet.2. • **unicast multicast**—Indicates *instance*.inet.0 and *instance*.inet.2.	**detail**
Import policy	(**instance** keyword only) Policy that **route export** uses to construct the import-export matrix. Not displayed if the instance type is **vrf**.	**detail**
Instance	(**instance** keyword only) Name of the routing instance.	**detail**
Type	(**instance** keyword only) Type of routing instance: **forwarding**, **non-forwarding**, or **vrf**.	**detail**

Sample Output

show route export

```
user@host> show route export
Table                            Export           Routes
inet.0                           N                     0
black.inet.0                     Y                     3
red.inet.0                       Y                     4
```

show route export detail

```
user@host> show route export detail
inet.0                                           Routes:        0
black.inet.0                                     Routes:        3
  Import: [ inet.0 ]
red.inet.0                                       Routes:        4
  Import: [ inet.0 ]
```


Copyright © 2017, Juniper Networks, Inc.

show route export instance detail

```
user@host> show route export instance detail
Instance: master                          Type: forwarding
  Flags: <config auto-policy> Options: <unicast multicast>
  Import policy: [ (ospf-master-from-red || isis-master-from-black) ]
Instance: black                           Type: non-forwarding
Instance: red                             Type: non-forwarding
```

Copyright © 2017, Juniper Networks, Inc.

show route extensive

List of Syntax
Syntax on page 324
Syntax (EX Series Switches) on page 324

Syntax
show route extensive
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route extensive
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display extensive information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route extensive on page 331
show route extensive (Access Route) on page 338
show route extensive (BGP PIC Edge) on page 338
show route extensive (FRR and LFA) on page 339
show route extensive (IS-IS) on page 340
show route extensive (Route Reflector) on page 340
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 340
show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 341

Output Fields
Table 22 on page 324 describes the output fields for the **show route extensive** command. Output fields are listed in the approximate order in which they appear.

Table 22: show route extensive Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.

Copyright © 2017, Juniper Networks, Inc.

Table 22: show route extensive Output Fields *(continued)*

Field Name	Field Description
***number* routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). • **hidden** (routes that are not used because of a routing policy).

***route-destination* (entry, announced)**	Route destination (for example: 10.0.0.1/24). The **entry** value is the number of route for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.

TSI	Protocol header information.
label stacking	(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.

Copyright © 2017, Juniper Networks, Inc.

Table 22: show route extensive Output Fields *(continued)*

Field Name	Field Description
Level	(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see the Output Field table in the show route detail command.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Offset	Whether the metric has been increased or decreased by an offset value.
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.

Copyright © 2017, Juniper Networks, Inc.

Table 22: show route extensive Output Fields *(continued)*

Field Name	Field Description
label-operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Indirect next hops	When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved. When BGP PIC Edge is enabled, the output lines that contain **Indirect next hop: weight** follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values: • 0x1 indicates active next hops. • 0x4000 indicates passive next hops.
State	State of the route (a route can be in more than one state). See the Output Field table in the show route detail command.
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
Weight	Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown. For sample output, see show route table.

Copyright © 2017, Juniper Networks, Inc.

Table 22: show route extensive Output Fields *(continued)*

Field Name	Field Description
Inactive reason	If the route is inactive, the reason for its current state is indicated. Typical reasons include: • **Active preferred**—Currently active route was selected over this route. • **Always compare MED**—Path with a lower multiple exit discriminator (MED) is available. • **AS path**—Shorter AS path is available. • **Cisco Non-deterministic MED selection**—Cisco nondeterministic MED is enabled and a path with a lower MED is available. • **Cluster list length**—Path with a shorter cluster list length is available. • **Forwarding use only**—Path is only available for forwarding purposes. • **IGP metric**—Path through the next hop with a lower IGP metric is available. • **IGP metric type**—Path with a lower OSPF link-state advertisement type is available. • **Interior > Exterior > Exterior via Interior**—Direct, static, IGP, or EBGP path is available. • **Local preference**—Path with a higher local preference value is available. • **Next hop address**—Path with a lower metric next hop is available. • **No difference**—Path from a neighbor with a lower IP address is available. • **Not Best in its group**—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). • **Number of gateways**—Path with a higher number of next hops is available. • **Origin**—Path with a lower origin code is available. • **OSPF version**—Path does not support the indicated OSPF version. • **RIB preference**—Route from a higher-numbered routing table is available. • **Route destinguisher**—64-bit prefix added to IP subnets to make them unique. • **Route metric or MED comparison**—Route with a lower metric or MED is available. • **Route preference**—Route with a lower preference value is available. • **Router ID**—Path through a neighbor with a lower ID is available. • **Unusable path**—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. • **Update source**—Last tiebreaker is the lowest IP address value.
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see show route table.

Copyright © 2017, Juniper Networks, Inc.

Table 22: show route extensive Output Fields *(continued)*

Field Name	Field Description
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that are consumers of the route. Using the following output as an example, **Announcement bits (3): 0-KRT 5-Resolve tree 2 8-BGP RT Background** there are (3) announcement bits to reflect the three clients (protocols) that have state for this route: Kernel (0-KRT), 5 (resolution tree process 2), and 8 (BGP). The notation ***n*-Resolve inet** indicates that the route is used for route resolution for next hops found in the routing table. ***n*** is an index used by Juniper Networks customer support only.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.

Copyright © 2017, Juniper Networks, Inc.

Table 22: show route extensive Output Fields *(continued)*

Field Name	Field Description
route status	Indicates the status of a BGP route: • **Accepted**—The specified BGP route is imported by the default BGP policy. • **Import**—The route is imported into a Layer 3 VPN routing instance. • **Import-Protect**—A remote instance egress that is protected. • **Multipath**—A BGP multipath active route. • **MultipathContrib**—The route is not active but contributes to the BGP multipath. • **Protect**—An egress route that is protected. • **Stale**—A route that is marked stale due to graceful restart.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding Equivalent Class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.

Copyright © 2017, Juniper Networks, Inc.

Table 22: show route extensive Output Fields *(continued)*

Field Name	Field Description
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.
Node path count	Number of nodes in the path.
Forwarding nexthops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route extensive

```
user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
203.0.113.10/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.10/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    64496
                Age: 1:34:06
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

203.0.113.30/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:32:40
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
```

Copyright © 2017, Juniper Networks, Inc.

```
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    64496
                Age: 1:32:40    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

203.0.113.103/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
                State: <Active NoReadvrt Int>
                Local AS:    644969
                Age: 1:32:43
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

203.0.113.203/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.203/30 -> {203.0.113.216}
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 203.0.113.216 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:32:19    Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.2/32 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.22/32 -> {}
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
```


Copyright © 2017, Juniper Networks, Inc.

```
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

203.0.113.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
                State: <Active Int>
                Local AS:    64496
                Age: 1:28:12    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

203.0.113.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    64496
                Age: 1:28:12    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0   (1 entry, 1 announced)
TSI:
KRT in-kernel 0      /36 -> {}
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:08    Metric: 1
```

Copyright © 2017, Juniper Networks, Inc.

```
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
  299840 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 299840 /52 -> {indirect(1048575)}
          *RSVP   Preference: 7/2
                  Next hop type: Flood
                  Address: 0x9174a30
                  Next-hop reference count: 4
                  Next hop type: Router, Next hop index: 798
                  Address: 0x9174c28
                  Next-hop reference count: 2
                  Next hop: 198.51.100.2 via lt-1/2/0.9 weight 0x1
                  Label-switched-path R2-to-R4-2p2mp
                  Label operation: Pop
                  Next hop type: Router, Next hop index: 1048574
                  Address: 0x92544f0
                  Next-hop reference count: 2
                  Next hop: 198.51.100.2 via lt-1/2/0.7 weight 0x1
                  Label-switched-path R2-to-R200-p2mp
                  Label operation: Pop
                  Next hop: 198.51.100.2 via lt-1/2/0.5 weight 0x8001
                  Label operation: Pop
                  State: <Active Int>
                  Age: 1:29       Metric: 1
                  Task: RSVP
                  Announcement bits (1): 0-KRT
                  AS path: I...

800010 (1 entry, 1 announced)

TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:31:53
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0      /16 -> {indirect(1048574)}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 203.0.113.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:31:53    Metric2: 2
```


Copyright © 2017, Juniper Networks, Inc.

```
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Indirect next hops: 1
                        Protocol next hop: 203.0.113.103 Metric: 2
                        Push 800012
                        Indirect next hop: 87272e4 1048574
                        Indirect path forwarding next hops: 1
                               Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1

                        203.0.113.103/32 Originating RIB: inet.3
                          Metric: 2                      Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 203.0.113.216 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

2001:db8::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
```

Copyright © 2017, Juniper Networks, Inc.

```
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
                Task: MLD
                Announcement bits (1): 0-KRT
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:34:07
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

203.0.113.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 203.0.113.103:1
                Next-hop reference count: 7
                Source: 203.0.113.103
                Protocol next hop: 203.0.113.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    64496 Peer AS:    64496
                Age: 1:28:12    Metric2: 1
                Task: BGP_69.203.0.113.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 203.0.113.103
                Primary Routing Table bgp.l2vpn.0

203.0.113.152:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
```


Copyright © 2017, Juniper Networks, Inc.

```
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8, status-vector: 0x9F

203.0.113.152:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

203.0.113.163:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 203.0.113.163 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    64499
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512


203.0.113.55/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
  198.51.100.0/24 -> {Push 300112}
        *BGP    Preference: 170/-101
                Next hop type: Router
                Address: 0x925c208
                Next-hop reference count: 2
                Source: 203.0.113.9
                Next hop: 203.0.113.9 via ge-1/2/0.15, selected
                Label operation: Push 300112
                Label TTL action: prop-ttl
                State: <Active Ext>
                Local AS:  64509 Peer AS: 65539
                Age: 1w0d 23:06:56
                AIGP: 25
                Task: BGP_65539.203.0.113.9+56732
                Announcement bits (1): 0-KRT
                AS path: 65539 64508 I
                Accepted
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Route Label: 300112
                    Localpref: 100
                    Router ID: 213.0.113.99
```

show route extensive (Access Route)

```
user@host> show route 203.0.113.102 extensive
inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
203.0.113.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.102/32 -> {192.0.2.2}
OSPF area : 0.0.0.0, LSA ID : 203.0.113.102, LSA type : Extern
        *Access Preference: 13
                Next-hop reference count: 78472
                Next hop: 192.0.2.2 via fe-0/0/0.0, selected
                State: <Active Int>
       Age: 12
                Task: RPD Unix Domain Server./var/run/rpd_serv.local
                Announcement bits (2): 0-KRT 1-OSPFv2
                AS path: I


user@host> show route 2001:db8:4641:1::/48 extensive

inet6.0: 75 destinations, 81 routes (75 active, 0 holddown, 0 hidden)
  2001:db8:4641:1::/48 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 2001:db8:4641:1::/48 -> {#0 0.13.1.0.0.1}
          *Access Preference: 13
                  Next hop type: Router, Next hop index: 74548
                  Address: 0x1638c1d8
                  Next-hop reference count: 6
                  Next hop: #0 0.13.1.0.0.1 via demux0.1073753267, selected
                  Session Id: 0x0
                  State: <Active Int>
                  Age: 4:17
                  Validation State: unverified
                  Task: RPD Unix Domain Server./var/run/rpd_serv.local
                  Announcement bits (2): 0-KRT 4-Resolve tree 2
                  AS path: I
2001:db8:4641:1::/128 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 2001:db8:4641:1::/128 -> {#0 0.13.1.0.0.1}
          *Access-internal Preference: 12
                  Next hop type: Router, Next hop index: 74548
                  Address: 0x1638c1d8
                  Next-hop reference count: 6
                  Next hop: #0 0.13.1.0.0.1 via demux0.1073753267, selected
                  Session Id: 0x0
                  State: <Active Int>
                  Age: 4:17
                  Validation State: unverified
                  Task: RPD Unix Domain Server./var/run/rpd_serv.local
                  Announcement bits (2): 0-KRT 4-Resolve tree 2
                  AS path: I
```

show route extensive (BGP PIC Edge)

```
user@host> show route 198.51.100.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  198.51.100.6/32 (3 entries, 2 announced)
```


Copyright © 2017, Juniper Networks, Inc.

```
          State: <CalcForwarding>
  TSI:
  KRT in-kernel 198.51.100.6/32 -> {indirect(1048574), indirect(1048577)}
  Page 0 idx 0 Type 1 val 9219e30
      Nexthop: Self
      AS path: [2] 3 I
      Communities: target:2:1
  Path 198.51.100.6 from 198.51.100.4 Vector len 4.  Val: 0
..
          #Multipath Preference: 255
                  Next hop type: Indirect
                  Address: 0x93f4010
                  Next-hop reference count: 2
..
                  Protocol next hop: 198.51.1001.4
                  Push 299824
                  Indirect next hop: 944c000 1048574 INH Session ID: 0x3
                  Indirect next hop: weight 0x1
                  Protocol next hop: 198.51.100.5
                  Push 299824
                  Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
                  Indirect next hop: weight 0x4000
                  State: <ForwardingOnly Int Ext)>
                  Inactive reason: Forwarding use only
                  Age: 25         Metric2: 15
                  Validation State: unverified
                  Task: RT
                  Announcement bits (1): 0-KRT
                  AS path: 3 I
                  Communities: target:2:1
```

show route extensive (FRR and LFA)

```
user@host> show route 203.0.113.20 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
203.0.113.20/24 (2 entries, 1 announced)
        State: FlashAll
TSI:
KRT in-kernel 203.0.113.20/24 -> {Push 299776, Push 299792}
        *RSVP   Preference: 7/1
                Next hop type: Router, Next hop index: 1048574
                Address: 0xbbbc010
                Next-hop reference count: 5
                Next hop: 203.0.113.112 via ge-2/1/8.0 weight 0x1, selected
                Label-switched-path europa-d-to-europa-e
                Label operation: Push 299776
                Label TTL action: prop-ttl
                Session Id: 0x201
                Next hop: 203.0.113.122 via ge-2/1/4.0 weight 0x4001
                Label-switched-path europa-d-to-europa-e
                Label operation: Push 299792
                Label TTL action: prop-ttl
                Session Id: 0x202
                State: Active Int
                Local AS:   64500
                Age: 5:31  Metric: 2
                Task: RSVP
                Announcement bits (1): 0-KRT
                AS path: I
         OSPF   Preference: 10
                Next hop type: Router, Next hop index: 615
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Address: 0xb9d78c4
                    Next-hop reference count: 7
                    Next hop: 203.0.113.112 via ge-2/1/8.0, selected
                    Session Id: 0x201
                    State: Int
                    Inactive reason: Route Preference
                    Local AS:   64500
                    Age: 5:35  Metric: 3
                    Area: 0.0.0.0
                    Task: OSPF
                    AS path: I
```

show route extensive (IS-IS)

```
user@host> show route extensive
IS-IS Preference: 15
                    Level: 1
                    Next hop type: Router, Next hop index: 1048577
                    Address: 0xXXXXXXXXXX
                    Next-hop reference count: YY
                    Next hop: 203.0.113.22 via ae1.0 balance 43%, selected
                    Session Id: 0x141
                    Next hop: 203.0.113.22 via ae0.0 balance 57%
```

show route extensive (Route Reflector)

```
user@host> show route extensive
203.0.113.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 203.0.113.0/8 -> {indirect(40)}
        *BGP    Preference: 170/-101
                Source: 192.168.4.214
                Protocol next hop: 198.51.100.192 Indirect next hop: 84ac908 40
                State: <Active Int Ext>
                Local AS: 65548 Peer AS: 65548
                Age: 3:09       Metric: 0       Metric2: 0
                Task: BGP_65548.192.168.4.214+1033
                Announcement bits (2): 0-KRT 4-Resolve inet.0
                AS path: 65544 64507 I <Originator>
                Cluster list:  198.51.100.1
                Originator ID: 203.0.113.88
                Communities: 7777:7777
                Localpref: 100
                Router ID: 203.0.113.4
                Indirect next hops: 1
                        Protocol next hop: 203.0.113.192 Metric: 0
                        Indirect next hop: 84ac908 40
                        Indirect path forwarding next hops: 0
                                Next hop type: Discard
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Next hop: via vt-0/1/0.1
                    Next-hop index: 661
                    Label operation: Pop
                    Address: 0x9172130
                    Next hop: via so-0/0/3.0
                    Next-hop index: 654
                    Label operation: Swap 299872
                    State: **Active Int>
                    Local AS:  64511
                    Age: 8:20       Metric: 1
                    Task: LDP
                    Announcement bits (1): 0-KRT
                    AS path: I
                   FECs bound to route: P2MP root-addr 203.0.113.166, grp 203.0.1.1,
 src 192.168.142.2
```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
                Next-hop reference count: 3
                Next hop: 203.0.113.82 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 203.0.113.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    64500
                Age: 54:05      Metric: 1
                Validation State: unverified
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
               FECs bound to route: P2MP root-addr 198.51.100.1, grp: 232.1.1.1,
 src: 192.168.219.11
                Primary Upstream : 198.51.100.3:0--198.51.100.2:0
                   RPF Nexthops  :
                        ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
                        ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
                Backup Upstream : 198.51.100.3:0--198.51.100.6:0
                   RPF Nexthops :
                       ge-1/2/20.0, 198.51.100.96, Label: 301584, weight: 0xfffe

                       ge-1/2/19.0, 198.51.100.36, Label: 301584, weight: 0xfffe
```

Copyright © 2017, Juniper Networks, Inc.

show route forwarding-table

List of Syntax

Syntax on page 342
Syntax (MX Series Routers) on page 342
Syntax (TX Matrix and TX Matrix Plus Routers) on page 342

Syntax

show route forwarding-table
<detail | extensive | summary>
<all>
<ccc *interface-name*>
<destination *destination-prefix*>
<family *family* | matching *matching*>
<interface-name *interface-name*>
<label *name*>
<matching *matching*>
<multicast>
<table (default | *logical-system-name/routing-instance-name* | *routing-instance-name*)>
<vlan (all | *vlan-name*)>
<vpn *vpn*>

Syntax (MX Series Routers)

show route forwarding-table
<detail | extensive | summary>
<all>
<bridge-domain (all | *domain-name*)>
<ccc *interface-name*>
<destination *destination-prefix*>
<family *family* | matching *matching*>
<interface-name *interface-name*>
<label *name*>
<learning-vlan-id *learning-vlan-id*>
<matching *matching*>
<multicast>
<table (default | *logical-system-name/routing-instance-name* | *routing-instance-name*)>
<vlan (all | *vlan-name*)>
<vpn *vpn*>

Syntax (TX Matrix and TX Matrix Plus Routers)

show route forwarding-table
<detail | extensive | summary>
<all>
<ccc *interface-name*>
<destination *destination-prefix*>
<family *family* | matching *matching*>
<interface-name *interface-name*>
<matching *matching*>
<label *name*>
<lcc *number*>
<multicast>
<table *routing-instance-name*>
<vpn *vpn*>

Release Information

Command introduced before Junos OS Release 7.4.
Option **bridge-domain** introduced in Junos OS Release 7.5

Copyright © 2017, Juniper Networks, Inc.

Option **learning-vlan-id** introduced in Junos OS Release 8.4
Options **all** and **vlan** introduced in Junos OS Release 9.6.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.

NOTE: The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

Options

none—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

detail | extensive | summary—(Optional) Display the specified level of output.

all—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

bridge-domain (all | *bridge-domain-name*)—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

ccc *interface-name*—(Optional) Display route entries for the specified circuit cross-connect interface.

destination *destination-prefix*—(Optional) Destination prefix.

family *family*—(Optional) Display routing table entries for the specified family: **bridge (ccc | destination | detail | extensive | interface-name | label | learning-vlan-id | matching | multicast | summary | table | vlan | vpn)**, **ethernet-switching**, **evpn**, **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mcsnoop-inet**, **mcsnoop-inet6**, **mpls**, **satellite-inet**, **satellite-inet6**, **satellite-vpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

interface-name *interface-name*—(Optional) Display routing table entries for the specified interface.

label *name*—(Optional) Display route entries for the specified label.

lcc *number*—(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers,

Copyright © 2017, Juniper Networks, Inc.

display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

learning-vlan-id *learning-vlan-id*—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

matching *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast—(Optional) Display routing table entries for multicast routes.

table —(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the show route instance command.

vlan (all | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

vpn *vpn*—(Optional) Display routing table entries for a specified VPN.

Required Privilege Level view

List of Sample Output

show route forwarding-table on page 349
show route forwarding-table detail on page 350
show route forwarding-table destination extensive (Weights and Balances) on page 350
show route forwarding-table extensive on page 351
show route forwarding-table extensive (RPF) on page 352
show route forwarding-table (dynamic list next hop) on page 353
show route forwarding-table family mpls on page 354
show route forwarding-table family mpls ccc ge-0/0/1.1004 on page 354
show route forwarding-table family vpls on page 354
show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled) on page 354
show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics) on page 355
show route forwarding-table family vpls extensive on page 355

Copyright © 2017, Juniper Networks, Inc.

show route forwarding-table table default on page 356
show route forwarding-table table
logical-system-name/routing-instance-name on page 357
show route forwarding-table vpn on page 358

Output Fields Table 23 on page 345 lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

Table 23: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Logical system	Name of the logical system. This field is displayed if you specify the **table** ***logical-system-name/routing-instance-name*** option on a device that is configured for and supports logical systems.	All levels
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels

Copyright © 2017, Juniper Networks, Inc.

Table 23: show route forwarding-table Output Fields *(continued)*

Field Name	Field Description	Level of Output
Enabled protocols	The features and protocols that have been enabled for a given routing table. This field can contain the following values: • BUM hashing—BUM hashing is enabled. • MAC Stats—Mac Statistics is enabled. • Bridging—Routing instance is a normal layer 2 bridge. • No VLAN—No VLANs are associated with the bridge domain. • All VLANs—The **vlan-id all** statement has been enabled for this bridge domain. • Single VLAN—Single VLAN ID is associated with the bridge domain. • MAC action drop—New MACs will be dropped when the MAC address limit is reached. • Dual VLAN—Dual VLAN tags are associated with the bridge domain • No local switching—No local switching is enabled for this routing instance.. • Learning disabled—Layer 2 learning is disabled for this routing instance. • MAC limit reached—The maximum number of MAC addresses that was configured for this routing instance has been reached. • VPLS—The VPLS protocol is enabled. • No IRB l2-copy—The no-irb-layer-2-copy feature is enabled for this routing instance. • ACKed by all peers—All peers have acknowledged this routing instance. • BUM Pruning—BUM pruning is enabled on the VPLS instance. • Def BD VXLAN—VXLAN is enabled for the default bridge domain. • EVPN—EVPN protocol is enabled for this routing instance. • Def BD OVSDB—Open vSwitch Database (OVSDB) is enabled on the default bridge domain. • Def BD Ingress replication—VXLAN ingress node replication is enabled on the default bridge domain. • L2 backhaul—Layer 2 backhaul is enabled. • FRR optimize—Fast reroute optimization • MAC pinning—MAC pinning is enabled for this bridge domain. • MAC Aging Timer—The MAC table aging time is set per routing instance. • EVPN VXLAN—This routing instance supports EVPN with VXLAN encapsulation. • PBBN—This routing instance is configured as a provider backbone bridged network. • PBN—This routing instance is configured as a provider bridge network. • ETREE—The ETREE protocol is enabled on this EVPN routing instance. • ARP/NDP suppression—EVPN ARP NDP suppression is enabled in this routing instance. • Def BD EVPN VXLAN—EVPN VXLAN is enabled for the default bridge domain. • MPLS control word—Control word is enabled for this MPLS routing instance.	All levels
Address family	Address family (for example, **IP**, **IPv6**, **ISO**, **MPLS**, and **VPLS**).	All levels
Destination	Destination of the route.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 23: show route forwarding-table Output Fields *(continued)*

Field Name	Field Description	Level of Output
Route Type (Type)	How the route was placed into the forwarding table. When the **detail** keyword is used, the route type might be abbreviated (as shown in parentheses): • **cloned** (**clon**)—(TCP or multicast only) Cloned route. • **destination** (**dest**)—Remote addresses directly reachable through an interface. • **destination down** (**iddn**)—Destination route for which the interface is unreachable. • **interface cloned** (**ifcl**)—Cloned route for which the interface is unreachable. • **route down** (**ifdn**)—Interface route for which the interface is unreachable. • **ignore** (**ignr**)—Ignore this route. • **interface** (**intf**)—Installed as a result of configuring an interface. • **permanent** (**perm**)—Routes installed by the kernel when the routing table is initialized. • **user**—Routes installed by the routing protocol process or as a result of the configuration.	All levels
Route Reference (RtRef)	Number of routes to reference.	**detail extensive**
Flags	Route type flags: • **none**—No flags are enabled. • **accounting**—Route has accounting enabled. • **cached**—Cache route. • **incoming-iface** ***interface-number***—Check against incoming interface. • **prefix load balance**—Load balancing is enabled for this prefix. • **rt nh decoupled**—Route has been decoupled from the next hop to the destination. • **sent to PFE**—Route has been sent to the Packet Forwarding Engine. • **static**—Static route.	**extensive**
Next hop	IP address of the next hop to the destination.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 23: show route forwarding-table Output Fields *(continued)*

Field Name	Field Description	Level of Output
Next hop Type (Type)	Next-hop type. When the **detail** keyword is used, the next-hop type might be abbreviated (as indicated in parentheses): • **broadcast (bcst)**—Broadcast. • **deny**—Deny. • **discard (dscd)** —Discard. • **hold**—Next hop is waiting to be resolved into a unicast or multicast type. • **indexed (idxd)**—Indexed next hop. • **indirect (indr)**—Indirect next hop. • **local (locl)**—Local address on an interface. • **routed multicast (mcrt)**—Regular multicast next hop. • **multicast (mcst)**—Wire multicast next hop (limited to the LAN). • **multicast discard (mdsc)**—Multicast discard. • **multicast group (mgrp)**—Multicast group member. • **receive (recv)**—Receive. • **reject (rjct)**—Discard. An ICMP unreachable message was sent. • **resolve (rslv)**—Resolving the next hop. • **unicast (ucst)**—Unicast. • **unilist (ulst)**—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.	**detail extensive**
Index	Software index of the next hop that is used to route the traffic for a given prefix.	**detail extensive** none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	**extensive**
Reference (NhRef)	Number of routes that refer to this next hop.	**detail extensive** none
Next-hop interface (Netif)	Interface used to reach the next hop.	**detail extensive** none
Weight	Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the **Balance** field description).	**extensive**
Balance	Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.	**extensive**
RPF interface	List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when **rpf-check** is configured on the interface.	**extensive**

Copyright © 2017, Juniper Networks, Inc.

Sample Output

show route forwarding-table

```
user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    46     4
0.0.0.0/32         perm     0                    dscd    44     1
172.16.1.0/24         ifdn     0                    rslv   608     1 ge-2/0/1.0
172.16.1.0/32         iddn     0 172.16.1.0           recv   606     1 ge-2/0/1.0
172.16.1.1/32         user     0                    rjct    46     4
172.16.1.1/32         intf     0 172.16.1.1           locl   607     2
172.16.1.1/32         iddn     0 172.16.1.1           locl   607     2
172.16.1.255/32       iddn     0 ff:ff:ff:ff:ff:ff  bcst   605     1 ge-2/0/1.0
10.0.0.0/24        intf     0                    rslv   616     1 ge-2/0/0.0
10.0.0.0/32        dest     0 10.0.0.0           recv   614     1 ge-2/0/0.0
10.0.0.1/32        intf     0 10.0.0.1           locl   615     2
10.0.0.1/32        dest     0 10.0.0.1           locl   615     2
10.0.0.255/32      dest     0 10.0.0.255         bcst   613     1 ge-2/0/0.0
10.1.1.0/24        ifdn     0                    rslv   612     1 ge-2/0/1.0
10.1.1.0/32        iddn     0 10.1.1.0           recv   610     1 ge-2/0/1.0
10.1.1.1/32        user     0                    rjct    46     4
10.1.1.1/32        intf     0 10.1.1.1           locl   611     2
10.1.1.1/32        iddn     0 10.1.1.1           locl   611     2
10.1.1.255/32      iddn     0 ff:ff:ff:ff:ff:ff  bcst   609     1 ge-2/0/1.0
10.206.0.0/16      user     0 10.209.63.254      ucst   419    20 fxp0.0
10.209.0.0/16      user     1 0:12:1e:ca:98:0    ucst   419    20 fxp0.0
10.209.0.0/18      intf     0                    rslv   418     1 fxp0.0
10.209.0.0/32      dest     0 10.209.0.0         recv   416     1 fxp0.0
10.209.2.131/32    intf     0 10.209.2.131       locl   417     2
10.209.2.131/32    dest     0 10.209.2.131       locl   417     2
10.209.17.55/32    dest     0 0:30:48:5b:78:d2   ucst   435     1 fxp0.0
10.209.63.42/32    dest     0 0:23:7d:58:92:ca   ucst   434     1 fxp0.0
10.209.63.254/32   dest     0 0:12:1e:ca:98:0    ucst   419    20 fxp0.0
10.209.63.255/32   dest     0 10.209.63.255      bcst   415     1 fxp0.0
10.227.0.0/16      user     0 10.209.63.254      ucst   419    20 fxp0.0

...

Routing table: iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                   rjct    27     1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf     0                   locl    28     1

Routing table: inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                   rjct     6     1
ff00::/8           perm     0                   mdsc     4     1
ff02::1/128        perm     0 ff02::1           mcst     3     1

Routing table: ccc
MPLS:
Interface.Label    Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 16     1
100004(top)fe-0/0/1.0
```

Copyright © 2017, Juniper Networks, Inc.

show route forwarding-table detail

```
user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            user     2 0:90:69:8e:b1:1b   ucst   132     4 fxp0.0
default            perm     0                    rjct    14     1
10.1.1.0/24        intf     0 ff.3.0.21          ucst   322     1 so-5/3/0.0
10.1.1.0/32        dest     0 10.1.1.0           recv   324     1 so-5/3/0.0
10.1.1.1/32        intf     0 10.1.1.1           locl   321     1
10.1.1.255/32      dest     0 10.1.1.255         bcst   323     1 so-5/3/0.0
10.21.21.0/24      intf     0 ff.3.0.21          ucst   326     1 so-5/3/0.0
10.21.21.0/32      dest     0 10.21.21.0         recv   328     1 so-5/3/0.0
10.21.21.1/32      intf     0 10.21.21.1         locl   325     1
10.21.21.255/32    dest     0 10.21.21.255       bcst   327     1 so-5/3/0.0
127.0.0.1/32       intf     0 127.0.0.1          locl   320     1
172.17.28.19/32    clon     1 192.168.4.254      ucst   132     4 fxp0.0
172.17.28.44/32    clon     1 192.168.4.254      ucst   132     4 fxp0.0

...

Routing table: private1__.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    46     1
10.0.0.0/8         intf     0                    rslv   136     1 fxp1.0
10.0.0.0/32        dest     0 10.0.0.0           recv   134     1 fxp1.0
10.0.0.4/32        intf     0 10.0.0.4           locl   135     2
10.0.0.4/32        dest     0 10.0.0.4           locl   135     2

...

Routing table: iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    38     1

Routing table: inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    22     1
ff00::/8           perm     0                    mdsc    21     1
ff02::1/128        perm     0 ff02::1            mcst    17     1

...

Routing table: mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 28     1
```

show route forwarding-table destination extensive (Weights and Balances)

```
user@host> show route forwarding-table destination 3.4.2.1 extensive
Routing table: inet [Index 0]
Internet:

Destination:  3.4.2.1/32
  Route type: user
```


Copyright © 2017, Juniper Networks, Inc.

```
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: unilist                Index: 262143   Reference: 1
  Nexthop: 172.16.4.4
  Next-hop type: unicast                Index: 335      Reference: 2
  Next-hop interface: so-1/1/0.0        Weight: 22      Balance: 3
  Nexthop: 145.12.1.2
  Next-hop type: unicast                Index: 337      Reference: 2
  Next-hop interface: so-0/1/2.0        Weight: 33      Balance: 33
```

show route forwarding-table extensive

```
user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:

Destination:  default
  Route type: user
  Route reference: 2                    Route interface-index: 0
  Flags: sent to PFE
  Nexthop: 00:00:5E:00:53:1b
  Next-hop type: unicast                Index: 132      Reference: 4
  Next-hop interface: fxp0.0

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: none
  Next-hop type: reject                 Index: 14       Reference: 1

Destination:  127.0.0.1/32
  Route type: interface
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Nexthop: 127.0.0.1
  Next-hop type: local                  Index: 320      Reference: 1

...

Routing table: private1__.inet [Index 1]
Internet:

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                 Index: 46       Reference: 1

Destination:  10.0.0.0/8
  Route type: interface
  Route reference: 0                    Route interface-index: 3
  Flags: sent to PFE
  Next-hop type: resolve                Index: 136      Reference: 1
  Next-hop interface: fxp1.0

...

Routing table: iso [Index 0]
ISO:

Destination:  default
```

Copyright © 2017, Juniper Networks, Inc.

```
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                 Index: 38       Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                 Index: 22       Reference: 1

Destination:  ff00::/8
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: multicast discard      Index: 21       Reference: 1

...

Routing table: private1__.inet6 [Index 1]
Internet6:

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                 Index: 54       Reference: 1

Destination:  fe80::2a0:a5ff:fe3d:375/128
  Route type: interface
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Nexthop: fe80::2a0:a5ff:fe3d:375
  Next-hop type: local                  Index: 75       Reference: 1

...
```

show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```
so-1/1/0 {
    unit 0 {
        family inet {
            rpf-check;
            address 192.0.2.2/30;
        }
    }
}
```

```
user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:
...
...
Destination:  192.0.2.3/32
```


Copyright © 2017, Juniper Networks, Inc.

```
  Route type: destination
  Route reference: 0                   Route interface-index: 67
  Flags: sent to PFE
  Nexthop: 192.0.2.3
  Next-hop type: broadcast             Index: 328      Reference: 1
  Next-hop interface: so-1/1/0.0
  RPF interface: so-1/1/0.0
```

show route forwarding-table (dynamic list next hop)

The **show route forwarding table** output shows the two next hop elements for a multihomed EVPN destination.

```
user@host> show route forwarding-table label 299952 extensive
MPLS:

Destination:  299952
  Route type: user
  Route reference: 0                   Route interface-index: 0
  Multicast RPF nh index: 0
  P2mpidx: 0
  Flags: sent to PFE, rt nh decoupled
  Next-hop type: indirect              Index: 1048575  Reference: 2
  Nexthop:
  Next-hop type: composite             Index: 601      Reference: 2
  Next-hop type: indirect              Index: 1048574  Reference: 3
  Nexthop: 1.0.0.4
  Next-hop type: Push 301632, Push 299776(top) Index: 600 Reference: 2
  Load Balance Label: None
  Next-hop interface: ge-0/0/1.0
  Next-hop type: indirect              Index: 1048577  Reference: 3
  Nexthop: 1.0.0.4
  Next-hop type: Push 301344, Push 299792(top) Index: 619 Reference: 2
  Load Balance Label: None
  Next-hop interface: ge-0/0/1.0
```

After one of the PE router has been disabled in the EVPN multihomed network, the same **show route forwarding table** output command shows one next hop element and one empty next hop element.

```
user@host> show route forwarding-table label 299952 extensive
Routing table: default.mpls [Index 0]
MPLS:

Destination:  299952
  Route type: user
  Route reference: 0                   Route interface-index: 0
  Multicast RPF nh index: 0
  P2mpidx: 0
  Flags: sent to PFE, rt nh decoupled
  Next-hop type: indirect              Index: 1048575  Reference: 2
  Nexthop:
  Next-hop type: composite             Index: 601      Reference: 2
  Next-hop type: indirect              Index: 1048577  Reference: 3
  Nexthop: 1.0.0.4
  Next-hop type: Push 301344, Push 299792(top) Index: 619 Reference: 2
  Load Balance Label: None
  Next-hop interface: ge-0/0/1.0
```

Copyright © 2017, Juniper Networks, Inc.

show route forwarding-table family mpls

```
user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            perm     0                   rjct    19     1
0                  user     0                   recv    18     3
1                  user     0                   recv    18     3
2                  user     0                   recv    18     3
100000             user     0 10.31.1.6         swap  100001     fe-1/1/0.0
800002             user     0                   Pop              vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
                   user     0                   indr   351     4
                                                Push 800000, Push 100002(top)
so-0/0/0.0
```

show route forwarding-table family mpls ccc ge-0/0/1.1004

```
user@host>show route forwarding-table mpls ccc ge-0/0/1.1004
Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop          Type  Index    NhRef Netif
ge-0/0/1.1004  (CCC) user     0                    ulst   1048577     2
                                                   comp      754        3
                                                  comp      755        3
                                                comp      756        3

Routing table: __mpls-oam__.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                    dscd      556     1
```

show route forwarding-table family vpls

```
user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            dynm     0                   flood   353     1
default            perm     0                   rjct    298     1
fe-0/1/0.0         dynm     0                   flood   355     1
00:00:5E:00:53:1f/48             <<<<<Remote CE

                   dynm     0                   indr   351     4
                                                Push 800000, Push 100002(top)
so-0/0/0.0
00:00:5E:00:53:1f/48             <<<<<<Local CE

                   dynm     0                   ucst   354     2 fe-0/1/0.0
```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled)

```
user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing
Destination        Type RtRef Next hop           Type Index    NhRef Netif
```


Copyright © 2017, Juniper Networks, Inc.

```
default            perm     0                   dscd      519     1
lsi.1048832        intf     0                   indr  1048574     4
                              172.16.3.2          Push 262145      621     2
ge-3/0/0.0
00:00:5E:00:53:01/48 user     0                  ucst      590     5 ge-2/3/9.0
0x30003/51         user     0                   comp      627     2
ge-2/3/9.0         intf     0                   ucst      590     5 ge-2/3/9.0
ge-3/1/3.0         intf     0                   ucst      619     4 ge-3/1/3.0
0x30002/51         user     0                   comp      600     2
0x30001/51         user     0                   comp      597     2
```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics)

```
user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                   dscd      519     1
lsi.1048834        intf     0                   indr  1048574     4
                              172.16.3.2          Push 262145      592     2
ge-3/0/0.0
00:19:e2:25:d0:01/48 user     0                  ucst      590     5 ge-2/3/9.0
0x30003/51         user     0                   comp      630     2
ge-2/3/9.0         intf     0                   ucst      590     5 ge-2/3/9.0
ge-3/1/3.0         intf     0                   ucst      591     4 ge-3/1/3.0
0x30002/51         user     0                   comp      627     2
0x30001/51         user     0                   comp      624     2
```

show route forwarding-table family vpls extensive

```
user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

Destination:  default
  Route type: dynamic
  Route reference: 0                    Route interface-index: 72
  Flags: sent to PFE
  Next-hop type: flood                  Index: 289      Reference: 1
  Next-hop type: unicast                Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0
  Next-hop type: unicast                Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: none
  Next-hop type: discard                Index: 341      Reference: 1

Destination:  fe-0/1/2.0
  Route type: dynamic
  Route reference: 0                    Route interface-index: 69
  Flags: sent to PFE
  Next-hop type: flood                  Index: 293      Reference: 1
  Next-hop type: indirect               Index: 363      Reference: 4
  Next-hop type: Push 800016
  Next-hop interface: at-1/0/1.0
```

Copyright © 2017, Juniper Networks, Inc.

```
  Next-hop type: indirect              Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
  Next-hop type: unicast               Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0

Destination:  fe-0/1/3.0
  Route type: dynamic
  Route reference: 0                   Route interface-index: 70
  Flags: sent to PFE
  Next-hop type: flood                 Index: 292      Reference: 1
  Next-hop type: indirect              Index: 363      Reference: 4
  Next-hop type: Push 800016
  Next-hop interface: at-1/0/1.0
  Next-hop type: indirect              Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
  Next-hop type: unicast               Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0

Destination:  00:00:5E:00:53:01/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 70
  Flags: sent to PFE, prefix load balance
  Next-hop type: unicast               Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0
  Route used as destination:
    Packet count:         6640    Byte count:      675786
  Route used as source
    Packet count:         6894    Byte count:      696424

Destination:  00:00:5E:00:53:04/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 69
  Flags: sent to PFE, prefix load balance
  Next-hop type: unicast               Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0
  Route used as destination:
    Packet count:           96    Byte count:        8079
  Route used as source:
    Packet count:          296    Byte count:       24955

Destination:  00:00:5E:00:53:05/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 74
  Flags: sent to PFE, prefix load balance
  Next-hop type: indirect              Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
```

show route forwarding-table table default

```
user@host> show route forwarding-table table default
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    36     2
0.0.0.0/32         perm     0                    dscd    34     1
```


Copyright © 2017, Juniper Networks, Inc.

```
10.0.60.0/30       user     0 10.0.60.13         ucst   713     5 fe-0/1/3.0
10.0.60.12/30      intf     0                    rslv   688     1 fe-0/1/3.0
10.0.60.12/32      dest     0 10.0.60.12         recv   686     1 fe-0/1/3.0
10.0.60.13/32      dest     0 0:5:85:8b:bc:22    ucst   713     5 fe-0/1/3.0
10.0.60.14/32      intf     0 10.0.60.14         locl   687     2
10.0.60.14/32      dest     0 10.0.60.14         locl   687     2
10.0.60.15/32      dest     0 10.0.60.15         bcst   685     1 fe-0/1/3.0
10.0.67.12/30      user     0 10.0.60.13         ucst   713     5 fe-0/1/3.0
10.0.80.0/30       ifdn     0 ff.3.0.21          ucst   676     1 so-0/0/1.0
10.0.80.0/32       dest     0 10.0.80.0          recv   678     1 so-0/0/1.0
10.0.80.2/32       user     0                    rjct    36     2
10.0.80.2/32       intf     0 10.0.80.2          locl   675     1
10.0.80.3/32       dest     0 10.0.80.3          bcst   677     1 so-0/0/1.0
10.0.90.12/30      intf     0                    rslv   684     1 fe-0/1/0.0
10.0.90.12/32      dest     0 10.0.90.12         recv   682     1 fe-0/1/0.0
10.0.90.14/32      intf     0 10.0.90.14         locl   683     2
10.0.90.14/32      dest     0 10.0.90.14         locl   683     2
10.0.90.15/32      dest     0 10.0.90.15         bcst   681     1 fe-0/1/0.0
10.5.0.0/16        user     0 192.168.187.126    ucst   324    15 fxp0.0
10.10.0.0/16       user     0 192.168.187.126    ucst   324    15 fxp0.0
10.13.10.0/23      user     0 192.168.187.126    ucst   324    15 fxp0.0
10.84.0.0/16       user     0 192.168.187.126    ucst   324    15 fxp0.0
10.150.0.0/16      user     0 192.168.187.126    ucst   324    15 fxp0.0
10.157.64.0/19     user     0 192.168.187.126    ucst   324    15 fxp0.0
10.209.0.0/16      user     0 192.168.187.126    ucst   324    15 fxp0.0

...

Routing table: default.iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    60     1

Routing table: default.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    44     1
::/128             perm     0                    dscd    42     1
ff00::/8           perm     0                    mdsc    43     1
ff02::1/128        perm     0 ff02::1            mcst    39     1

Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd    50     1
```

show route forwarding-table table logical-system-name/routing-instance-name

```
user@host> show route forwarding-table table R4/vpn-red
Logical system: R4
Routing table: vpn-red.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   563     1
0.0.0.0/32         perm     0                    dscd   561     2
172.16.0.1/32         user     0                    dscd   561     2
172.16.2.0/24         intf     0                    rslv   771     1 ge-1/2/0.3
172.16.2.0/32        dest     0 172.16.2.0           recv   769     1 ge-1/2/0.3
172.16.2.1/32         intf     0 172.16.2.1           locl   770     2
172.16.2.1/32         dest     0 172.16.2.1           locl   770     2
172.16.2.2/32         dest     0 0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0
```

```
                                                ucst   789     1 ge-1/2/0.3
172.16.2.255/32       dest     0 172.16.2.255          bcst   768     1 ge-1/2/0.3
172.16.233.0/4         perm     1                     mdsc   562     1
172.16.233.1/32        perm     0 172.16.233.1           mcst   558     1
255.255.255.255/32 perm     0                     bcst   559     1

Logical system: R4
Routing table: vpn-red.iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   608     1

Logical system: R4
Routing table: vpn-red.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   708     1
::/128             perm     0                    dscd   706     1
ff00::/8           perm     0                    mdsc   707     1
ff02::1/128        perm     0 ff02::1            mcst   704     1

Logical system: R4
Routing table: vpn-red.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd   638
```

show route forwarding-table vpn

```
user@host> show route forwarding-table vpn VPN-A
      Routing table:: VPN-A.inet
Internet:
      Destination        Type RtRef Nexthop           Type Index NhRef Netif
      default            perm     0                   rjct     4     4
      10.39.10.20/30     intf     0 ff.3.0.21         ucst    40     1
      so-0/0/0.0
      10.39.10.21/32     intf     0 10.39.10.21       locl    36     1
      10.255.14.172/32   user     0                   ucst    69     2
      so-0/0/0.0
      10.255.14.175/32   user     0                   indr    81     3
                                                      Push 100004, Push
      100004(top) so-1/0/0.0
      172.16.233.0/4         perm     2                   mdsc     5     3
      172.16.233.1/32        perm     0 172.16.233.1         mcst     1     8
      172.16.233.5/32        user     1 172.16.233.5         mcst     1     8
      255.255.255.255/32 perm     0                   bcst     2     3
```

On QFX5200, the results for this command look like this:

```
show route forwarding-table family mpls
Routing table: default.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 65 1
0 user 0 recv 64 4
1 user 0 recv 64 4
2 user 0 recv 64 4
13 user 0 recv 64 4
300384 user 0 9.1.1.1 Pop 1711 2 xe-0/0/34.0
300384(S=0) user 0 9.1.1.1 Pop 1712 2 xe-0/0/34.0
```


Copyright © 2017, Juniper Networks, Inc.

```
300400 user 0 ulst 131071 2
                              10.1.1.2 Pop 1713 1 xe-0/0/38.0
                              172.16.11.2 Pop 1714 1 xe-0/0/40.0
300400(S=0) user 0 ulst 131072 2
                              10.1.1.2 Pop 1715 1 xe-0/0/38.0
                              172.16.11.2 Pop 1716 1 xe-0/0/40.0

Routing table: __mpls-oam__.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 1681 1
```

Copyright © 2017, Juniper Networks, Inc.

show route hidden

Syntax show route hidden
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Release Information Command introduced before Junos OS Release 7.4.

Description Display only hidden route information. A hidden route is unusable, even if it is the best path.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level view

Related Documentation

- *Understanding Hidden Routes*

List of Sample Output

show route hidden on page 360
show route hidden detail on page 361
show route hidden extensive on page 361
show route hidden terse on page 361

Output Fields For information about output fields, see the output field table for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route hidden

```
user@host> show route hidden
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
127.0.0.1/32        [Direct/0] 04:26:38
                    > via lo0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.5.5.5/32         [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                      AS path: 100 I
```


Copyright © 2017, Juniper Networks, Inc.

```
                      Unusable
10.12.1.0/24        [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                      AS path: 100 I
                      Unusable
10.12.80.4/30       [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                      AS path: I
                      Unusable
...
```

show route hidden detail

```
user@host> show route hidden detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
127.0.0.1/32 (1 entry, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Hidden Martian Int>
                Local AS:     1
                Age: 4:27:37
                Task: IF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.5.5.5/32 (1 entry, 0 announced)
         BGP    Preference: 170/-101
                Route Distinguisher: 10.4.4.4:4
                Next hop type: Unusable
                Next-hop reference count: 6
                State: <Secondary Hidden Int Ext>
                Local AS:     1 Peer AS:     1
                Age: 3:45:09
                Task: BGP_1.10.4.4.4+2493
                AS path: 100 I
                Communities: target:1:999
                VPN Label: 100064
                Localpref: 100
                Router ID: 10.4.4.4
                Primary Routing Table bgp.l3vpn.0

...
```

show route hidden extensive

The output for the **show route hidden extensive** command is identical to that of the **show route hidden detail** command. For sample output, see show route hidden detail on page 361.

show route hidden terse

```
user@host> show route hidden terse
```

Copyright © 2017, Juniper Networks, Inc.

```
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  127.0.0.1/32       D   0                     >lo0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.5.5.5/32        B 170        100             Unusable        100 I
  10.12.1.0/24       B 170        100             Unusable        100 I
  10.12.80.4/30      B 170        100             Unusable        I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.4.4.4:4:10.5.5.5/32
                     B 170        100             Unusable        100 I
  10.4.4.4:4:10.12.1.0/24
                     B 170        100             Unusable        100 I
  10.4.4.4:4:10.12.80.4/30
                     B 170        100             Unusable        I

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

show route inactive-path

List of Syntax
Syntax on page 363
Syntax (EX Series Switches) on page 363

Syntax
show route inactive-path
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route inactive-path
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.

Options
none—Display all inactive routes.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show route active-path on page 281

List of Sample Output
show route inactive-path on page 363
show route inactive-path detail on page 364
show route inactive-path extensive on page 365
show route inactive-path terse on page 365

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-path

```
user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

```
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30     [OSPF/10] 03:57:28, metric 1
                    > via so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/8          [Direct/0] 04:39:56
                    > via fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30       [BGP/170] 04:38:17, localpref 100
                      AS path: 100 I
                    > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route inactive-path detail

```
user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:     1
                Age: 3:58:24    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp1.0, selected
                State: <NotBest Int>
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Inactive reason: No difference
                    Age: 4:40:52
                    Task: IF
                    AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)
         BGP    Preference: 170/-101
                Next-hop reference count: 6
                Source: 10.12.80.1
                Next hop: 10.12.80.1 via ge-6/3/2.0, selected
                State: <Ext>
                Inactive reason: Route Preference
                Peer AS:   100
                Age: 4:39:13
                Task: BGP_100.10.12.80.1+179
                AS path: 100 I
                Localpref: 100
                Router ID: 10.0.0.0
```

show route inactive-path extensive

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see show route inactive-path detail on page 364.

show route inactive-path terse

```
user@host> show route inactive-path terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.100.12/30    O  10          1            >so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.0.0.0/8         D   0                     >fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.80.0/30      B 170        100            >10.12.80.1      100 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
```

Copyright © 2017, Juniper Networks, Inc.

```
bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

show route instance

List of Syntax
Syntax on page 367
Syntax (EX Series Switches and QFX Series) on page 367

Syntax

```
show route instance
<brief | detail | summary>
<instance-name>
<logical-system (all | logical-system-name)>
<operational>
```

Syntax (EX Series Switches and QFX Series)

```
show route instance
<brief | detail | summary>
<instance-name>
<operational>
```

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description

Display routing instance information.

Options

none—(Same as **brief**) Display standard information about all routing instances.

brief | detail | summary—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**. (These options are not available with the **operational** keyword.)

instance-name—(Optional) Display information for all routing instances whose name begins with this string (for example, **cust1**, **cust11**, and **cust111** are all displayed when you run the **show route instance cust1** command).

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

operational—(Optional) Display operational routing instances.

Required Privilege Level

view

Related Documentation

- *Example: Transporting IPv6 Traffic Across IPv4 Using Filter-Based Tunneling*
- *Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart*

List of Sample Output

show route instance on page 369
show route instance detail (Graceful Restart Complete) on page 369
show route instance detail (Graceful Restart Incomplete) on page 371

Copyright © 2017, Juniper Networks, Inc.

show route instance detail (VPLS Routing Instance) on page 373
show route instance operational on page 373
show route instance summary on page 373

Output Fields Table 24 on page 368 lists the output fields for the **show route instance** command. Output fields are listed in the approximate order in which they appear.

Table 24: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or *instance-name*	Name of the routing instance.	All levels
Operational Routing Instances	(**operational** keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: **forwarding**, **l2vpn**, **no-forwarding**, **vpls**, **virtual-router**, or **vrf**.	All levels
State	State of the routing instance: **active** or **inactive**.	**brief detail** none
Interfaces	Name of interfaces belonging to this routing instance.	**brief detail** none
Restart State	Status of graceful restart for this instance: **Pending** or **Complete**.	**detail**
Path selection timeout	Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is **300**.	**detail**
Tables	Tables (and number of routes) associated with this routing instance.	**brief detail** none
Route-distinguisher	Unique route distinguisher associated with this routing instance.	**detail**
Vrf-import	VPN routing and forwarding instance import policy name.	**detail**
Vrf-export	VPN routing and forwarding instance export policy name.	**detail**
Vrf-import-target	VPN routing and forwarding instance import target community name.	**detail**
Vrf-export-target	VPN routing and forwarding instance export target community name.	**detail**
Vrf-edge-protection-id	Context identifier configured for edge-protection.	**detail**
Fast-reroute-priority	Fast reroute priority setting for a VPLS routing instance: **high**, **medium**, or **low**. The default is **low**.	**detail**
Restart State	Restart state: • **Pending:*protocol-name***—List of protocols that have not yet completed graceful restart for this routing table. • **Complete**—All protocols have restarted for this routing table.	**detail**

Copyright © 2017, Juniper Networks, Inc.

Table 24: show route instance Output Fields *(continued)*

Field Name	Field Description	Level of Output
Primary rib	Primary table for this routing instance.	**brief** none **summary**
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

Sample Output

show route instance

```
user@host> show route instance
Instance             Type
         Primary RIB                                 Active/holddown/hidden
master               forwarding
         inet.0                                      16/0/1
         iso.0                                       1/0/0
         mpls.0                                      0/0/0
         inet6.0                                     2/0/0
         l2circuit.0                                 0/0/0
__juniper_private1__ forwarding
         __juniper_private1__.inet.0                 12/0/0
         __juniper_private1__.inet6.0                1/0/0
```

show route instance detail (Graceful Restart Complete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding        State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0                : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0                 : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0           : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0               : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0           : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0       : 4 routes (4 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

```
    Restart Complete
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0           : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
    BGP-L.mpls.0           : 3 routes (3 active, 0 holddown, 0 hidden)
    Restart Complete
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn             State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
LDP:
  Router ID: 10.69.105.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0             : 5 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
OSPF:
  Router ID: 10.69.101.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Vrf-import-target: [ target:11111
  Tables:
    OSPF.inet.0            : 8 routes (7 active, 0 holddown, 0 hidden)
    Restart Complete
RIP:
  Router ID: 10.69.102.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
```


Copyright © 2017, Juniper Networks, Inc.

```
  Tables:
    RIP.inet.0              : 6 routes (6 active, 0 holddown, 0 hidden)
    Restart Complete
STATIC:
  Router ID: 10.69.100.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0           : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
```

show route instance detail (Graceful Restart Incomplete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding        State: Active
  Restart State: Pending  Path selection timeout: 300
  Tables:
    inet.0                  : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                  : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                   : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                  : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0             : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0             : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0         : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0            : 6 routes (5 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

```
    Restart Pending: VPN
    BGP-L.mpls.0           : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn              State: Active
  Restart State: Pending   Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN
LDP:
  Router ID: 10.69.105.1
  Type: vrf                State: Active
  Restart State: Pending   Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0             : 5 routes (4 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF LDP VPN
OSPF:
  Router ID: 10.69.101.1
  Type: vrf                State: Active
  Restart State: Pending   Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0            : 8 routes (7 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF VPN
RIP:
  Router ID: 10.69.102.1
  Type: vrf                State: Active
  Restart State: Pending   Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0             : 8 routes (6 active, 2 holddown, 0 hidden)
    Restart Pending: RIP VPN
STATIC:
  Router ID: 10.69.100.1
  Type: vrf                State: Active
  Restart State: Pending   Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
```


Copyright © 2017, Juniper Networks, Inc.

```
    STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
```

show route instance detail (VPLS Routing Instance)

```
user@host> show route instance detail test-vpls
test-vpls:
  Router ID: 0.0.0.0
  Type: vpls              State: Active
  Interfaces:
    lsi.1048833
    lsi.1048832
    fe-0/1/0.513
  Route-distinguisher: 10.255.37.65:1
  Vrf-import: [ __vrf-import-test-vpls-internal__ ]
  Vrf-export: [ __vrf-export-test-vpls-internal__ ]
  Vrf-import-target: [ target:300:1 ]
  Vrf-export-target: [ target:300:1 ]
  Vrf-edge-protection-id:  166.1.3.1  Fast-reroute-priority: high
  Tables:
    test-vpls.l2vpn.0          : 3 routes (3 active, 0 holddown, 0 hidden)
```

show route instance operational

```
user@host> show route instance operational
Operational Routing Instances:

master
default
```

show route instance summary

```
user@host> show route instance summary
Instance           Type           Primary rib            Active/holddown/hidden
master             forwarding
                                  inet.0                 15/0/1
                                  iso.0                  1/0/0
                                  mpls.0                 35/0/0
                                  l3vpn.0                0/0/0
                                  inet6.0                2/0/0
                                  l2vpn.0                0/0/0
                                  l2circuit.0            0/0/0
BGP-INET           vrf
                                  BGP-INET.inet.0        5/0/0
                                  BGP-INET.iso.0         0/0/0
                                  BGP-INET.inet6.0       0/0/0
BGP-L              vrf
                                  BGP-L.inet.0           5/0/0
                                  BGP-L.iso.0            0/0/0
                                  BGP-L.mpls.0           4/0/0
                                  BGP-L.inet6.0          0/0/0
L2VPN              l2vpn
                                  L2VPN.inet.0           0/0/0
                                  L2VPN.iso.0            0/0/0
                                  L2VPN.inet6.0          0/0/0
                                  L2VPN.l2vpn.0          2/0/0
LDP                vrf
                                  LDP.inet.0             4/0/0
                                  LDP.iso.0              0/0/0
                                  LDP.mpls.0             0/0/0
```

Copyright © 2017, Juniper Networks, Inc.

```
                                    LDP.inet6.0             0/0/0
                                    LDP.l2circuit.0         0/0/0
OSPF                vrf
                                    OSPF.inet.0             7/0/0
                                    OSPF.iso.0              0/0/0
                                    OSPF.inet6.0            0/0/0
RIP                 vrf
                                    RIP.inet.0              6/0/0
                                    RIP.iso.0               0/0/0
                                    RIP.inet6.0             0/0/0
STATIC              vrf
                                    STATIC.inet.0           4/0/0
                                    STATIC.iso.0            0/0/0
                                    STATIC.inet6.0          0/0/0
```


Copyright © 2017, Juniper Networks, Inc.

show route next-hop

List of Syntax
Syntax on page 375
Syntax (EX Series Switches) on page 375

Syntax
show route next-hop *next-hop*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route next-hop *next-hop*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the entries in the routing table that are being sent to the specified next-hop address.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of ouput.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

next-hop—Next-hop address.

Required Privilege Level
view

List of Sample Output
show route next-hop on page 375
show route next-hop detail on page 376
show route next-hop extensive on page 378
show route next-hop terse on page 379

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route next-hop

```
user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 06:26:25
```

Copyright © 2017, Juniper Networks, Inc.

```
                    > to 192.168.71.254 via fxp0.0
172.16.0.0/12       *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16      *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
192.168.102.0/23    *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
207.17.136.0/24     *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
207.17.136.192/32   *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop detail

```
user@host> show route next-hop 192.168.71.254 detail

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
```


Copyright © 2017, Juniper Networks, Inc.

```
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

192.168.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

192.168.102.0/23 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

207.17.136.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

207.17.136.192/32 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
```

Copyright © 2017, Juniper Networks, Inc.

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop extensive

```
user@host> show route next-hop 192.168.71.254 extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

192.168.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```


Copyright © 2017, Juniper Networks, Inc.

```
192.168.102.0/23 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.192/32 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop terse

```
user@host> show route next-hop 192.168.71.254 terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
```

Copyright © 2017, Juniper Networks, Inc.

```
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                       >192.168.71.254
* 10.209.0.0/16      S   5                       >192.168.71.254
* 172.16.0.0/12      S   5                       >192.168.71.254
* 192.168.0.0/16     S   5                       >192.168.71.254
* 192.168.102.0/23   S   5                       >192.168.71.254
* 207.17.136.0/24    S   5                       >192.168.71.254
* 207.17.136.192/32  S   5                       >192.168.71.254

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

show route output

List of Syntax
Syntax on page 381
Syntax (EX Series Switches) on page 381

Syntax
show route output (address *ip-address* | interface *interface-name*)
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route output (address *ip-address* | interface *interface-name*)
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the entries in the routing table learned through static routes and interior gateway protocols that are to be sent out the interface with either the specified IP address or specified name.

To view routes advertised to a neighbor or received from a neighbor for the BGP protocol, use the **show route advertising-protocol bgp** and **show route receive-protocol bgp** commands instead.

Options
address *ip-address*—Display entries in the routing table that are to be sent out the interface with the specified IP address.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

interface *interface-name*—Display entries in the routing table that are to be sent out the interface with the specified name.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route output address on page 382
show route output address detail on page 382
show route output address extensive on page 383
show route output address terse on page 383
show route output interface on page 383
show route output interface detail on page 384
show route output interface extensive on page 384
show route output interface terse on page 384

Copyright © 2017, Juniper Networks, Inc.

Output Fields For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route output address

```
user@host> show route output address 172.16.36.1/24

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.36.0/24         *[Direct/0] 00:19:56
                     > via so-0/1/2.0
                     [OSPF/10] 00:19:55, metric 1
                     > via so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output address detail

```
user@host> show route output address 172.16.36.1 detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
172.16.36.0/24 (2 entries, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via so-0/1/2.0, selected
                State: <Active Int>
                Age: 23:00
                Task: IF
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/1/2.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Age: 22:59      Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

```
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output address extensive

The output for the **show route output address extensive** command is identical to that of the **show route output address detail** command. For sample output, see show route output address detail on page 382.

show route output address terse

```
user@host> show route output address 172.16.36.1 terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 172.16.36.0/24        D   0                      >so-0/1/2.0
                      O  10          1            >so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output interface

```
user@host> show route output interface so-0/1/2.0

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.240/32   *[OSPF/10] 00:13:00, metric 2
                      via so-0/1/2.0
                    > via so-0/3/2.0
10.255.71.241/32   *[OSPF/10] 00:13:10, metric 1
                    > via so-0/1/2.0
172.16.14.0/24        *[OSPF/10] 00:05:11, metric 3
                      to 35.1.1.2 via ge-3/1/0.0
                    > via so-0/1/2.0
                      via so-0/3/2.0
172.16.16.0/24        *[OSPF/10] 00:13:10, metric 2
                    > via so-0/1/2.0
172.16.36.0/24        *[Direct/0] 00:13:21
                    > via so-0/1/2.0
                    [OSPF/10] 00:13:20, metric 1
                    > via so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output interface detail

```
user@host> show route output interface so-0/1/2.0 detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.255.71.240/32 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 2
                Next hop: via so-0/1/2.0
                Next hop: via so-0/3/2.0, selected
                State: <Active Int>
                Age: 14:52      Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (1): 0-KRT
                AS path: I

10.255.71.241/32 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 4
                Next hop: via so-0/1/2.0, selected
                State: <Active Int>
                Age: 15:02      Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (1): 0-KRT
                AS path: I
...
```

show route output interface extensive

The output for the **show route output interface extensive** command is identical to that of the **show route output interface detail** command. For sample output, see show route output interface detail on page 384.

show route output interface terse

```
user@host> show route output interface so-0/1/2.0 terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.71.240/32   O  10          2             so-0/1/2.0
                                                 >so-0/3/2.0
* 10.255.71.241/32   O  10          1            >so-0/1/2.0
* 172.16.14.0/24        O  10          3             35.1.1.2
                                                 >so-0/1/2.0
                                                  so-0/3/2.0
* 172.16.16.0/24        O  10          2            >so-0/1/2.0
* 172.16.36.0/24        D   0                       >so-0/1/2.0
                      O  10          1            >so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

show route protocol

List of Syntax
Syntax on page 386
Syntax (EX Series Switches) on page 386

Syntax
show route protocol *protocol*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route protocol *protocol*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
ospf2 and **ospf3** options introduced in Junos OS Release 9.2.
ospf2 and **ospf3** options introduced in Junos OS Release 9.2 for EX Series switches.
flow option introduced in Junos OS Release 10.0.
flow option introduced in Junos OS Release 10.0 for EX Series switches.

Description
Display the route entries in the routing table that were learned from a particular protocol.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

protocol—Protocol from which the route was learned:

- **access**—Access route for use by DHCP application
- **access-internal**—Access-internal route for use by DHCP application
- **aggregate**—Locally generated aggregate route
- **arp**—Route learned through the Address Resolution Protocol
- **atmvpn**—Asynchronous Transfer Mode virtual private network
- **bgp**—Border Gateway Protocol
- **ccc**—Circuit cross-connect
- **direct**—Directly connected route
- **dvmrp**—Distance Vector Multicast Routing Protocol
- **esis**—End System-to-Intermediate System
- **flow**—Locally defined flow-specification route
- **frr**—Precomputed protection route or backup route used when a link goes down
- **isis**—Intermediate System-to-Intermediate System

Copyright © 2017, Juniper Networks, Inc.

- **ldp**—Label Distribution Protocol
- **l2circuit**—Layer 2 circuit
- **l2vpn**—Layer 2 virtual private network
- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network

NOTE: EX Series switches run a subset of these protocols. See the switch CLI for details.

Required Privilege Level view

List of Sample Output

show route protocol access on page 388
show route protocol access-internal extensive on page 388
show route protocol arp on page 388
show route protocol bgp on page 389
show route protocol bgp detail on page 389
show route protocol bgp detail (Labeled Unicast) on page 390
show route protocol bgp detail (Aggregate Extended Community Bandwidth) on page 390
show route protocol bgp extensive on page 391
show route protocol bgp terse on page 392
show route protocol direct on page 392
show route protocol frr on page 392
show route protocol l2circuit detail on page 393

Copyright © 2017, Juniper Networks, Inc.

show route protocol l2vpn extensive on page 394
show route protocol ldp on page 394
show route protocol ldp extensive on page 395
show route protocol ospf (Layer 3 VPN) on page 396
show route protocol ospf detail on page 397
show route protocol rip on page 397
show route protocol rip detail on page 397
show route protocol ripng table inet6 on page 397
show route protocol static detail on page 398

Output Fields For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route protocol access

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
        *Access-internal Preference: 12
                Next-hop reference count: 200000
                Next hop: 13.160.0.2 via fe-0/0/0.0, selected
                State: <Active Int>
       Age: 36
                Task: RPD Unix Domain Server./var/run/rpd_serv.local
                Announcement bits (1): 0-KRT
                AS path: I
```

show route protocol arp

```
user@host>  show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
```


Copyright © 2017, Juniper Networks, Inc.

```
                      Unusable
20.20.1.4/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.5/32        [ARP/4294967293] 00:04:32, from 20.20.1.1
                      Unusable
20.20.1.6/32        [ARP/4294967293] 00:04:34, from 20.20.1.1
                      Unusable
20.20.1.7/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.8/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.9/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.10/32       [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.11/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
20.20.1.12/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
20.20.1.13/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
...
```

show route protocol bgp

```
user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21       *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
                      AS path: 10458 14203 2914 4788 4788 I
                    > to 192.168.167.254 via fxp0.0
```

show route protocol bgp detail

```
user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24  (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006436
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                State: <Active Ext>
                Local AS:    69 Peer AS: 10458
                Age: 6d 10:42:42        Metric2: 0
                Task: BGP_10458.192.168.69.71+179
                Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
                AS path: 10458 14203 2914 4788 4788 I
                Communities: 2914:410 2914:2403 2914:3400
                Accepted
                Localpref: 100
                Router ID: 207.17.136.192
```

Copyright © 2017, Juniper Networks, Inc.

show route protocol bgp detail (Labeled Unicast)

```
user@host> show route protocol bgp 1.1.1.8/32 detail
inet.0: 45 destinations, 46 routes (45 active, 0 holddown, 0 hidden)
1.1.1.8/32 (2 entries, 2 announced)
State:
*BGP Preference: 1/-101
Next hop type: Indirect, Next hop index: 0
Address: 0xc007f30
Next-hop reference count: 2
Source: 1.1.1.1
Next hop type: Router, Next hop index: 614
Next hop: 20.1.1.2 via ge-0/0/1.0, selected
Label-switched-path lsp1
Label operation: Push 1000126, Push 1000125, Push 1000124, Push 1000123, Push
299872(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 1000126: None; Label 1000125: None; Label 1000124: None;
 Label 1000123: None; Label 299872: None;
Label element ptr: 0xc007860
Label parent element ptr: 0xc0089a0
Label element references: 1
Label element child references: 0
Label element lsp id: 0
Session Id: 0x140
Protocol next hop: 1.1.1.4
Label operation: Push 1000126, Push 1000125, Push 1000124, Push 1000123(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl, prop-ttl
Load balance label: Label 1000126: None; Label 1000125: None; Label 1000124: None;
 Label 1000123: None;
Indirect next hop: 0xae8d300 1048576 INH Session ID: 0x142
State:
Local AS: 5 Peer AS: 5
Age: 22:43 Metric2: 2
Validation State: unverified
Task: BGP_5.1.1.1.1
Announcement bits (2): 0-KRT 7-Resolve tree 2
AS path: I
Accepted
Route Labels: 1000123(top) 1000124 1000125 1000126
Localpref: 100
Router ID: 1.1.1.1
```

show route protocol bgp detail (Aggregate Extended Community Bandwidth)

```
user@host> show route 10.0.2.0 protocol bgp detail
inet.0: 20 destinations, 26 routes (20 active, 0 holddown, 0 hidden)
10.0.2.0/30 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 0
                Address: 0xb618990
                Next-hop reference count: 3
                Source: 10.0.1.1
                Next hop: 10.0.0.2 via ge-0/0/0.0 balance 40%
                Session Id: 0x0
                Next hop: 10.0.1.1 via ge-0/0/1.0 balance 60%, selected
                Session Id: 0x0
                State: <Active Ext>
                Local AS: 65000 Peer AS: 65001
                Age: 20:33
```


Copyright © 2017, Juniper Networks, Inc.

```
                Validation State: unverified
                Task: BGP_65001.10.0.1.1
                Announcement bits (3): 0-KRT 2-BGP_Listen.0.0.0.0+179
3-BGP_RT_Background
                AS path: 65001 I
                Communities: bandwidth:65000:60000000
          Accepted Multipath
                Localpref: 100
                Router ID: 128.49.121.137
         BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 595
                Address: 0xb7a1330
                Next-hop reference count: 9
                Source: 10.0.0.2
                Next hop: 10.0.0.2 via ge-0/0/0.0, selected
                Session Id: 0x141
                State: <NotBest Ext>
                Inactive reason: Not Best in its group - Active preferred
                Local AS: 65000 Peer AS: 65001
                Age: 20:33
                Validation State: unverified
                Task: BGP_65001.10.0.0.2
                AS path: 65001 I
                Communities: bandwidth:65000:40000000
                Accepted MultipathContrib
                Localpref: 100
                Router ID: 128.49.121.132
```

show route protocol bgp extensive

```
user@host> show route protocol bgp 192.168.64.0/21 extensive

inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
    Nexthop: Self
    AS path: [69] 10458 14203 2914 4788 4788 I
    Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4.  Val: 1
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006502
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                State: <Active Ext>
                Local AS:    69 Peer AS: 10458
                Age: 6d 10:44:45        Metric2: 0
                Task: BGP_10458.192.168.69.71+179
                Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
                AS path: 10458 14203 2914 4788 4788 I
                Communities: 2914:410 2914:2403 2914:3400
                Accepted
                Localpref: 100
                Router ID: 207.17.136.192
                Indirect next hops: 1
```

Copyright © 2017, Juniper Networks, Inc.

```
                                    Protocol next hop: 192.168.69.71
                                    Indirect next hop: 8e166c0 342
                                    Indirect path forwarding next hops: 1
                                            Next hop type: Router
                                            Next hop: 192.168.167.254 via fxp0.0
                                    192.168.0.0/16 Originating RIB: inet.0
                                      Node path count: 1
                                      Forwarding nexthops: 1
                                            Nexthop: 192.168.167.254 via fxp0.0
```

show route protocol bgp terse

```
user@host> show route protocol bgp 192.168.64.0/21 terse

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  192.168.64.0/21    B 170        100            >172.16.100.1    10023 21 I
```

show route protocol direct

```
user@host> show route protocol direct

inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.8.0/24          *[Direct/0] 17w0d 10:31:49
                        > via fe-1/3/1.0
10.255.165.1/32     *[Direct/0] 25w4d 04:13:18
                        > via lo0.0
172.16.30.0/24       *[Direct/0] 17w0d 23:06:26
                        > via fe-1/3/2.0
192.168.164.0/22    *[Direct/0] 25w4d 04:13:20
                        > via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
                    *[Direct/0] 25w4d 04:13:21
                        > via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::10:255:165:1/128
                    *[Direct/0] 25w4d 04:13:21
                        > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
                    *[Direct/0] 25w4d 04:13:21
                        > via lo0.0
```

show route protocol frr

```
user@host> show route protocol frr
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

```
cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.3 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.4 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32       *[FRR/200] 00:05:35, from 20.20.1.1
                    > to 20.20.1.5 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32       *[FRR/200] 00:05:37, from 20.20.1.1
                    > to 20.20.1.6 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.7 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.8 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.9 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32      *[FRR/200] 00:05:38, from 20.20.1.1
...
```

show route protocol l2circuit detail

```
user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via ge-2/0/0.0, selected
                Label operation: Pop       Offset: 4
                State: <Active Int>
                Local AS:    99
                Age: 9:52
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

ge-2/0/0.0 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000, Push 100000(top)[0] Offset: -4
                Protocol next hop: 10.245.255.63
                Push 100000 Offset: -4
                 Indirect next hop: 86af0c0 298
                State: <Active Int>
                Local AS:    99
                Age: 9:52
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
```

Copyright © 2017, Juniper Networks, Inc.

```
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    99
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512
```

show route protocol l2vpn extensive

```
user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
        *L2VPN  Preference: 7
                Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
                Label operation: Pop       Offset: 4
                State: <Active Int>
                Local AS:    69
                Age: 7:48
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0.0      /16 -> {indirect(288)}
        *L2VPN  Preference: 7
                Next hop: via so-0/0/1.0, selected
                Label operation: Push 800000 Offset: -4
                Protocol next hop: 10.255.14.220
                Push 800000 Offset: -4
                 Indirect next hop: 85142a0 288
                State: <Active Int>
                Local AS:    69
                Age: 7:48
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:69:1 Layer2-info: encaps:PPP,
                control flags:2, mtu: 0
```

show route protocol ldp

```
user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

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

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0

private1__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100080             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Swap 100000
```

show route protocol ldp extensive

```
user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *LDP    Preference: 9
                Next-hop reference count: 3
                Next hop: via t1-4/0/0.0, selected
                Label operation: Push 100000
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
                AS path: I

192.168.17.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *LDP    Preference: 9
                Next-hop reference count: 3
                Next hop: via t1-4/0/0.0, selected
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
```

Copyright © 2017, Juniper Networks, Inc.

```
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                Label operation: Pop
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                Label operation: Swap 100000
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 192.168.16.1/32
```

show route protocol ospf (Layer 3 VPN)

```
user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30       *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.39.1.8/30        [OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
10.255.14.171/32   *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.255.14.179/32   *[OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
172.16.233.5/32       *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30       [OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
10.255.14.173/32   *[OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
172.16.233.5/32       *[OSPF/10] 20:26:20, metric 1
```


Copyright © 2017, Juniper Networks, Inc.

show route protocol ospf detail

```
user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
         OSPF   Preference:  10
                Nexthop: via so-0/2/2.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Age: 6:25       Metric: 1
                Area: 0.0.0.0
                Task: VPN-AB-OSPF
                AS path: I
                Communities: Route-Type:0.0.0.0:1:0

...
```

show route protocol rip

```
user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32   *[RIP/100] 20:24:34, metric 2
                    > to 10.39.1.22 via t3-0/2/2.0
172.16.233.9/32      *[RIP/100] 00:03:59, metric 1
```

show route protocol rip detail

```
user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
        *RIP    Preference: 100
                Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
                State: <Active Int>
                Age: 20:25:02   Metric: 2
                Task: VPN-AB-RIPv2
                Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
                AS path: I
                Route learned from 10.39.1.22 expires in 96 seconds
```

show route protocol ripng table inet6

```
user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128        *[RIPng/100] 02:13:33, metric 2
```

Copyright © 2017, Juniper Networks, Inc.

```
                     > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128        *[RIPng/100] 02:13:33, metric 2
                     > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128        *[RIPng/100] 02:13:33, metric 2
                     > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128        *[RIPng/100] 02:13:33, metric 2
                     > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

show route protocol static detail

```
user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.13.10.0/23 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```


Copyright © 2017, Juniper Networks, Inc.

show route receive-protocol

List of Syntax
Syntax on page 399
Syntax (EX Series Switches) on page 399

Syntax
show route receive-protocol *protocol neighbor-address*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)

Syntax (EX Series Switches)
show route receive-protocol *protocol neighbor-address*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

protocol neighbor-address—Protocol transmitting the route (**bgp**, **dvmrp**, **msdp**, **pim**, **rip**, or **ripng**) and address of the neighboring router from which the route entry was received.

Additional Information
The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes.

Required Privilege Level
view

List of Sample Output
show route receive-protocol bgp on page 402
show route receive-protocol bgp extensive on page 402
show route receive-protocol bgp table extensive on page 403
show route receive-protocol bgp logical-system extensive on page 403
show route receive-protocol bgp detail (Layer 2 VPN) on page 404
show route receive-protocol bgp extensive (Layer 2 VPN) on page 404
show route receive-protocol bgp (Layer 3 VPN) on page 405
show route receive-protocol bgp detail (Layer 3 VPN) on page 405
show route receive-protocol bgp detail (Long-Lived Graceful Restart) on page 406
show route receive-protocol bgp detail (Labeled Unicast) on page 406
show route receive-protocol bgp extensive (Layer 3 VPN) on page 407
Show route receive protocol (Segment Routing Traffic Engineering on page 407

Copyright © 2017, Juniper Networks, Inc.

Output Fields Table 25 on page 400 describes the output fields for the **show route receive-protocol** command. Output fields are listed in the approximate order in which they appear.

Table 25: show route receive-protocol Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example, inet.0.	All levels
number destinations	Number of destinations for which there are routes in the routing table.	All levels
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** • **holddown** (routes that are in pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)	All levels
Prefix	Destination prefix.	none **brief**
MED	Multiple exit discriminator value included in the route.	none **brief**
destination-prefix (entry, announced)	Destination prefix. The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination.	**detail extensive**
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.	**detail extensive**
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table	**detail extensive**
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.	**detail extensive**
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.	**detail extensive**
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 25: show route receive-protocol Output Fields *(continued)*

Field Name	Field Description	Level of Output
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either an RSVP or an LDP label-switched path (LSP) tunnel.	**detail extensive**
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	All levels
Localpref or Lclpref	Local preference value included in the route.	All levels
AS path	Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.	All levels
Route Labels	Stack of labels carried in the BGP route update.	**detail extensive**
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	**detail extensive**
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	**detail extensive**
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.	**detail extensive**
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	**detail extensive**
Attrset AS	Number, local preference, and path of the AS that originated the route. These values are stored in the **Attrset** attribute at the originating routing device.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 25: show route receive-protocol Output Fields *(continued)*

Field Name	Field Description	Level of Output
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	**detail extensive**
control flags	Control flags: **none** or **Site Down**.	**detail extensive**
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	**detail extensive**

Sample Output

show route receive-protocol bgp

```
user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix                  Next hop              MED     Lclpref    AS path
10.22.1.0/24            10.255.245.215        0       100        I
10.22.2.0/24            10.255.245.215        0       100        I
```

show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
Prefix             Next hop               MED    Lclpref AS path
172.16.1.0/24 (1 entry, 1 announced)
     Next hop: 10.0.50.3
     Localpref: 100
     AS path: I <Orginator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.45
172.16.163.0/16 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.68
172.16.164.0/16 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.45
172.16.195.0/24 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix             Next hop               MED    Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Next hop               MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Next hop               MED    Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

show route receive-protocol bgp table extensive

```
user@host> show route receive-protocol bgp 207.17.136.192 table inet.0 66.117.68.0/24 extensive
inet.0: 227315 destinations, 227316 routes (227302 active, 0 holddown, 13 hidden)
* 66.117.63.0/24 (1 entry, 1 announced)
     Nexthop: 207.17.136.29
     Localpref: 100
     AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
     AS path: AS4 PA[2]: 33437 393219
     AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
     Communities: 2914:420
```

show route receive-protocol bgp logical-system extensive

```
user@host> show route receive-protocol bgp 10.0.0.9 logical-system PE4 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
* 10.0.0.0/30 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 10.0.0.4/30 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

  10.0.0.8/30 (2 entries, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 10.9.9.1/32 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 10.100.1.1/32 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 172.16.44.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300096
     Nexthop: 10.0.0.9
     AS path: 13979 I
     AIGP: 203

* 172.16.55.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300112
     Nexthop: 10.0.0.9
     AS path: 13979 7018 I
     AIGP: 25
```

Copyright © 2017, Juniper Networks, Inc.

```
* 172.16.66.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300144
     Nexthop: 10.0.0.9
     AS path: 13979 7018 I

* 172.16.99.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300160
     Nexthop: 10.0.0.9
     AS path: 13979 7018 I
```

show route receive-protocol bgp detail (Layer 2 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix             Nexthop               MED    Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
```

show route receive-protocol bgp extensive (Layer 2 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix             Nexthop               MED    Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
     Route Distinguisher: 10.255.245.35:1
```


Copyright © 2017, Juniper Networks, Inc.

```
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
```

show route receive-protocol bgp (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171
      inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      10.255.14.175/32   10.255.14.171                    100 2 I
      10.255.14.179/32   10.255.14.171           2        100 I
      VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      10.255.14.175/32   10.255.14.171                    100 2 I
      10.255.14.177/32   10.255.14.171                    100 I
      iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      10.255.14.171:300:10.255.14.177/32
                         10.255.14.171                    100 I
      10.255.14.171:100:10.255.14.179/32
                         10.255.14.171           2        100 I
      10.255.14.171:200:10.255.14.175/32
                         10.255.14.171                    100 2 I
```

show route receive-protocol bgp detail (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
     Route Distinguisher: 10.255.14.176:2
     VPN Label: 101264
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
```

Copyright © 2017, Juniper Networks, Inc.

```
            AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
     Route Distinguisher: 10.255.14.176:2
     VPN Label: 101280
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
         AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
     Route Distinguisher: 10.255.14.174:2
     VPN Label: 101264
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
         AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
     Route Distinguisher: 10.255.14.174:2
     VPN Label: 101280
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
         AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp detail (Long-Lived Graceful Restart)

```
user@host> show route receive-protocol bgp 10.4.12.11 detail

  bgp.l2vpn.0: 38 destinations, 39 routes (37 active, 0 holddown, 1 hidden)
  * 172.16.1.4:100:172.16.1.4/96 AD (1 entry, 1 announced)
    Accepted LongLivedStale LongLivedStaleImport
    Nexthop: 10.4.12.11
    Localpref: 100
    AS path: I
```

show route receive-protocol bgp detail (Labeled Unicast)

```
user@host> show route receive-protocol bgp 1.1.1.1 detail
inet.0: 45 destinations, 46 routes (45 active, 0 holddown, 0 hidden)
* 1.1.1.8/32 (2 entries, 2 announced)
Accepted
Route Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: 1.1.1.4
Localpref: 100
AS path: I
Entropy label capable, next hop field matches route next hop

inet.3: 15 destinations, 21 routes (6 active, 0 holddown, 14 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

inet6.0: 26 destinations, 28 routes (26 active, 0 holddown, 0 hidden)

* 100::1/128 (2 entries, 2 announced)
Accepted
Route Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: ::ffff:1.1.1.4
Localpref: 100
AS path: I

inet6.3: 22 destinations, 23 routes (22 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp extensive (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     172.16.1.0/24 (1 entry, 1 announced)
          Nexthop: 10.0.50.3
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     172.16.163.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68
     172.16.164.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     172.16.195.0/24 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68
     inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

Show route receive protocol (Segment Routing Traffic Engineering

```
show route receive protocol bgp 10.1.1.4
bgp.inetcolor.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

* 50-4.4.4.4-1234(sr6)/96 (1 entry, 0 announced)
     Import Accepted
```

Copyright © 2017, Juniper Networks, Inc.

```
     Distinguisher: 50
     Color: 1234
     Nexthop: 10.1.1.4
     Localpref: 100
     AS path: 3 I
     Communities: target:1.1.1.1:1


inetcolor.0: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
* 4.4.4.4-1234(c6)/64 (1 entry, 1 announced)
     Import Accepted
     Color: 1234
     Nexthop: 10.1.1.4
     Localpref: 100
     AS path: 3 I
     Communities: target:1.1.1.1:1


user@host# run show route receive-protocol bgp 5001:1::4
bgp.inet6color.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

* 50-2001:1::4-1234(sr6)/192 (1 entry, 0 announced)
     Import Accepted
     Distinguisher: 50
     Color: 1234
     Nexthop: ::ffff:1.1.1.4
     Localpref: 100
     AS path: 3 I
     Communities: target:1.1.1.1:1


inet6color.0: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
* 2001::5-1234(c6)/160 (1 entry, 1 announced)
     Import Accepted
     Color: 1234
     Nexthop: ::ffff:1.1.1.5
     Localpref: 100
     AS path: 3 I
     Communities: target:2:1
```


Copyright © 2017, Juniper Networks, Inc.

show route table

List of Syntax

Syntax on page 409
Syntax (EX Series Switches, QFX Series Switches) on page 409

Syntax

show route table *routing-table-name*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches, QFX Series Switches)

show route table *routing-table-name*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.
Show route table evpn statement introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.

Description

Display the route entries in a particular routing table.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—Display route entries for all routing tables whose names begin with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route table inet** command).

Required Privilege Level

view

Related Documentation

- *show route summary*

List of Sample Output

show route table bgp.l2.vpn on page 420
show route table bgp.l3vpn.0 on page 420
show route table bgp.l3vpn.0 detail on page 421
show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 422
show route table bgp.evpn.0 on page 422
show route table evpna.evpn.0 on page 423
show route table inet.0 on page 423
show route table inet.3 on page 424
show route table inet.3 protocol ospf on page 424
show route table inet6.0 on page 424
show route table inet6.3 on page 424

Copyright © 2017, Juniper Networks, Inc.

show route table inetflow detail on page 425
show route table lsdist.0 extensive on page 425
show route table l2circuit.0 on page 427
show route table mpls on page 427
show route table mpls extensive on page 427
show route table mpls.0 on page 428
show route table mpls.0 detail (PTX Series) on page 429
show route table mpls.0 ccc ge-0/0/1.1004 detail on page 429
show route table mpls.0 protocol evpn on page 430
show route table mpls.0 protocol ospf on page 436
show route table mpls.0 extensive (PTX Series) on page 437
show route table mpls.0 (RSVP Route—Transit LSP) on page 437
show route table vpls_1 detail on page 438
show route table vpn-a on page 438
show route table vpn-a.mdt.0 on page 438
show route table VPN-A detail on page 439
show route table VPN-AB.inet.0 on page 439
show route table VPN_blue.mvpn-inet6.0 on page 440
show route table vrf1.mvpn.0 extensive on page 440
show route table inetflow detail on page 440
show route table bgp.evpn.0 extensive |no-more (EVPN) on page 443

Output Fields Table 15 on page 261 describes the output fields for the **show route table** command. Output fields are listed in the approximate order in which they appear.

Table 26: show route table Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).

Restart complete	All protocols have restarted for this routing table.
	
	Restart state:
	
	• **Pending:*protocol-name***—List of protocols that have not yet completed graceful restart for this routing table.
	• **Complete**—All protocols have restarted for this routing table.
	
	For example, if the output shows-
	
	• `LDP.inet.0 : 5 routes (4 active, 1 holddown, 0 hidden)`
	`Restart Pending: OSPF LDP VPN`
	
	This indicates that **OSPF**, **LDP**, and **VPN** protocols did not restart for the **LDP.inet.0** routing table.
	• `vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)`
	`Restart Complete`
	
	This indicates that all protocols have restarted for the **vpls_1.l2vpn.0** routing table.

number destinations	Number of destinations for which there are routes in the routing table.

Copyright © 2017, Juniper Networks, Inc.

Table 26: show route table Output Fields *(continued)*

Field Name	Field Description
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)

route-destination (entry, announced)	Route destination (for example:10.0.0.1/24). The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.
	• ***inclusive multicast Ethernet tag route***—Type of route destination represented by (for example, 3:100.100.100.10:100::0::10::100.100.100.10/384):
	  ◦ *route distinguisher*—(8 octets) Route distinguisher (RD) must be the RD of the EVPN instance (EVI) that is advertising the NLRI.
	  ◦ *Ethernet tag ID*—(4 octets) Identifier of the Ethernet tag. Can set to 0 or to a valid Ethernet tag value.
	  ◦ *IP address length*—(1 octet) Length of IP address in bits.
	  ◦ *originating router's IP address*—(4 or 16 octets) Must set to the provider edge (PE) device's IP address. This address should be common for all EVIs on the PE device, and may be the PE device's loopback address.

label stacking	**(**Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).

Copyright © 2017, Juniper Networks, Inc.

Table 26: show route table Output Fields *(continued)*

Field Name	Field Description
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • ***—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 18 on page 305.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.

Copyright © 2017, Juniper Networks, Inc.

Table 26: show route table Output Fields *(continued)*

Field Name	Field Description
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 19 on page 307.
Local AS	AS number of the local routing devices.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric*n*	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.
Task	Name of the protocol that has added the route.
Announcement bits	The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the kernel routing table (KRT) for installing the route into the Packet Forwarding Engine, to a resolve tree, a Layer 2 VC, or even a VPN. For example, ***n*-Resolve inet** indicates that the specified route is used for route resolution for next hops found in the routing table. • ***n***—An index used by Juniper Networks customer support only.

Copyright © 2017, Juniper Networks, Inc.

Table 26: show route table Output Fields *(continued)*

Field Name	Field Description
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Indicates point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, indicates the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, indicates the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.

Copyright © 2017, Juniper Networks, Inc.

Table 26: show route table Output Fields *(continued)*

Field Name	Field Description
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 20 on page 309 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or **Site Down**.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag.
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.

Copyright © 2017, Juniper Networks, Inc.

Table 26: show route table Output Fields *(continued)*

Field Name	Field Description
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 18 on page 305 describes all possible values for the Next-hop Types output field.

Table 27: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.

Copyright © 2017, Juniper Networks, Inc.

Table 27: Next-hop Types Output Field Values *(continued)*

Next-Hop Type	Description
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.
Router	A specific node or set of nodes to which the routing device forwards packets that match the route prefix. To qualify as a next-hop type router, the route must meet the following criteria: • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 19 on page 307 describes all possible values for the State output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).

Table 28: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.

Copyright © 2017, Juniper Networks, Inc.

Table 28: State Output Field Values *(continued)*

Value	Description
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.

Copyright © 2017, Juniper Networks, Inc.

Table 28: State Output Field Values *(continued)*

Value	Description
Pending	Route pending because of a hold-down configured on another route.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available.
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 20 on page 309 describes the possible values for the Communities output field.

Table 29: Communities Output Field Values

Value	Description
area-number	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is **0**. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: ***local AS number:link-bandwidth-number***	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
link-bandwidth-number	Link-bandwidth number: from **0** through **4,294,967,295** (bytes per second).
local AS number	Local AS number: from **1** through **65,535**.

Copyright © 2017, Juniper Networks, Inc.

Table 29: Communities Output Field Values *(continued)*

Value	Description
options	1 byte. Currently this is only used if the route type is **5** or **7**. Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	1 byte, encoded as **1** or **2** for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); **3** for summary routes; **5** for external routes (area number must be**0**); **7** for NSSA routes; or 1**29** for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x8000**. The format is ***area-number:ospf-route-type:options***.
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x0306**. The format is ***area-number:ospf-route-type:options***.
target	Defines which VPN the route participates in; **target** has the format ***32-bit IP address:16-bit number***. For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between **0x1** and **0x7fff**. This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above **0x8000**. This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route table bgp.l2.vpn

```
user@host> show route table bgp.l2.vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
                   *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
      bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.255.71.15:100:10.255.71.17/32
                         *[BGP/170] 00:03:59, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.15:200:10.255.71.18/32
```


Copyright © 2017, Juniper Networks, Inc.

```
                                *[BGP/170] 00:03:59, MED 1, localpref 100, from
      10.255.71.15
                                  AS path: I
                                > via so-2/1/0.0, Push 100021, Push 100011(top)
```

show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:172.16.4.0/8 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
         Route Distinguisher: 10.255.245.12:1
         Source: 10.255.245.12
         Next hop: 192.168.208.66 via fe-0/0/0.0, selected
         Label operation: Push 182449
         Protocol next hop: 10.255.245.12
         Push 182449
         Indirect next hop: 863a630 297
         State: <Active Int Ext>
         Local AS:    35 Peer AS:    35
         Age: 12:19      Metric2: 1
         Task: BGP_35.10.255.245.12+179
         Announcement bits (1): 0-BGP.0.0.0.0+179
         AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

         Communities: 2914:420 target:11111:1 origin:56:78
         VPN Label: 182449
         Localpref: 100
          Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
         Route Distinguisher: 10.255.245.12:1
         Source: 10.255.245.12
         Next hop: 192.168.208.66 via fe-0/0/0.0, selected
         Label operation: Push 182465
         Protocol next hop: 10.255.245.12
         Push 182465
         Indirect next hop: 863a8f0 305
         State: <Active Int Ext>
         Local AS:    35 Peer AS:    35
         Age: 12:19      Metric2: 1
         Task: BGP_35.10.255.245.12+179
         Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
                Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                VPN Label: 182465
                Localpref: 100
                Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
         Route Distinguisher: 10.255.245.12:1
         Source: 10.255.245.12
         Next hop: 192.168.208.66 via fe-0/0/0.0, selected
         Label operation: Push 182465
         Protocol next hop: 10.255.245.12
         Push 182465
         Indirect next hop: 86bd210 330
         State: <Active Int Ext>
```

Copyright © 2017, Juniper Networks, Inc.

```
                Local AS:    35 Peer AS:    35
                Age: 12:19      Metric2: 1
                Task: BGP_35.10.255.245.12+179
                Announcement bits (1): 0-BGP.0.0.0.0+179
                AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

                6496 I
                Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                VPN Label: 182465
                Localpref: 100
                Router ID: 10.255.245.12

10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
                Route Distinguisher: 10.255.245.12:1
                Source: 10.255.245.12
                Next hop: 192.168.208.66 via fe-0/0/0.0, selected
                Label operation: Push 182465
                Protocol next hop: 10.255.245.12
                Push 182465
                Indirect next hop: 86bd210 330
                State: <Active Int Ext>
                Local AS:    35 Peer AS:    35
                Age: 12:19      Metric2: 1
                Task: BGP_35.10.255.245.12+179
                Announcement bits (1): 0-BGP.0.0.0.0+179
                AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

                6496 I
                Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                VPN Label: 182465
                Localpref: 100
```

show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```
user@host> show route table bgp.rtarget.o
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100:100:100/96
                   *[RTarget/5] 00:03:14
                      Type Proxy
                        for 10.255.165.103
                        for 10.255.166.124
                      Local
```

show route table bgp.evpn.0

```
user@host> show route table bgp.evpn.0
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:100.100.100.2:100::0::00:26:88:5f:67:b0/304
                   *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.2:100::0::00:51:51:51:51:51/304
                   *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
```


Copyright © 2017, Juniper Networks, Inc.

```
2:100.100.100.3:100::0::00:52:52:52:52:52/304
                   *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                      AS path: I, validation-state: unverified
                    > to 100.64.13.3 via ge-2/0/8.0, label-switched-path R0toR2
2:100.100.100.3:100::0::a8:d0:e5:5b:01:c8/304
                   *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                      AS path: I, validation-state: unverified
                    > to 100.64.13.3 via ge-2/0/8.0, label-switched-path R0toR2
3:100.100.100.2:100::1000::100.100.100.2/304
                   *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
3:100.100.100.2:100::2000::100.100.100.2/304
                   *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
```

show route table evpna.evpn.0

```
user@host> show route table evpna.evpn.0
evpna.evpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

3:100.100.100.10:100::0::10::100.100.100.10/384
                   *[EVPN/170] 01:37:09
                      Indirect
3:100.100.100.2:100::2000::100.100.100.2/304
                   *[EVPN/170] 01:37:12
                      Indirect
```

show route table inet.0

```
user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:51:57
                    > to 172.16.5.254 via fxp0.0
10.0.0.1/32         *[Direct/0] 00:51:58
                    > via at-5/3/0.0
10.0.0.2/32         *[Local/0] 00:51:58
                      Local
10.12.12.21/32     *[Local/0] 00:51:57
                      Reject
10.13.13.13/32     *[Direct/0] 00:51:58
                    > via t3-5/2/1.0
10.13.13.14/32     *[Local/0] 00:51:58
                      Local
10.13.13.21/32     *[Local/0] 00:51:58
                      Local
10.13.13.22/32     *[Direct/0] 00:33:59
                    > via t3-5/2/0.0
127.0.0.1/32        [Direct/0] 00:51:58
                    > via lo0.0
10.222.5.0/24     *[Direct/0] 00:51:58
                    > via fxp0.0
10.222.5.81/32    *[Local/0] 00:51:58
                      Local
```

Copyright © 2017, Juniper Networks, Inc.

show route table inet.3

```
user@host> show route table inet.3
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32        *[LDP/9] 00:25:43, metric 10, tag 200
                      to 10.2.94.2 via lt-1/2/0.49
                    > to 10.2.3.2 via lt-1/2/0.23
```

show route table inet.3 protocol ospf

```
user@host> show route table inet.3 protocol ospf
inet.3: 9 destinations, 18 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.20/32         [L-OSPF/10] 1d 00:00:56, metric 2
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800020
                     to 10.0.6.60 via lt-1/2/0.12, Push 800020, Push 800030(top)
1.1.1.30/32         [L-OSPF/10] 1d 00:01:01, metric 3
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800030
                      to 10.0.6.60 via lt-1/2/0.12, Push 800030
1.1.1.40/32         [L-OSPF/10] 1d 00:01:01, metric 4
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800040
                      to 10.0.6.60 via lt-1/2/0.12, Push 800040
1.1.1.50/32         [L-OSPF/10] 1d 00:01:01, metric 5
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800050
                      to 10.0.6.60 via lt-1/2/0.12, Push 800050
1.1.1.60/32         [L-OSPF/10] 1d 00:01:01, metric 6
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800060
                      to 10.0.6.60 via lt-1/2/0.12, Pop
```

show route table inet6.0

```
user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0
```

show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
                   *[LDP/9] 00:00:22, metric 1
                    > via so-1/0/0.0
::10.255.245.196/128
                   *[LDP/9] 00:00:08, metric 1
                    > via so-1/0/0.0, Push 100008
```


Copyright © 2017, Juniper Networks, Inc.

show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 2
                State: <Active Ext>
                Local AS: 64502 Peer AS: 64500
                Age: 4
                Task: BGP_64500.10.12.99.5+3792
                Announcement bits (1): 0-Flow
                AS path: 64500 I
                Communities: traffic-rate:0:0
                Validation state: Accept, Originator: 10.12.99.5
                Via: 10.12.44.0/24, Active
                Localpref: 100
                Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
        *Flow   Preference: 5
                Next-hop reference count: 2
                State: <Active>
                Local AS: 64502
                Age: 6:30
                Task: RT Flow
                Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
                AS path: I
                Communities: 1:1
```

show route table lsdist.0 extensive

```
user@host> show route table lsdist.0 extensive
lsdist.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
NODE { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 ISIS-L1:0 }/1152
 (1 entry, 1 announced)
TSI:
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xa62f378 (adv_entry)
   Advertised metrics:
     Nexthop: Self
     Localpref: 100
     AS path: [4170512532] I
     Communities:
Path NODE { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 ISIS-L1:0 }
 Vector len 4.  Val: 0
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0x95dfc64
                Next-hop reference count: 9
                State: <Active NotInstall>
                Local AS: 4170512532
                Age: 6:05
                Validation State: unverified
                Task: IS-IS
                Announcement bits (1): 0-BGP_RT_Background
                AS path: I
                IPv4 Router-ids:
                   128.220.11.197
                Area membership:
```

Copyright © 2017, Juniper Networks, Inc.

```
                47 00 05 80 ff f8 00 00 00 01 08 00 01
             SPRING-Capabilities:                 - SRGB block [Start: 800000,
 Range: 256, Flags: 0xc0]
              SPRING-Algorithms:
                 - Algo: 0
 LINK { Local { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 }.{
IPv4:8.65.1.105 } Remote { AS:4170512532 BGP-LS ID:4170512532 ISO:4284.3300.5067)
TSI:
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xa62f3cc (adv_entry)
    Advertised metrics:
      Nexthop: Self
      Localpref: 100
      AS path: [4170512532] I
      Communities:
Path LINK { Local { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 }.{
 IPv4:8.65.1.105 } Remote { AS:4170512532 BGP-LS ID:4170512532 ISO:4284.33000
      *IS-IS  Preference: 15
               Level: 1
               Next hop type: Fictitious, Next hop index: 0
               Address: 0x95dfc64
               Next-hop reference count: 9
               State: <Active NotInstall>
               Local AS: 4170512532
               Age: 6:05
               Validation State: unverified
               Task: IS-IS
               Announcement bits (1): 0-BGP_RT_Background
               AS path: I
               Color: 32768
               Maximum bandwidth: 1000Mbps
               Reservable bandwidth: 1000Mbps
               Unreserved bandwidth by priority:
                  0   1000Mbps
                  1   1000Mbps
                  2   1000Mbps
                  3   1000Mbps
                  4   1000Mbps
                  5   1000Mbps
                  6   1000Mbps
                  7   1000Mbps
                Metric: 10
                TE Metric: 10
                LAN IPV4 Adj-SID -  Label: 299776, Flags: 0x30,
Weight: 0, Nbr: 10.220.1.83

PREFIX { Node { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 } {
IPv4:128.220.11.197/32 } ISIS-L1:0 }/1152 (1 entry, 1 announced) TSI: Page 0 idx
 0, (group ibgp type Internal) Type 1 val 0xa62f43c (adv_entry)
    Advertised metrics:
      Nexthop: Self
      Localpref: 100
      AS path: [4170512532] I
      Communities:
Path PREFIX { Node { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 }
{ IPv4:128.220.11.197/32 } ISIS-L1:0 } Vector len 4.  Val: 0
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0x95dfc64
                Next-hop reference count: 9
                State:<Active NotInstall>
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Local AS: 4170512532
                    Age: 6:05
                    Validation State: unverified
                    Task: IS-IS
                    Announcement bits (1): 0-BGP_RT_Background
                    AS path: I
                                 Prefix SID: 67, Flags: 0x40, Algo: 0
```

show route table l2circuit.0

```
user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
10.1.1.195:CtrlWord:1:2:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
```

show route table mpls

```
user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1                  *[MPLS/0] 00:13:55, metric 1
                      Receive
2                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1024               *[VPN/0] 00:04:18
                      to table red.inet.0, Pop
```

show route table mpls extensive

```
user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
        *LDP    Preference:   9
                Next hop:  via so-1/0/0.0, selected
                Pop
                State: <Active Int>
                Age: 29:50      Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 10.0.0.194/32
```

Copyright © 2017, Juniper Networks, Inc.

show route table mpls.0

```
user@host> show route table mpls.0
mpls.0: 18 destinations, 19 routes (18 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 11:39:56, metric 1
                      to table inet.0
0(S=0)             *[MPLS/0] 11:39:56, metric 1
                      to table mpls.0
1                  *[MPLS/0] 11:39:56, metric 1
                      Receive
2                  *[MPLS/0] 11:39:56, metric 1
                      to table inet6.0
2(S=0)             *[MPLS/0] 11:39:56, metric 1
                      to table mpls.0
13                 *[MPLS/0] 11:39:56, metric 1
                      Receive
303168             *[EVPN/7] 11:00:49, routing-instance pbbn10, route-type
Ingress-MAC, ISID 0
                      to table pbbn10.evpn-mac.0
303184             *[EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 1000
                      to table pbbn10.evpn-mac.0
                    [EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 2000
                      to table pbbn10.evpn-mac.0
303264             *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-IM, ISID 1000
                     > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303280             *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-IM, ISID 2000
                     > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303328             *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                     > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303344             *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                     > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303360             *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:26:88:5f:67:b0
                     > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303376             *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:51:51:51:51:51
                     > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303392             *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                     > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303408             *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                     > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303424             *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC a8:d0:e5:5b:01:c8
                     > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303440             *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:52:52:52:52:52
                     > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
```


Copyright © 2017, Juniper Networks, Inc.

show route table mpls.0 detail (PTX Series)

```
user@host> show route table mpls.0 detail
ge-0/0/2.600 (1 entry, 1 announced)
        *L2VPN  Preference: 7
                Next hop type: Indirect
                Address: 0x9438f34
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 567
                Next hop: 10.0.0.1 via ge-0/0/1.0, selected
                Label operation: Push 299808
                Label TTL action: prop-ttl
                Load balance label: Label 299808:None;
                Session Id: 0x1
                Protocol next hop: 10.255.255.1
                Label operation: Push 299872 Offset: 252
                Label TTL action: no-prop-ttl
                Load balance label: Label 299872:Flow label PUSH;
                Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                State: <Active Int>
                Age: 21         Metric2: 1
                Validation State: unverified
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 2-Common L2 VC
                AS path: I
```

show route table mpls.0 ccc ge-0/0/1.1004 detail

```
user@host>show route table mpls.0 ccc ge-0/0/1.1004 detail
mpls.0: 121 destinations, 121 routes (121 active, 0 holddown, 0 hidden)
ge-0/0/1.1004 (1 entry, 1 announced)
        *EVPN   Preference: 7
                Next hop type: List, Next hop index: 1048577
                Address: 0xdc14770
                Next-hop reference count: 3
                Next hop: ELNH Address 0xd011e30
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0xd011e30
                    Next-hop reference count: 3
                    Protocol next hop: 100.100.100.1
                    Label operation: Push 301952
                    Composite next hop: 0xd011dc0 754 INH Session ID: 0x146
                    Indirect next hop: 0xb69a890 1048615 INH Session ID: 0x146
                        Next hop type: Router, Next hop index: 735
                        Address: 0xd00e530
                        Next-hop reference count: 23
                        Next hop: 100.46.1.2 via ge-0/0/5.0
                        Label-switched-path pe4_to_pe1
                        Label operation: Push 300320
                        Label TTL action: prop-ttl
                        Load balance label: Label 300320: None;
                        Label element ptr: 0xd00e580
                        Label parent element ptr: 0x0
                        Label element references: 18
                        Label element child references: 16
                        Label element lsp id: 5
                Next hop: ELNH Address 0xd012070
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0xd012070
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Next-hop reference count: 3
                    Protocol next hop: 100.100.100.2
                    Label operation: Push 301888
                    Composite next hop: 0xd012000 755 INH Session ID: 0x143
                    Indirect next hop: 0xb69a9a0 1048641 INH Session ID: 0x143
                        Next hop type: Router, Next hop index: 716
                        Address: 0xd00e710
                        Next-hop reference count: 23
                        Next hop: 100.46.1.2 via ge-0/0/5.0
                        Label-switched-path pe4_to_pe2
                        Label operation: Push 300304
                        Label TTL action: prop-ttl
                        Load balance label: Label 300304: None;
                        Label element ptr: 0xd00e760
                        Label parent element ptr: 0x0
                        Label element references: 15
                        Label element child references: 13
                        Label element lsp id: 6
                Next hop: ELNH Address 0xd0121f0, selected
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0xd0121f0
                    Next-hop reference count: 3
                    Protocol next hop: 100.100.100.3
                    Label operation: Push 301984
                    Composite next hop: 0xd012180 756 INH Session ID: 0x145
                    Indirect next hop: 0xb69aab0 1048642 INH Session ID: 0x145
                        Next hop type: Router, Next hop index: 801
                        Address: 0xd010ed0
                        Next-hop reference count: 32
                        Next hop: 100.46.1.2 via ge-0/0/5.0
                        Label-switched-path pe4_to_pe3
                        Label operation: Push 300336
                        Label TTL action: prop-ttl
                        Load balance label: Label 300336: None;
                        Label element ptr: 0xd0108c0
                        Label parent element ptr: 0x0
                        Label element references: 22
                        Label element child references: 20
                        Label element lsp id: 7
                State: < Active Int >
                Age: 2:06:50
                Validation State: unverified
                Task: evpn global task
                Announcement bits (1): 1-KRT
                AS path: I
```

show route table mpls.0 protocol evpn

```
user@host>show route table mpls.0 protocol evpn
mpls.0: 121 destinations, 121 routes (121 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299872             *[EVPN/7] 02:30:58, routing-instance mhevpn, route-type
Ingress-IM, vlan-id 10
                      to table mhevpn.evpn-mac.0
300016             *[EVPN/7] 02:30:38, routing-instance VS-1, route-type
Ingress-IM, vlan-id 110
                      to table VS-1.evpn-mac.0
300032             *[EVPN/7] 02:30:38, routing-instance VS-1, route-type
Ingress-IM, vlan-id 120
                      to table VS-1.evpn-mac.0
```


Copyright © 2017, Juniper Networks, Inc.

```
300048             *[EVPN/7] 02:30:38, routing-instance VS-1, route-type
Ingress-IM, vlan-id 130
                      to table VS-1.evpn-mac.0
300064             *[EVPN/7] 02:30:38, routing-instance VS-2, route-type
Ingress-IM, vlan-id 210
                      to table VS-2.evpn-mac.0
300080             *[EVPN/7] 02:30:38, routing-instance VS-2, route-type
Ingress-IM, vlan-id 220
                      to table VS-2.evpn-mac.0
300096             *[EVPN/7] 02:30:38, routing-instance VS-2, route-type
Ingress-IM, vlan-id 230
                      to table VS-2.evpn-mac.0
300112             *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:44:44:44:44:44:44:44:44:44
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300128             *[EVPN/7] 02:29:22, routing-instance mhevpn, route-type
Ingress-Aliasing
                      to table mhevpn.evpn-mac.0
300144             *[EVPN/7] 02:27:06, routing-instance VS-1, route-type
Egress-MAC, ESI 00:44:44:44:44:44:44:44:44:44
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300160             *[EVPN/7] 02:29:22, routing-instance VS-1, route-type
Ingress-Aliasing
                      to table VS-1.evpn-mac.0
300176             *[EVPN/7] 02:27:07, routing-instance VS-2, route-type
Egress-MAC, ESI 00:44:44:44:44:44:44:44:44:44
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300192             *[EVPN/7] 02:29:22, routing-instance VS-2, route-type
Ingress-Aliasing
                      to table VS-2.evpn-mac.0
300208             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-IM, vlan-id 120
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300224             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 mhevpn, route-type Egress-IM, vlan-id 10
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300240             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-IM, vlan-id 110
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300256             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-IM, vlan-id 130
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300272             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-IM, vlan-id 210
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300288             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-IM, vlan-id 220
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300304             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-IM, vlan-id 230
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300320             *[EVPN/7] 02:27:06, routing-instance VS-1, route-type
Egress-MAC, ESI 00:11:11:11:11:11:11:11:11:11
                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300336             *[EVPN/7] 02:27:06, routing-instance VS-1, route-type
Egress-MAC, ESI 00:33:33:33:33:33:33:33:33:33
                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
```

Copyright © 2017, Juniper Networks, Inc.

```
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300368             *[EVPN/7] 02:27:07, routing-instance VS-2, route-type
Egress-MAC, ESI 00:33:33:33:33:33:33:33:33:33
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300384             *[EVPN/7] 02:27:07, routing-instance VS-2, route-type
Egress-MAC, ESI 00:11:11:11:11:11:11:11:11:11
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300416             *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:33:33:33:33:33:33:33:33:33
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300432             *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:11:11:11:11:11:11:11:11:11
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300480             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-MAC
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300496             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-MAC
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300560             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-MAC
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300592             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-MAC
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300608             *[EVPN/7] 02:29:23
                    > via ge-0/0/1.1001, Pop
300624             *[EVPN/7] 02:29:23
                    > via ge-0/0/1.2001, Pop
301232             *[EVPN/7] 02:29:17
                    > via ge-0/0/1.1002, Pop
301296             *[EVPN/7] 02:29:10
                    > via ge-0/0/1.1003, Pop
301312             *[EVPN/7] 02:27:06
                    > via ae10.2003, Pop
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301360             *[EVPN/7] 02:29:01
                    > via ge-0/0/1.1004, Pop
301408             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 vpws1004, route-type Egress, vlan-id 2004
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301456             *[EVPN/7] 02:27:06
                    > via ae10.1010, Pop
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301552             *[EVPN/7] 02:27:07, routing-instance VS-1, route-type
Egress-MAC, ESI 00:22:22:22:22:22:22:22:22:22
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301568             *[EVPN/7] 02:27:07, routing-instance VS-2, route-type
```


Copyright © 2017, Juniper Networks, Inc.

```
Egress-MAC, ESI 00:22:22:22:22:22:22:22:22:22
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301648              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 vpws1010, route-type Egress, vlan-id 2010
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301664              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301680              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301696              *[EVPN/7] 02:27:07, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:22:22:22:22:22:22:22:22:22
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301712              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301728              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301744              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-IM, vlan-id 230
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301760              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 vpws1010, route-type Egress, vlan-id 2010
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301776              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301792              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-IM, vlan-id 130
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301808              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 vpws1004, route-type Egress, vlan-id 2004
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301824              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 mhevpn, route-type Egress-IM, vlan-id 10
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301840              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 vpws1002, route-type Egress, vlan-id 2002
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301856              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 vpws1003, route-type Egress, vlan-id 2003
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301872              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 vpws1003, route-type Egress Protection, vlan-id 2003
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301888              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 vpws1010, route-type Egress Protection, vlan-id 1010
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301904              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-IM, vlan-id 220
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301920              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-IM, vlan-id 210
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301936              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-IM, vlan-id 230
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301952              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-SH, vlan-id 230
```

Copyright © 2017, Juniper Networks, Inc.

```
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301968              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-IM, vlan-id 220
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301984              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-SH, vlan-id 220
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302000              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-IM, vlan-id 210
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302016              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-SH, vlan-id 210
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302032              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302048              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302064              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302080              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302096              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302112              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302128              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302144              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302160              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-IM, vlan-id 120
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302176              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-IM, vlan-id 110
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302192              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-IM, vlan-id 130
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302208              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-SH, vlan-id 130
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302224              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-IM, vlan-id 120
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302240              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-SH, vlan-id 120
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302256              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-IM, vlan-id 110
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302272              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-SH, vlan-id 110
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
```


Copyright © 2017, Juniper Networks, Inc.

```
302288             *[EVPN/7] 02:27:06, remote-pe 100.100.100.1, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302304             *[EVPN/7] 02:27:06, remote-pe 100.100.100.1, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302320             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302336             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 mhevpn, route-type Egress-MAC
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302352             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 vpws1004, route-type Egress, vlan-id 2004
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302368             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 mhevpn, route-type Egress-IM, vlan-id 10
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302384             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 mhevpn, route-type Egress-SH, vlan-id 10
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302400             *[EVPN/7] 02:26:21
                    > via ge-0/0/1.3001, Pop
302432             *[EVPN/7] 02:26:21, remote-pe 100.100.100.3, routing-instance
 vpws3001, route-type Egress, vlan-id 40000
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302448             *[EVPN/7] 02:26:21, remote-pe 100.100.100.1, routing-instance
 vpws3001, route-type Egress, vlan-id 40000
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302464             *[EVPN/7] 02:26:20, remote-pe 100.100.100.2, routing-instance
 vpws3001, route-type Egress, vlan-id 40000
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302480             *[EVPN/7] 02:26:14
                    > via ge-0/0/1.3016, Pop
302512             *[EVPN/7] 02:26:14, remote-pe 100.100.100.1, routing-instance
 vpws3016, route-type Egress, vlan-id 40016
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302528             *[EVPN/7] 02:26:14, remote-pe 100.100.100.2, routing-instance
 vpws3016, route-type Egress, vlan-id 40016
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302560             *[EVPN/7] 02:26:06
                    > via ae10.3011, Pop
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302592             *[EVPN/7] 02:26:07, remote-pe 100.100.100.1, routing-instance
 vpws3011, route-type Egress, vlan-id 401100
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302608             *[EVPN/7] 02:26:07, remote-pe 100.100.100.2, routing-instance
 vpws3011, route-type Egress, vlan-id 401100
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302624             *[EVPN/7] 02:26:07, remote-pe 100.100.100.3, routing-instance
 vpws3011, route-type Egress Protection, vlan-id 301100
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302656             *[EVPN/7] 02:25:59
                    > via ae10.3006, Pop
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302688             *[EVPN/7] 02:26:00, remote-pe 100.100.100.2, routing-instance
 vpws3006, route-type Egress, vlan-id 400600
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302704             *[EVPN/7] 02:26:00, remote-pe 100.100.100.1, routing-instance
 vpws3006, route-type Egress, vlan-id 400600
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
```

Copyright © 2017, Juniper Networks, Inc.

```
302720             *[EVPN/7] 02:25:59, remote-pe 100.100.100.3, routing-instance
 vpws3006, route-type Egress, vlan-id 400600
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302736             *[EVPN/7] 02:25:59, remote-pe 100.100.100.3, routing-instance
 vpws3006, route-type Egress Protection, vlan-id 300600
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ge-0/0/1.1001      *[EVPN/7] 02:29:23
                    > via ge-0/0/1.2001
ge-0/0/1.2001      *[EVPN/7] 02:29:23
                    > via ge-0/0/1.1001
ge-0/0/1.1002      *[EVPN/7] 02:27:06
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ae10.2003          *[EVPN/7] 02:29:10
                    > via ge-0/0/1.1003
ge-0/0/1.1003      *[EVPN/7] 02:27:06
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3

                     > via ae10.2003
ge-0/0/1.1004      *[EVPN/7] 02:27:06
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ae10.1010          *[EVPN/7] 02:27:06
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ge-0/0/1.3001      *[EVPN/7] 02:26:20
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ge-0/0/1.3016      *[EVPN/7] 02:26:13
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ae10.3011          *[EVPN/7] 02:26:06
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ae10.3006          *[EVPN/7] 02:25:59
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
```

show route table mpls.0 protocol ospf

```
user@host> show route table mpls.0 protocol ospf
mpls.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299952             *[L-OSPF/10] 23:59:42, metric 0
                    > to 10.0.10.70 via lt-1/2/0.14, Pop
                      to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299952(S=0)        *[L-OSPF/10] 23:59:42, metric 0
                    > to 10.0.10.70 via lt-1/2/0.14, Pop
                      to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299968             *[L-OSPF/10] 23:59:48, metric 0
                    > to 10.0.6.60 via lt-1/2/0.12, Pop
```


Copyright © 2017, Juniper Networks, Inc.

show route table mpls.0 extensive (PTX Series)

```
user@host> show route table mpls.0 extensive
ge-0/0/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel ge-0/0/2.600.0      /32 -> {composite(570)}
        *L2VPN  Preference: 7
                Next hop type: Indirect
                Address: 0x9438f34
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 567
                Next hop: 10.0.0.1 via ge-0/0/1.0, selected
                Label operation: Push 299808
                Label TTL action: prop-ttl
                Load balance label: Label 299808:None;
                Session Id: 0x1
                Protocol next hop: 10.255.255.1
                Label operation: Push 299872 Offset: 252
                Label TTL action: no-prop-ttl
                Load balance label: Label 299872:Flow label PUSH;
                Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                State: <Active Int>
                Age: 47         Metric2: 1
                Validation State: unverified
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 2-Common L2 VC
                AS path: I
                Composite next hops: 1
                        Protocol next hop: 10.255.255.1 Metric: 1
                        Label operation: Push 299872 Offset: 252
                        Label TTL action: no-prop-ttl
                        Load balance label: Label 299872:Flow label PUSH;
                        Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                        Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.0.0.1 via ge-0/0/1.0
                                Session Id: 0x1
                        10.255.255.1/32 Originating RIB: inet.3
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 10.0.0.1 via ge-0/0/1.0
```

show route table mpls.0 (RSVP Route—Transit LSP)

```
user@host> show route table mpls.0

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:37:31, metric 1
                      Receive
1                  *[MPLS/0] 00:37:31, metric 1
                      Receive
2                  *[MPLS/0] 00:37:31, metric 1
                      Receive
13                 *[MPLS/0] 00:37:31, metric 1
                      Receive
300352             *[RSVP/7/1] 00:08:00, metric 1
```

Copyright © 2017, Juniper Networks, Inc.

```
                    > to 10.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0)         *[RSVP/7/1] 00:08:00, metric 1
                    > to 10.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384              *[RSVP/7/2] 00:05:20, metric 1
                     > to 10.64.1.106 via ge-1/0/0.0, Pop
300384(S=0)         *[RSVP/7/2] 00:05:20, metric 1
                     > to 10.64.1.106 via ge-1/0/0.0, Pop
```

show route table vpls_1 detail

```
user@host> show route table vpls_1 detail
vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

172.16.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF
```

show route table vpn-a

```
user@host> show route table vpn-a
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
                    *[VPN/7] 05:48:27
                      Discard
192.168.24.1:1:2:1/96
                    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                       AS path: I
                     > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
                    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                       AS path: I
                     > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table vpn-a.mdt.0

```
user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
                    *[MVPN/70] 01:23:05, metric2 1
                       Indirect
1:1:1:10.255.14.218:232.1.1.1/144
                    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
                       AS path: I
                     > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
                    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
```


Copyright © 2017, Juniper Networks, Inc.

```
                        AS path: I
> via so-0/0/1.0, label-switched-path r0-to-r2
```

show route table VPN-A detail

```
user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.179.13:200
                Next hop type: Indirect
                Next-hop reference count: 5
                Source: 10.255.179.13
                Next hop type: Router, Next hop index: 732
                Next hop: 10.39.1.14 via fe-0/3/0.0, selected
                Label operation: Push 299824, Push 299824(top)
                Protocol next hop: 10.255.179.13
                Push 299824
                Indirect next hop: 8f275a0 1048574
                State: (Secondary Active Int Ext)
                Local AS: 1 Peer AS: 1
                Age: 3:41:06 Metric: 1 Metric2: 1
                Task: BGP_1.10.255.179.13+64309
                Announcement bits (2): 0-KRT 1-BGP RT Background
                AS path: I
                Communities: target:1:200 rte-type:0.0.0.0:1:0
                Import Accepted
                VPN Label: 299824 TTL Action: vrf-ttl-propagate
                Localpref: 100
                Router ID: 10.255.179.13
                Primary Routing Table bgp.l3vpn.0
```

show route table VPN-AB.inet.0

```
user@host> show route table VPN-AB.inet.0
      VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.39.1.0/30       *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
      10.39.1.4/30       *[Direct/0] 00:08:42
                          > via so-5/1/0.0
      10.39.1.6/32       *[Local/0] 00:08:46
                            Local
      10.255.71.16/32    *[Static/5] 00:07:24
                          > via so-2/0/0.0
      10.255.71.17/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                             AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.18/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                             AS path: I
                          > via so-2/1/0.0, Push 100021, Push 100011(top)
      10.255.245.245/32  *[BGP/170] 00:08:35, localpref 100
                             AS path: 2 I
                          > to 10.39.1.5 via so-5/1/0.0
      10.255.245.246/32  *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
```

Copyright © 2017, Juniper Networks, Inc.

show route table VPN_blue.mvpn-inet6.0

```
user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65536:10.255.2.202/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                    > via so-0/1/3.0
1:10.255.2.203:65536:10.255.2.203/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
                      AS path: I
                    > via so-0/1/0.0
1:10.255.2.204:65536:10.255.2.204/432
                   *[MVPN/70] 00:57:23, metric2 1
                      Indirect
5:10.255.2.202:65536:128:::192.168.90.2:128:ffff::1/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                    > via so-0/1/3.0
6:10.255.2.203:65536:64500:128:::10.12.53.12:128:ffff::1/432
                   *[PIM/105] 00:02:37
                      Multicast (IPv6)
7:10.255.2.202:65536:64500:128:::192.168.90.2:128:ffff::1/432
                   *[MVPN/70] 00:02:37, metric2 1
                      Indirect
```

show route table vrf1.mvpn.0 extensive

```
user@host> show route table vrf1.mvpn.0 extensive
1:10.255.50.77:1:10.255.50.77/240 (1 entry, 1 announced)
        *MVPN   Preference: 70
                PMSI: Flags 0x0: Label 0: RSVP-TE:
Session_13[10.255.50.77:0:25624:10.255.50.77]
                Next hop type: Indirect
                Address: 0xbb2c944
                Next-hop reference count: 360
                Protocol next hop: 10.255.50.77
                Indirect next hop: 0x0 - INH Session ID: 0x0
                State: <Active Int Ext>
                Age: 53:03      Metric2: 1
                Validation State: unverified
                Task: mvpn global task
                Announcement bits (3): 0-PIM.vrf1 1-mvpn global task 2-rt-export

                AS path: I
```

show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 2
                State: <Active Ext>
                Local AS: 64502 Peer AS: 64500
                Age: 4
                Task: BGP_64500.10.12.99.5+3792
                Announcement bits (1): 0-Flow
```


Copyright © 2017, Juniper Networks, Inc.

```
                AS path: 64500 I
                Communities: traffic-rate:0:0
                Validation state: Accept, Originator: 10.12.99.5
                Via: 10.12.44.0/24, Active
                Localpref: 100
                Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
        *Flow   Preference: 5
                Next-hop reference count: 2
                State: <Active>
                Local AS: 64502
                Age: 6:30
                Task: RT Flow
                Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
                AS path: I
                Communities: 1:1


user@host> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.2:100:10.1.1.2/96 AD
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
10.1.1.4:100:10.1.1.4/96 AD
                   *[BGP/170] 1d 03:11:02, localpref 100, from 10.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
10.1.1.2:100:1:0/96 MH
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
10.1.1.4:100:1:0/96 MH
                   *[BGP/170] 1d 03:11:02, localpref 100, from 10.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
10.1.1.4:NoCtrlWord:5:100:100:10.1.1.2:10.1.1.4/176
                   *[VPLS/7] 1d 03:11:02, metric2 1
                    >    via ge-1/2/1.5
10.1.1.4:NoCtrlWord:5:100:100:10.1.1.4:10.1.1.2/176
                   *[LDP/9] 1d 03:11:02
                         Discard


user@host> show route table red extensive
red.inet.0: 364481 destinations, 714087 routes (364480 active, 48448 holddown, 1
 hidden)
10.0.0.0/32 (3 entries, 1 announced)
        State: <OnList CalcForwarding>
TSI:
KRT in-kernel 10.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

    Nexthop: Self
    AS path: [2] I
    Communities: target:2:1
Path 10.0.0.0 from 10.3.0.0 Vector len 4.  Val: 1
        @BGP    Preference: 170/-1
                Route Distinguisher: 2:1
                Next hop type: Indirect
                Address: 0x258059e4
                Next-hop reference count: 2
```

```
                    Source: 2.2.0.0
                    Next hop type: Router
                    Next hop: 10.1.1.1 via ge-1/1/9.0, selected
                    Label operation: Push 707633
                    Label TTL action: prop-ttl
                    Session Id: 0x17d8
                    Protocol next hop: 10.2.0.0
                    Push 16
                    Composite next hop: 0x25805988 - INH Session ID: 0x193c
                    Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                    State: <Secondary Active Int Ext ProtectionPath ProtectionCand>
                    Local AS:     2 Peer AS:     2
                    Age: 23         Metric2: 35
                    Validation State: unverified
                    Task: BGP_172.16.2.0.0+34549
                    AS path: I
                    Communities: target:2:1
                    Import Accepted
                    VPN Label: 16
                    Localpref: 0
                    Router ID: 10.2.0.0
                    Primary Routing Table bgp.l3vpn.0
                    Composite next hops: 1
                            Protocol next hop: 10.2.0.0 Metric: 35
                            Push 16
                            Composite next hop: 0x25805988 - INH Session ID: 0x193c
                            Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                            Indirect path forwarding next hops: 1
                                    Next hop type: Router
                                    Next hop: 10.1.1.1 via ge-1/1/9.0
                                    Session Id: 0x17d8
                            2.2.0.0/32 Originating RIB: inet.3
                              Metric: 35                      Node path count: 1
                              Forwarding nexthops: 1
                                    Nexthop: 10.1.1.1 via ge-1/1/9.0
            BGP     Preference: 170/-1
                    Route Distinguisher: 2:1
                    Next hop type: Indirect
                    Address: 0x9347028
                    Next-hop reference count: 3
                    Source: 10.3.0.0
                    Next hop type: Router, Next hop index: 702
                    Next hop: 10.1.4.2 via ge-1/0/0.0, selected
                    Label operation: Push 634278
                    Label TTL action: prop-ttl
                    Session Id: 0x17d9
                    Protocol next hop: 10.3.0.0
                    Push 16
                    Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
                    Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da
                    State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

                    Inactive reason: Not Best in its group - IGP metric
                    Local AS:     2 Peer AS:     2
                    Age: 3:34       Metric2: 70
                    Validation State: unverified
                    Task: BGP_172.16.3.0.0+32805
                    Announcement bits (2): 0-KRT 1-BGP_RT_Background
                    AS path: I
                    Communities: target:2:1
                    Import Accepted
```


Copyright © 2017, Juniper Networks, Inc.

```
                    VPN Label: 16
                    Localpref: 0
                    Router ID: 10.3.0.0
                    Primary Routing Table bgp.l3vpn.0
                    Composite next hops: 1
                            Protocol next hop: 10.3.0.0 Metric: 70
                            Push 16
                            Composite next hop: 0x93463a0 1048575 INH Session ID:
0x17da
                            Indirect next hop: 0x91e8800 1048574 INH Session ID:
0x17da
                            Indirect path forwarding next hops: 1
                                    Next hop type: Router
                                    Next hop: 10.1.4.2 via ge-1/0/0.0
                                    Session Id: 0x17d9
                            10.3.0.0/32 Originating RIB: inet.3
                              Metric: 70                    Node path count: 1
                              Forwarding nexthops: 1
                                    Nexthop: 10.1.4.2 via ge-1/0/0.0
            #Multipath Preference: 255
                    Next hop type: Indirect
                    Address: 0x24afca30
                    Next-hop reference count: 1
                    Next hop type: Router
                    Next hop: 10.1.1.1 via ge-1/1/9.0, selected
                    Label operation: Push 707633
                    Label TTL action: prop-ttl
                    Session Id: 0x17d8
                    Next hop type: Router, Next hop index: 702
                    Next hop: 10.1.4.2 via ge-1/0/0.0
                    Label operation: Push 634278
                    Label TTL action: prop-ttl
                    Session Id: 0x17d9
                    Protocol next hop: 10.2.0.0
                    Push 16
                    Composite next hop: 0x25805988 - INH Session ID: 0x193c
                   Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

                    Protocol next hop: 10.3.0.0
                    Push 16
                    Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
                  Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da Weight
 0x4000
                    State: <ForwardingOnly Int Ext>
                    Inactive reason: Forwarding use only
                    Age: 23         Metric2: 35
                    Validation State: unverified
                    Task: RT
                    AS path: I
                    Communities: target:2:1
```

show route table bgp.evpn.0 extensive |no-more (EVPN)

```
show route table bgp.evpn.0 extensive | no-more
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
2:1000:10::100::00:aa:aa:aa:aa:aa/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
```

```
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
         State: Local AS: 17  Peer AS:17  Age:21:12  Metric2:1  Validation State:
 unverified
                Task: BGP_17.1.2.3.4+50756
                AS path: I
                Communities: target:1111:8388708 encapsulation0:0:0:0:3
                Import Accepted
                Route Label: 100
                ESI: 00:00:00:00:00:00:00:00:00:00
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        1.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

2:1000:10::200::00:bb:bb:bb:bb:bb/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS:17  Peer AS:17  Age:19:43  Metric2:1  Validation
State:unverified
                Task: BGP_17.1.2.3.4+50756
                AS path: I
                Communities: target:2222:22 encapsulation0:0:0:0:3
                Import Accepted
                Route Label: 200
                ESI: 00:00:00:00:00:00:00:00:00:00
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

2:1000:10::300::00:cc:cc:cc:cc:cc/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
```


Copyright © 2017, Juniper Networks, Inc.

```
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS:17  Peer AS:17  Age:17:21  Metric2:1  Validation State:
 unverified  Task: BGP 17,1,2,3,4+50756
           AS path: I
                Communities: target:3333:33 encapsulation0:0:0:0:3
                Import Accepted
                Route Label: 300
                ESI: 00:00:00:00:00:00:00:00:00:00
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::100::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                PMSI: Flags 0x0: Label 100: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS:17  Peer AS:17 Age:37:01  Metric2:1  Validation State:
 unverified  Task: BGP 17.1.2.3.4+50756
                AS path: I
                Communities: target:1111:8388708 encapsulation0:0:0:0:3
                Import Accepted
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::200::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
```

```
                PMSI: Flags 0x0: Label 200: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS: 17 Peer AS: 17   Age:35:22  Metric2:1  Validation
State:unverified  Task: BGP 17.1.2.3.4+50756
           AS path:I  Communities: target:2222:22 encapsulation):0:0:0:0:3

Import Accepted
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::300::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                PMSI: Flags 0x0: Label 300: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
         State: Local AS: 17 Peer AS:  17  Age 35:22 Metric2:1  Validation State:
 unverified   Task: BGP 17.1.2.3.4+5075
          6  AS path: I   Communities: target:3333:33  encapsulation0:0:0:0:3
Import Accepted   Localpref:100
          Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0
```


Copyright © 2017, Juniper Networks, Inc.

show route terse

List of Syntax
Syntax on page 447
Syntax (EX Series Switches) on page 447

Syntax
show route terse
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route terse

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display a high-level summary of the routes in the routing table.

NOTE: For BGP routes, the **show route terse** command displays the local preference attribute and MED instead of the metric1 and metric2 values. This is mostly due to historical reasons.

To display the metric1 and metric2 value of a BGP route, use the show route extensive command.

Options
none—Display a high-level summary of the routes in the routing table.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route terse on page 449

Output Fields
Table 30 on page 447 describes the output fields for the **show route terse** command. Output fields are listed in the approximate order in which they appear.

Table 30: show route terse Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number destinations	Number of destinations for which there are routes in the routing table.

Copyright © 2017, Juniper Networks, Inc.

Table 30: show route terse Output Fields *(continued)*

Field Name	Field Description
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)
route key	Key for the state of the route: • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route.
A	Active route. An asterisk (*) indicates this is the active route.
V	Validation status of the route: • **?**—Not evaluated. Indicates that the route was not learned through BGP. • **I**—Invalid. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **N**—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **V**—Valid. Indicates that the prefix and autonomous system pair are found in the database.
Destination	Destination of the route.
P	Protocol through which the route was learned: • **A**—Aggregate • **B**—BGP • **C**—CCC • **D**—Direct • **G**—GMPLS • **I**—IS-IS • **L**—L2CKT, L2VPN, LDP, Local • **K**—Kernel • **M**—MPLS, MSDP • **O**—OSPF • **P**—PIM • **R**—RIP, RIPng • **S**—Static • **T**—Tunnel
Prf	Preference value of the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.

Copyright © 2017, Juniper Networks, Inc.

Table 30: show route terse Output Fields *(continued)*

Field Name	Field Description
Metric 1	First metric value in the route. For routes learned from BGP, this is the MED metric.
Metric 2	Second metric value in the route. For routes learned from BGP, this is the IGP metric.
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated.

Sample Output

show route terse

```
user@host> show route terse
inet.0: 10 destinations, 12 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 172.16.1.1/32        O  10          1            >10.0.0.2
  ?                     B 170        100                          I
  unverified                                      >10.0.0.2
* ? 172.16.1.1/32        D   0                       >lo0.2
* V 2.2.0.2/32          B 170        110                          200 I
  valid                                           >10.0.0.2
* ? 10.0.0.0/30         D   0                       >lt-1/2/0.1
  ?                     B 170        100                          I
  unverified                                      >10.0.0.2
* ? 10.0.0.1/32         L   0                        Local
* ? 10.0.0.4/30         B 170        100                          I
  unverified                                      >10.0.0.2
* ? 10.0.0.8/30         B 170        100                          I
  unverified                                      >10.0.0.2
* I 172.16.1.1/32       B 170         90                          200 I
  invalid                                         >10.0.0.2
* N 192.168.2.3/32      B 170        100                          200 I
  unknown                                         >10.0.0.2
* ? 172.16.233.5/32       O  10          1             MultiRecv
```

Copyright © 2017, Juniper Networks, Inc.

test policy

Syntax test policy *policy-name prefix*

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Test a policy configuration to determine which prefixes match routes in the routing table.

NOTE: If you are using the **test policy** command on a logical system, you must first set the CLI to the logical system context. For example, if you want to test a routing policy that is configured on logical system R2, first run the **set cli logical-system R2** command.

Options ***policy-name***—Name of a policy.

prefix—Destination prefix to match.

Additional Information All prefixes in the default unicast routing table (inet.0) that match prefixes that are the same as or longer than the specific prefix are processed by the **from** clause in the specified policy. All prefixes accepted by the policy are displayed. The **test policy** command evaluates a policy differently from the BGP import process. When testing a policy that contains an **interface** match condition in the **from** clause, the **test policy** command uses the match condition. In contrast, BGP does not use the **interface** match condition when evaluating the policy against routes learned from internal BGP (IBGP) or external BGP (EGBP) multihop peers.

When testing a policy, you can see the length of time (in microseconds) required to evaluate the policy and the number of times it has been executed by running the show policy ***policy-name*** **statistics** command.

Required Privilege Level view

Related Documentation

- *Understanding Routing Policy Tests*
- Example: Testing a Routing Policy with Complex Regular Expressions on page 129
- show policy on page 248

List of Sample Output test policy on page 451

Copyright © 2017, Juniper Networks, Inc.

Output Fields For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

test policy

```
user@host> test policy test-statics 172.16.0.1/8
inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
Prefixes passing policy:

172.16.3.0/8          *[BGP/170] 16:22:46, localpref 100, from 10.255.255.41
                    AS Path: 50888 I
                     > to 10.11.4.32 via en0.2, label-switched-path l2
172.16.3.1/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
172.16.3.2/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
172.16.3.3/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
172.16.3.4/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
Policy test-statics: 5 prefixes accepted, 0 prefixes rejected
```

Copyright © 2017, Juniper Networks, Inc.

Copyright © 2017, Juniper Networks, Inc.

CHAPTER 23

RIPng Operational Commands

- clear ripng general-statistics
- clear ripng statistics
- restart
- show policy
- show policy conditions
- show ripng general-statistics
- show ripng neighbor
- show ripng statistics
- show route
- show route active-path
- show route advertising-protocol
- show route all
- show route best
- show route brief
- show route detail
- show route exact
- show route export
- show route extensive
- show route forwarding-table
- show route hidden
- show route inactive-path
- show route instance
- show route next-hop
- show route output
- show route protocol
- show route receive-protocol
- show route table

Copyright © 2017, Juniper Networks, Inc.

- show route terse
- test policy

Copyright © 2017, Juniper Networks, Inc.

clear ripng general-statistics

List of Syntax
Syntax on page 455
Syntax (EX Series Switches) on page 455

Syntax
clear ripng general-statistics
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
clear ripng general-statistics

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Clear RIP next generation (RIPng) general statistics.

Options
none—Clear RIPng general statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
clear

Related Documentation
- show ripng general-statistics on page 472

List of Sample Output
clear ripng general-statistics on page 455

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ripng general-statistics

```
user@host> clear ripng general-statistics
```

Copyright © 2017, Juniper Networks, Inc.

clear ripng statistics

List of Syntax
Syntax on page 456
Syntax (EX Series Switch) on page 456

Syntax
clear ripng statistics
<*instance* | *name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch)
clear ripng statistics
<*instance* | *name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Clear RIP next-generation (RIPng) statistics.

Options
none—Reset RIPng counters for all neighbors for all routing instances.

instance—(Optional) Reset RIPng counters for the specified instance.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Reset RIPng counters for the specified neighbor.

Required Privilege Level
clear

Related Documentation
- show ripng statistics on page 476

List of Sample Output
clear ripng statistics on page 456

Output Fields
When you enter this command, you are provided feedback on the status of your request.

Sample Output

clear ripng statistics

```
user@host> clear ripng statistics
```


Copyright © 2017, Juniper Networks, Inc.

restart

List of Syntax

Syntax on page 457
Syntax (ACX Series Routers) on page 457
Syntax (EX Series Switches) on page 457
Syntax (MX Series Routers) on page 458
Syntax (QFX Series) on page 458
Syntax (Routing Matrix) on page 458
Syntax (TX Matrix Routers) on page 458
Syntax (TX Matrix Plus Routers) on page 459
Syntax (QFX Series) on page 459

Syntax

```
restart
<adaptive-services |ancpd-service | application-identification |audit-process |
  auto-configuration |captive-portal-content-delivery |ce-l2tp-service |chassis-control |
  class-of-service |clksyncd-service |database-replication|datapath-trace-service
  |dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture |
  ecc-error-logging | ethernet-connectivity-fault-management
  |ethernet-link-fault-management |event-processing | firewall
  |general-authentication-service | gracefully | iccp-service |idp-policy | immediately
  |interface-control | ipsec-key-management | kernel-replication | l2-learning | l2cpd-service
  | l2tp-service | l2tp-universal-edge | lacp | license-service |link-management
  |local-policy-decision-function |mac-validation |mib-process | mountd-service
  |mpls-traceroute |mspd | multicast-snooping |named-service | nfsd-service |
  packet-triggered-subscribers |peer-selection-service |pgm | pic-services-logging | pki-service
  |ppp | ppp-service |pppoe | protected-system-domain-service |
  redundancy-interface-process | remote-operations | root-system-domain-service | routing
  <logical-system logical-system-name> | sampling | sbc-configuration-process | sdk-service
  |service-deployment | services | snmp |soft |static-subscribers |statistics-service|
  subscriber-management | subscriber-management-helper | tunnel-oamd |usb-control|
  vrrp |web-management>
<gracefully | immediately | soft>
```

Syntax (ACX Series Routers)

```
restart
<adaptive-services |audit-process | auto-configuration | autoinstallation |chassis-control |
  class-of-service |clksyncd-service |database-replication| dhcp-service | diameter-service
  | disk-monitoring | dynamic-flow-capture | ethernet-connectivity-fault-management
  |ethernet-link-fault-management |event-processing | firewall
  |general-authentication-service | gracefully | immediately |interface-control |
  ipsec-key-management | l2-learning | lacp |link-management |mib-process | mountd-service
  |mpls-traceroute |mspd | named-service | nfsd-service | pgm | pki-service |ppp | pppoe |
  redundancy-interface-process | remote-operations | routing | sampling | sdk-service
  |secure-neighbor-discovery | service-deployment | services | snmp |soft |statistics-service|
  subscriber-management | subscriber-management-helper | tunnel-oamd | vrrp>
```

Syntax (EX Series Switches)

```
restart
<autoinstallation | chassis-control | class-of-service | database-replication | dhcp |
  dhcp-service | diameter-service | dot1x-protocol | ethernet-link-fault-management |
  ethernet-switching | event-processing | firewall | general-authentication-service |
  interface-control | kernel-replication | l2-learning | lacp | license-service | link-management
  | lldpd-service | mib-process | mountd-service | multicast-snooping | pgm |
```

Copyright © 2017, Juniper Networks, Inc.

redundancy-interface-process | remote-operations | routing | secure-neighbor-discovery | service-deployment | sflow-service | snmp | vrrp | web-management>

Syntax (MX Series Routers)

```
restart
<adaptive-services | ancpd-service |application-identification|audit-process |
  auto-configuration |captive-portal-content-delivery |ce-l2tp-service |chassis-control |
  class-of-service |clksyncd-service |database-replication| datapath-trace-service
  |dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture |
  ecc-error-logging | ethernet-connectivity-fault-management
  |ethernet-link-fault-management |event-processing | firewall |
  general-authentication-service | gracefully |iccp-service |idp-policy |immediately
  |interface-control | ipsec-key-management | kernel-replication | l2-learning |l2cpd-service
  |l2tp-service | l2tp-universal-edge | lacp | license-service |link-management
  |local-policy-decision-function |mac-validation |mib-process | mountd-service
  |mpls-traceroute |mspd | multicast-snooping |named-service | nfsd-service |
  packet-triggered-subscribers |peer-selection-service | pgm | pic-services-logging |
  pki-service |ppp | ppp-service |pppoe | protected-system-domain-service |
  redundancy-interface-process | remote-operations |root-system-domain-service | routing
  |routing <logical-system logical-system-name>  | sampling | sbc-configuration-process |
  sdk-service |service-deployment |services | snmp |soft |static-subscribers |statistics-service|
  subscriber-management | subscriber-management-helper | tunnel-oamd | usb-control|
  vrrp |web-management>
<all-members>
<gracefully | immediately | soft>
<local>
<member member-id>
```

Syntax (QFX Series)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services |
  diameter-service | dlsw | ethernet-connectivity | event-processing | fibre-channel | firewall
  | general-authentication-service | igmp-host-services | interface-control |
  ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process |
  named-service | network-access-service | nstrace-process | pgm | ppp | pppoe |
  redundancy-interface-process | remote-operations |logical-system-name>  | routing |
  sampling |secure-neighbor-discovery | service-deployment | snmp | usb-control |
  web-management>
<gracefully | immediately | soft>
```

Syntax (Routing Matrix)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | disk-monitoring |
  dynamic-flow-capture | ecc-error-logging | event-processing | firewall | interface-control
  | ipsec-key-management | kernel-replication | l2-learning | l2tp-service | lacp |
  link-management | mib-process | pgm | pic-services-logging | ppp | pppoe |
  redundancy-interface-process | remote-operations | routing <logical-system
  logical-system-name>  | sampling | service-deployment | snmp>
<all | all-lcc | lcc number>
<gracefully | immediately | soft>
```

Syntax (TX Matrix Routers)

```
restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp-service |
  diameter-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging |
  event-processing | firewall | interface-control | ipsec-key-management | kernel-replication
  | l2-learning | l2tp-service | lacp | link-management | mib-process |pgm | pic-services-logging
```


Copyright © 2017, Juniper Networks, Inc.

| ppp | pppoe | redundancy-interface-process | remote-operations | routing <logical-system *logical-system-name*> | sampling | service-deployment | snmp| statistics-service>
<all-chassis | all-lcc | lcc *number* | scc>
<gracefully | immediately | soft>

Syntax (TX Matrix Plus Routers)

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dhcp-service | diameter-service | disk-monitoring | dynamic-flow-capture | ecc-error-logging | event-processing | firewall | interface-control | ipsec-key-management | kernel-replication | l2-learning | l2tp-service | lacp | link-management | mib-process | pgm | pic-services-logging | ppp | pppoe | redundancy-interface-process | remote-operations | routing <logical-system *logical-system-name*> | sampling | service-deployment | snmp| statistics-service>
<all-chassis | all-lcc | all-sfc | lcc *number* | sfc *number*>
<gracefully | immediately | soft>

Syntax (QFX Series)

restart
<adaptive-services | audit-process | chassis-control | class-of-service | dialer-services | diameter-service | dlsw | ethernet-connectivity | event-processing | fibre-channel | firewall | general-authentication-service | igmp-host-services | interface-control | ipsec-key-management | isdn-signaling | l2ald | l2-learning | l2tp-service | mib-process | named-service | network-access-service | nstrace-process | pgm | ppp | pppoe | redundancy-interface-process | remote-operations |*logical-system-name*> | routing | sampling |secure-neighbor-discovery | service-deployment | snmp | usb-control | web-management>
<gracefully | immediately | soft>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.1 for the QFX Series.
Command introduced in Junos OS Release 12.2 for ACX Series routers.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.
Options added:

- **dynamic-flow-capture** in Junos OS Release 7.4.
- **dlsw** in Junos OS Release 7.5.
- **event-processing** in Junos OS Release 7.5.
- **ppp** in Junos OS Release 7.5.
- **l2ald** in Junos OS Release 8.0.
- **link-management** in Release 8.0.
- **pgcp-service** in Junos OS Release 8.4.
- **sbc-configuration-process** in Junos OS Release 9.5.
- **services pgcp gateway** in Junos OS Release 9.6.
- **sfc** and **all-sfc** for the TX Matrix Router in Junos OS Release 9.6.

Description

Restart a Junos OS process.

Copyright © 2017, Juniper Networks, Inc.

CAUTION: Never restart a software process unless instructed to do so by a customer support engineer. A restart might cause the router or switch to drop calls and interrupt transmission, resulting in possible loss of data.

Options **none**—Same as **gracefully**.

adaptive-services—(Optional) Restart the configuration management process that manages the configuration for stateful firewall, Network Address Translation (NAT), intrusion detection services (IDS), and IP Security (IPsec) services on the Adaptive Services PIC.

all-chassis—(TX Matrix and TX Matrix Plus routers only) (Optional) Restart the software process on all chassis.

all-lcc—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process on all T640 routers connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process on all T1600 routers connected to the TX Matrix Plus router.

all-members—(MX Series routers only) (Optional) Restart the software process for all members of the Virtual Chassis configuration.

all-sfc—(TX Matrix Plus routers only) (Optional) For a TX Matrix Plus router, restart the software processes for the TX Matrix Plus router (or switch-fabric chassis).

ancpd-service—(Optional) Restart the Access Node Control Protocol (ANCP) process, which works with a special Internet Group Management Protocol (IGMP) session to collect outgoing interface mapping events in a scalable manner.

application-identification—(Optional) Restart the process that identifies an application using intrusion detection and prevention (IDP) to allow or deny traffic based on applications running on standard or nonstandard ports.

audit-process—(Optional) Restart the RADIUS accounting process that gathers statistical data that can be used for general network monitoring, analyzing, and tracking usage patterns, for billing a user based on the amount of time or type of services accessed.

auto-configuration—(Optional) Restart the Interface Auto-Configuration process.

autoinstallation—(EX Series switches only) (Optional) Restart the autoinstallation process.

captive-portal-content-delivery—(Optional) Restart the HTTP redirect service by specifying the location to which a subscriber's initial Web browser session is redirected, enabling initial provisioning and service selection for the subscriber.

ce-l2tp-service—(M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Universal Edge Layer 2 Tunneling Protocol (L2TP) process, which establishes L2TP tunnels and Point-to-Point Protocol (PPP) sessions through L2TP tunnels.

Copyright © 2017, Juniper Networks, Inc.

chassis-control—(Optional) Restart the chassis management process.

class-of-service—(Optional) Restart the class-of-service (CoS) process, which controls the router's or switch's CoS configuration.

clksyncd-service—(Optional) Restart the external clock synchronization process, which uses synchronous Ethernet (SyncE).

database-replication—(EX Series switches and MX Series routers only) (Optional) Restart the database replication process.

datapath-trace-service—(Optional) Restart the packet path tracing process.

dhcp—(EX Series switches only) (Optional) Restart the software process for a Dynamic Host Configuration Protocol (DHCP) server. A DHCP server allocates network IP addresses and delivers configuration settings to client hosts without user intervention.

dhcp-service—(Optional) Restart the Dynamic Host Configuration Protocol process.

dialer-services—(EX Series switches only) (Optional) Restart the ISDN dial-out process.

diameter-service—(Optional) Restart the diameter process.

disk-monitoring—(Optional) Restart disk monitoring, which checks the health of the hard disk drive on the Routing Engine.

dlsw—(QFX Series only) (Optional) Restart the data link switching (DLSw) service.

dot1x-protocol—(EX Series switches only) (Optional) Restart the port-based network access control process.

dynamic-flow-capture—(Optional) Restart the dynamic flow capture (DFC) process, which controls DFC configurations on Monitoring Services III PICs.

ecc-error-logging—(Optional) Restart the error checking and correction (ECC) process, which logs ECC parity errors in memory on the Routing Engine.

ethernet-connectivity-fault-management—(Optional) Restart the process that provides IEEE 802.1ag Operation, Administration, and Management (OAM) connectivity fault management (CFM) database information for CFM maintenance association end points (MEPs) in a CFM session.

ethernet-link-fault-management—(EX Series switches and MX Series routers only) (Optional) Restart the process that provides the OAM link fault management (LFM) information for Ethernet interfaces.

ethernet-switching—(EX Series switches only) (Optional) Restart the Ethernet switching process.

event-processing—(Optional) Restart the event process (eventd).

fibre-channel—(QFX Series only) (Optional) Restart the Fibre Channel process.

Copyright © 2017, Juniper Networks, Inc.

firewall—(Optional) Restart the firewall management process, which manages the firewall configuration and enables accepting or rejecting packets that are transiting an interface on a router or switch.

general-authentication-service—(EX Series switches and MX Series routers only) (Optional) Restart the general authentication process.

gracefully—(Optional) Restart the software process.

iccp-service—(Optional) Restart the Inter-Chassis Communication Protocol (ICCP) process.

idp-policy—(Optional) Restart the intrusion detection and prevention (IDP) protocol process.

immediately—(Optional) Immediately restart the software process.

interface-control—(Optional) Restart the interface process, which controls the router's or switch's physical interface devices and logical interfaces.

ipsec-key-management—(Optional) Restart the IPsec key management process.

isdn-signaling—(QFX Series only) (Optional) Restart the ISDN signaling process, which initiates ISDN connections.

kernel-replication—(Optional) Restart the kernel replication process, which replicates the state of the backup Routing Engine when graceful Routing Engine switchover (GRES) is configured.

l2-learning—(Optional) Restart the Layer 2 address flooding and learning process.

l2cpd-service—(Optional) Restart the Layer 2 Control Protocol process, which enables features such as Layer 2 protocol tunneling and nonstop bridging.

l2tp-service— (M10, M10i, M7i, and MX Series routers only) (Optional) Restart the Layer 2 Tunneling Protocol (L2TP) process, which sets up client services for establishing Point-to-Point Protocol (PPP) tunnels across a network and negotiating Multilink PPP if it is implemented.

l2tp-universal-edge—(MX Series routers only) (Optional) Restart the L2TP process, which establishes L2TP tunnels and PPP sessions through L2TP tunnels.

lacp—(Optional) Restart the Link Aggregation Control Protocol (LACP) process. LACP provides a standardized means for exchanging information between partner systems on a link to allow their link aggregation control instances to reach agreement on the identity of the LAG to which the link belongs, and then to move the link to that LAG, and to enable the transmission and reception processes for the link to function in an orderly manner.

lcc *number*—(TX Matrix and TX Matrix Plus routers only) (Optional) For a TX Matrix router, restart the software process for a specific T640 router that is connected to the TX Matrix router. For a TX Matrix Plus router, restart the software process for a specific router that is connected to the TX Matrix Plus router.

Copyright © 2017, Juniper Networks, Inc.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

license-service—(EX Series switches only) (Optional) Restart the feature license management process.

link-management— (TX Matrix and TX Matrix Plus routers and EX Series switches only) (Optional) Restart the Link Management Protocol (LMP) process, which establishes and maintains LMP control channels.

lldpd-service—(EX Series switches only) (Optional) Restart the Link Layer Discovery Protocol (LLDP) process.

local—(MX Series routers only) (Optional) Restart the software process for the local Virtual Chassis member.

local-policy-decision-function— (Optional) Restart the process for the Local Policy Decision Function, which regulates collection of statistics related to applications and application groups and tracking of information about dynamic subscribers and static interfaces.

mac-validation— (Optional) Restart the Media Access Control (MAC) validation process, which configures MAC address validation for subscriber interfaces created on demux interfaces in dynamic profiles on MX Series routers.

member *member-id*—(MX Series routers only) (Optional) Restart the software process for a specific member of the Virtual Chassis configuration. Replace ***member-id*** with a value of **0** or **1**.

mib-process—(Optional) Restart the Management Information Base (MIB) version II process, which provides the router's MIB II agent.

mobile-ip—(Optional) Restart the Mobile IP process, which configures Junos OS Mobile IP features.

mountd-service—(EX Series switches and MX Series routers only) (Optional) Restart the service for NFS mount requests.

mpls-traceroute—(Optional) Restart the MPLS Periodic Traceroute process.

mspd—(Optional) Restart the Multiservice process.

Copyright © 2017, Juniper Networks, Inc.

multicast-snooping—(EX Series switches and MX Series routers only) (Optional) Restart the multicast snooping process, which makes Layer 2 devices, such as VLAN switches, aware of Layer 3 information, such as the media access control (MAC) addresses of members of a multicast group.

named-service—(Optional) Restart the DNS Server process, which is used by a router or a switch to resolve hostnames into addresses.

network-access-service—(QFX Series only) (Optional) Restart the network access process, which provides the router's Challenge Handshake Authentication Protocol (CHAP) authentication service.

nfsd-service—(Optional) Restart the Remote NFS Server process, which provides remote file access for applications that need NFS-based transport.

packet-triggered-subscribers—(Optional) Restart the packet-triggered subscribers and policy control (PTSP) process, which allows the application of policies to dynamic subscribers that are controlled by a subscriber termination device.

peer-selection-service—(Optional) Restart the Peer Selection Service process.

pgcp-service—(Optional) Restart the pgcpd service process running on the Routing Engine. This option does not restart pgcpd processes running on mobile station PICs. To restart pgcpd processes running on mobile station PICs, use the **services pgcp gateway** option.

pgm—(Optional) Restart the process that implements the Pragmatic General Multicast (PGM) protocol for assisting in the reliable delivery of multicast packets.

pic-services-logging—(Optional) Restart the logging process for some PICs. With this process, also known as fsad (the file system access daemon), PICs send special logging information to the Routing Engine for archiving on the hard disk.

pki-service—(Optional) Restart the PKI Service process.

ppp—(Optional) Restart the Point-to-Point Protocol (PPP) process, which is the encapsulation protocol process for transporting IP traffic across point-to-point links.

ppp-service—(Optional) Restart the Universal edge PPP process, which is the encapsulation protocol process for transporting IP traffic across universal edge routers.

pppoe—(Optional) Restart the Point-to-Point Protocol over Ethernet (PPPoE) process, which combines PPP that typically runs over broadband connections with the Ethernet link-layer protocol that allows users to connect to a network of hosts over a bridge or access concentrator.

protected-system-domain-service—(Optional) Restart the Protected System Domain (PSD) process.

redundancy-interface-process—(Optional) Restart the ASP redundancy process.

Copyright © 2017, Juniper Networks, Inc.

remote-operations—(Optional) Restart the remote operations process, which provides the ping and traceroute MIBs.

root-system-domain-service—(Optional) Restart the Root System Domain (RSD) service.

routing—(ACX Series routers, QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the routing protocol process.

routing <logical-system *logical-system-name*>—(Optional) Restart the routing protocol process, which controls the routing protocols that run on the router or switch and maintains the routing tables. Optionally, restart the routing protocol process for the specified logical system only.

sampling—(Optional) Restart the sampling process, which performs packet sampling based on particular input interfaces and various fields in the packet header.

sbc-configuration-process—(Optional) Restart the session border controller (SBC) process of the border signaling gateway (BSG).

scc—(TX Matrix routers only) (Optional) Restart the software process on the TX Matrix router (or switch-card chassis).

sdk-service—(Optional) Restart the SDK Service process, which runs on the Routing Engine and is responsible for communications between the SDK application and Junos OS. Although the SDK Service process is present on the router, it is turned off by default.

secure-neighbor-discovery—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the secure Neighbor Discovery Protocol (NDP) process, which provides support for protecting NDP messages.

sfc *number*—(TX Matrix Plus routers only) (Optional) Restart the software process on the TX Matrix Plus router (or switch-fabric chassis). Replace ***number*** with **0**.

service-deployment—(Optional) Restart the service deployment process, which enables Junos OS to work with the Session and Resource Control (SRC) software.

services—(Optional) Restart a service.

services pgcp gateway *gateway-name*—(Optional) Restart the pgcpd process for a specific border gateway function (BGF) running on an MS-PIC. This option does not restart the pgcpd process running on the Routing Engine. To restart the pgcpd process on the Routing Engine, use the **pgcp-service** option.

sflow-service—(EX Series switches only) (Optional) Restart the flow sampling (sFlow technology) process.

snmp—(Optional) Restart the SNMP process, which enables the monitoring of network devices from a central location and provides the router's or switch's SNMP master agent.

Copyright © 2017, Juniper Networks, Inc.

soft—(Optional) Reread and reactivate the configuration without completely restarting the software processes. For example, BGP peers stay up and the routing table stays constant. Omitting this option results in a graceful restart of the software process.

static-subscribers—(Optional) Restart the static subscribers process, which associates subscribers with statically configured interfaces and provides dynamic service activation and activation for these subscribers.

statistics-service—(Optional) Restart the process that manages the Packet Forwarding Engine statistics.

subscriber-management—(Optional) Restart the Subscriber Management process.

subscriber-management-helper—(Optional) Restart the Subscriber Management Helper process.

tunnel-oamd—(Optional) Restart the Tunnel OAM process, which enables the Operations, Administration, and Maintenance of Layer 2 tunneled networks. Layer 2 protocol tunneling (L2PT) allows service providers to send Layer 2 protocol data units (PDUs) across the provider's cloud and deliver them to Juniper Networks EX Series Ethernet Switches that are not part of the local broadcast domain.

usb-control—(MX Series routers) (Optional) Restart the USB control process.

vrrp—(ACX Series routers, EX Series switches, and MX Series routers only) (Optional) Restart the Virtual Router Redundancy Protocol (VRRP) process, which enables hosts on a LAN to make use of redundant routing platforms on that LAN without requiring more than the static configuration of a single default route on the hosts.

web-management—(QFX Series, EX Series switches, and MX Series routers only) (Optional) Restart the Web management process.

Required Privilege Level reset

Related Documentation
- *Overview of Junos OS CLI Operational Mode Commands*

List of Sample Output restart interface-control gracefully on page 466

Output Fields When you enter this command, you are provided feedback on the status of your request.

Sample Output

restart interface-control gracefully

```
user@host> restart interface-control gracefully
Interface control process started, pid 41129
```


Copyright © 2017, Juniper Networks, Inc.

show policy

List of Syntax
Syntax on page 467
Syntax (EX Series Switches) on page 467

Syntax

show policy
<logical-system (all | *logical-system-name*)>
<*policy-name*>
<*statistics* >

Syntax (EX Series Switches)

show policy
<*policy-name*>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
statistics option introduced in Junos OS Release 16.1 for MX Series routers.

Description

Display information about configured routing policies.

Options

none—List the names of all configured routing policies.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

policy-name—(Optional) Show the contents of the specified policy.

statistics—(Optional) Use in conjunction with the **test policy** command to show the length of time (in microseconds) required to evaluate a given policy and the number of times it has been executed. This information can be used, for example, to help structure a policy so it is evaluated efficiently. Timers shown are per route; times are not cumulative. Statistics are incremented even when the router is learning (and thus evaluating) routes from peering routers.

Required Privilege Level

view

Related Documentation

- *show policy damping*
- test policy on page 450

List of Sample Output

show policy on page 468
show policy policy-name on page 468
show policy statistics policy-name on page 468
show policy (Multicast Scoping) on page 469
show policy (Route Filter and source Address Filter Lists) on page 469

Copyright © 2017, Juniper Networks, Inc.

Output Fields Table 10 on page 249 lists the output fields for the **show policy** command. Output fields are listed in the approximate order in which they appear.

Table 31: show policy Output Fields

Field Name	Field Description
policy-name	Name of the policy listed.
term	Name of the user-defined policy term. The term name **unnamed** is used for policy elements that occur outside of user defined terms
from	Match condition for the policy.
then	Action for the policy.

Sample Output

show policy

```
user@host> show policy
Configured policies:
__vrf-export-red-internal__
__vrf-import-red-internal__
red-export
rf-test-policy
multicast-scoping
```

show policy policy-name

```
user@host> show policy vrf-import-red-internal
Policy vrf-import-red-internal:
    from
         203.0.113.0/28  accept
         203.0.113.32/28  accept
    then reject
```

show policy statistics policy-name

```
user@host> show policy statistics iBGP-v4-RR-Import
Policy iBGP-v4-RR-Import:
    [1243328] Term Lab-Infra:
        from [1243328 0]  proto BGP
             [28 0] route filter:
                 10.11.0.0/8 orlonger
                 10.13.0.0/8 orlonger
        then [28 0] accept
    [1243300] Term External:
        from [1243300 1]  proto BGP
         [1243296 0]  community Ext-Com1 [64496:1515 ]
         [1243296 0]  prefix-list-filter Customer-Routes
         [1243296 0]  aspath AS6221
             [1243296 1] route filter:
                 172.16.49.0/12 orlonger
                 172.16.50.0/12 orlonger
                 172.16.51.0/12 orlonger
```


Copyright © 2017, Juniper Networks, Inc.

```
                    172.16.52.0/12 orlonger
                    172.16.56.0/12 orlonger
                    172.16.60.0/12 orlonger
                then [1243296 2] community + Ext-Com2 [64496:2000 ] [1243296 0] accept
            [4] Term Final:
                then [4 0] reject
```

show policy (Multicast Scoping)

```
user@host> show policy multicast-scoping
Policy multicast-scoping:
    from
         multicast-scope == 8
  then
    accept
```

show policy (Route Filter and source Address Filter Lists)

```
user@host> show policy rf-test-policy
Policy rf-test-policy:
    Term term1:
        from source-address-filter-list saf-list-1
             source-address filter:
                 192.0.2.0/29 longer
                 192.0.2.64/28 exact
                 192.0.2.128/28 exact
                 192.0.2.160/28 orlonger
    Term term2:
        from route-filter-list rf-list-1
             route filter:
                 198.51.100.0/29 upto 198.51.100.0/30
                 198.51.100.8/29 upto 198.51.100.8/30 accept
    Term unnamed:
        then reject
```

Copyright © 2017, Juniper Networks, Inc.

show policy conditions

Syntax

show policy conditions
<*condition-name*>
<detail>
<dynamic>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)

show policy conditions
<*condition-name*>
<detail>
<dynamic>

Release Information

Command introduced in Junos OS Release 9.0.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description

Display all the configured conditions as well as the routing tables with which the configuration manager is interacting. If the **detail** keyword is included, the output also displays dependent routes for each condition.

Options

none—Display all configured conditions and associated routing tables.

condition-name—(Optional) Display information about the specified condition only.

detail—(Optional) Display the specified level of output.

dynamic—(Optional) Display information about the conditions in the dynamic database.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level

view

List of Sample Output

show policy conditions detail on page 471

Output Fields

Table 11 on page 251 lists the output fields for the **show policy conditions** command. Output fields are listed in the approximate order in which they appear.

Table 32: show policy conditions Output Fields

Field Name	Field Description	Level of Output
Condition	Name of configured condition.	All levels
event	Condition type. If the **if-route-exists** option is configured, the event type is: **Existence of a route in a specific routing table**.	All levels
Dependent routes	List of routes dependent on the condition, along with the latest generation number.	**detail**

Copyright © 2017, Juniper Networks, Inc.

Table 32: show policy conditions Output Fields *(continued)*

Field Name	Field Description	Level of Output
Condition tables	List of routing tables associated with the condition, along with the latest generation number and number of dependencies.	All levels
If-route-exists conditions	List of conditions configured to look for a route in the specified table.	All levels

Sample Output

show policy conditions detail

```
user@host> show policy conditions detail
Configured conditions:
Condition cond1, event: Existence of a route in a specific routing table
Dependent routes:
 172.16.4.4/32, generation 3
 6.6.6.6/32, generation 3
 10.10.10.10/32, generation 3

Condition cond2, event: Existence of a route in a specific routing table
Dependent routes:
None

Condition tables:
Table inet.0, generation 4, dependencies 3, If-route-exists conditions: cond1
(static) cond2 (static)
```

Copyright © 2017, Juniper Networks, Inc.

show ripng general-statistics

List of Syntax
Syntax on page 472
Syntax (EX Series Switch) on page 472

Syntax
show ripng general-statistics
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switch)
show ripng general-statistics

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display general RIP next-generation (RIPng) statistics.

Options
none—Display general RIPng statistics.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- clear ripng general-statistics on page 455

List of Sample Output
show ripng general-statistics on page 473

Output Fields
Table 33 on page 472 lists the output fields for the **show ripng general-statistics** command. Output fields are listed in the approximate order in which they appear.

Table 33: show ripng general-statistics Output Fields

Field Name	Field Description
bad msgs	Number of invalid messages received.
no recv intf	Number of packets received with no matching interface.
curr memory	Amount of memory currently used by RIPng.
max memory	Most memory used by RIPng.

Copyright © 2017, Juniper Networks, Inc.

Sample Output

show ripng general-statistics

```
user@host> show ripng general-statistics
RIPng I/O info:
     bad msgs       :         0
     no recv intf   :         0
     curr memory    :         0
     max memory     :         0
```

Copyright © 2017, Juniper Networks, Inc.

show ripng neighbor

List of Syntax
Syntax on page 474
Syntax (EX Series Switch) on page 474

Syntax
show ripng neighbor
<logical-system (all | *logical-system-name*)>
<*name*>

Syntax (EX Series Switch)
show ripng neighbor
<*name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display information about RIP next-generation (RIPng) neighbors.

Options
none—Display information about all RIPng neighbors.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display detailed information about a specific RIPng neighbor.

Required Privilege Level
view

List of Sample Output
show ripng neighbor on page 475

Output Fields
Table 34 on page 474 lists the output fields for the **show ripng neighbor** command. Output fields are listed in the approximate order in which they appear.

Table 34: show ripng neighbor Output Fields

Field Name	Field Description
Neighbor	Name of RIPng neighbor.
State	State of the connection: **Up** or **Dn** (Down).
Source Address	Source address.
Destination Address	Destination address.
Send	Send options: **broadcast**, **multicast**, **none**, **version 1**, or **yes.**
Recv	Type of packets to accept: **both**, **none**, **version 1**, or **yes.**

Copyright © 2017, Juniper Networks, Inc.

Table 34: show ripng neighbor Output Fields *(continued)*

Field Name	Field Description
In Met	Metric added to incoming routes when advertising into RIPng routes that were learned from other protocols.

Sample Output

show ripng neighbor

```
user@host> show ripng neighbor
                    Source                          Dest                In
Neighbor     State  Address                         Address  Send Recv Met
--------     -----  -------                         -------  ---- ---- ---
fe-0/0/2.0      Up  fe80::290:69ff:fe68:b002        ff02::9  yes  yes    1
```

show ripng statistics

List of Syntax
Syntax on page 476
Syntax (EX Series Switch) on page 476

Syntax
show ripng statistics
<logical-system (all | *logical-system-name*)>
<name>

Syntax (EX Series Switch)
show ripng statistics
<name>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display RIP next generation (RIPng) statistics about messages sent and received on an interface, as well as information received from advertisements from other routing devices.

Options
none—Display RIPng statistics for all neighbors.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

name—(Optional) Display detailed information about a specific RIPng neighbor.

Required Privilege Level
view

Related Documentation
- clear ripng statistics on page 456

List of Sample Output
show ripng statistics on page 477

Output Fields
Table 35 on page 477 lists the output fields for the **show ripng statistics** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 35: show ripng statistics Output Fields

Field Name	Field Description
RIPng info	Information about RIPng on the specified interface: • **port**—UDP port number used for RIPng. • **holddown**—Hold-down interval, in seconds. • **rts learned**—Number of routes learned through RIPng. • **rts held down**—Number of routes held down by RIPng. • **rqsts dropped**—Number of received request packets that were dropped. • **resps dropped**—Number of received response packets that were dropped. • **restart**—Graceful restart status. Displayed when RIPng is or has been in the process of graceful restart.
logical-interface	Name of the logical interface and its statistics: • **routes learned**—Number of routes learned on the logical interface. • **routes advertised**—Number of routes advertised by the logical interface. • **timeout**—Timeout interval, in seconds. • **update interval**—Interval between routing table updates, in seconds.
Counter	List of counter types: • **Updates Sent**—Number of update messages sent. • **Triggered Updates Sent**—Number of triggered update messages sent. • **Responses Sent**—Number of response messages sent. • **Bad Messages**—Number of invalid messages received. • **Updates Received**—Number of RIPng update messages received. • **Bad Route Entries**—Number of RIPng invalid route entry messages received. • **Updates Ignored**—Number of RIPng update messages ignored. • **RIPng Requests Received**—Number of RIPng request messages received. • **RIPng Requests Ignored**—Number of RIPng request messages ignored.
Total	Total number of packets for the selected counter.
Last 5 min	Number of packets for the selected counter in the most recent 5-minute period.
Last minute	Number of packets for the selected counter in the most recent 1-minute period.

Sample Output

show ripng statistics

```
user@host> show ripng statistics
RIPng info: port 521; holddown 120s;
    rts learned  rts held down  rqsts dropped  resps dropped
              0             0              0              0

so-0/1/3.0:  0 routes learned; 1 routes advertised; timeout 180s; update interval
 20s
Counter                         Total   Last 5 min  Last minute
-------                   -----------  -----------  -----------
```

Copyright © 2017, Juniper Networks, Inc.

Updates Sent	934	16	4
Triggered Updates Sent	1	0	0
Responses Sent	0	0	0
Bad Messages	0	0	0
Updates Received	0	0	0
Bad Route Entries	0	0	0
Updates Ignored	0	0	0
RIPng Requests Received	0	0	0
RIPng Requests Ignored	0	0	0

Copyright © 2017, Juniper Networks, Inc.

show route

List of Syntax
Syntax on page 479
Syntax (EX Series Switches) on page 479

Syntax

show route
<all>
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>
<private>
<te-ipv4-prefix-ip *te-ipv4-prefix-ip*>
<te-ipv4-prefix-node-ip *te-ipv4-prefix-node-ip*>
<te-ipv4-prefix-node-iso *te-ipv4-prefix-node-iso*>

Syntax (EX Series Switches)

show route
<all>
<*destination-prefix*>
<private>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Option **private** introduced in Junos OS Release 9.5.
Option **private** introduced in Junos OS Release 9.5 for EX Series switches.
Command introduced in Junos OS Release 15.1R3 on MX Series routers for enhanced subscriber management.
Option **display-client-data** introduced in Junos OS Release 16.2R1 on MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series routers.
Options **te-ipv4-prefix-ip**, **te-ipv4-prefix-node-ip**, and **te-ipv4-prefix-node-iso** introduced in Junos OS Release 17.2R1 on MX Series and PTX Series.

Description

Display the active entries in the routing tables.

Options

none—Display brief information about all active entries in the routing tables.

all—(Optional) Display information about all routing tables, including private, or internal, routing tables.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

private—(Optional) Display information only about all private, or internal, routing tables.

display-client-data —(Optional) Display client id and cookie information for routes installed by the routing protocol process client applications.

te-ipv4-prefix-ip ***te-ipv4-prefix-ip***—(Optional) Display IPv4 address of the traffic-engineering prefix, without the mask length if present in the routing table.

Copyright © 2017, Juniper Networks, Inc.

te-ipv4-prefix-node-ip ***te-ipv4-prefix-node-ip***—(Optional) Display all prefixes that have originated from the traffic-engineering node. You can filter IPv4 node addresses from the traffic-engineered routes in the **lsdist.0** table.

te-ipv4-prefix-node-iso ***te-ipv4-prefix-node-iso***—(Optional) Display all prefixes that have originated from the traffic-engineering node. You can filter IPv4 routes with the specified ISO circuit ID from the **lsdist.0** table.

Required Privilege Level view

Related Documentation

- *Understanding IS-IS Configuration*
- *Example: Configuring IS-IS*
- *Examples: Configuring Internal BGP Peering*
- *Examples: Configuring External BGP Peering*
- *Examples: Configuring OSPF Routing Policy*
- *Verifying and Managing Junos OS Enhanced Subscriber Management*

List of Sample Output

show route on page 483
show route (VPN) on page 484
show route (with Destination Prefix) on page 484
show route destination-prefix detail on page 484
show route extensive on page 484
show route extensive (ECMP) on page 485
show route extensive (Multipath Resolution) on page 485
show route (Enhanced Subscriber Management) on page 489
show route (IPv6 Flow Specification) on page 490
show route display-client-data detail on page 490
show route te-ipv4-prefix-ip on page 491
show route te-ipv4-prefix-ip extensive on page 491
show route te-ipv4-prefix-node-iso on page 494
show route te-ipv4-prefix-node-iso extensive on page 494
show route te-ipv4-prefix-node-iso detail on page 497

Output Fields Table 15 on page 261 describes the output fields for the **show route** command. Output fields are listed in the approximate order in which they appear.

Table 36: show route Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.

Copyright © 2017, Juniper Networks, Inc.

Table 36: show route Output Fields *(continued)*

Field Name	Field Description
***number* routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). A holddown route was once the active route and is no longer the active route. The route is in the holddown state because a protocol still has interest in the route, meaning that the interest bit is set. A protocol might have its interest bit set on the previously active route because the protocol is still advertising the route. The route will be deleted after all protocols withdraw their advertisement of the route and remove their interest bit. A persistent holddown state often means that the interested protocol is not releasing its interest bit properly. However, if you have configured advertisement of multiple routes (with the **add-path** or **advertise-inactive** statement), the holddown bit is most likely set because BGP is advertising the route as an active route. In this case, you can ignore the holddown state because nothing is wrong. • **hidden** (routes that are not used because of a routing policy).
destination-prefix	Route destination (for example:10.0.0.1/24). Sometimes the route information is presented in another format, such as: • ***MPLS-label*** (for example, 80001). • ***interface-name*** (for example, ge-1/0/2). • ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id* :*source*** (Layer 2 circuit only. For example, 10.1.1.195:NoCtrlWord:1:1:Local/96): ◦ ***neighbor-address***—Address of the neighbor. ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**. ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. ◦ ***vc-id***—Virtual circuit identifier. ◦ ***source***—Source of the advertisement: **Local** or **Remote**.
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
weeks:days hours:minutes:seconds	How long the route been known (for example, **2w4d 13:11:14**, or 2 weeks, 4 days, 13 hours, 11 minutes, and 14 seconds).
metric	Cost value of the indicated route. For routes within an AS, the cost is determined by the IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.

Copyright © 2017, Juniper Networks, Inc.

Table 36: show route Output Fields *(continued)*

Field Name	Field Description
localpref	Local preference value included in the route.
from	Interface from which the route was received.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
encapsulated	Extended next-hop encoding capability enabled for the specified BGP community for routing IPv4 traffic over IPv6 tunnels. When BGP receives routes without the tunnel community, IPv4-0ver IPv6 tunnels are not created and BGP routes are resolved without encapsulation.
Route Labels	Stack of labels carried in the BGP route update.
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
to	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the destination is **Discard**, traffic is dropped.

Copyright © 2017, Juniper Networks, Inc.

Table 36: show route Output Fields *(continued)*

Field Name	Field Description
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • **Weight**—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • **Balance**—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing. • ***lsp-path-name***—Name of the LSP used to reach the next hop. • ***label-action***—MPLS label and operation occurring at the next hop. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label). For VPNs, expect to see multiple **push** operations, corresponding to the inner and outer labels required for VPN routes (in the case of a direct PE-to-PE connection, the VPN route would have the inner label push only).
Private unicast	(Enhanced subscriber management for MX Series routers) Indicates that an access-internal route is managed by enhanced subscriber management. By contrast, access-internal routes *not* managed by enhanced subscriber management are displayed with associated next-hop and media access control (MAC) address information.
balance	Distribution of the load based on the underlying operational interface bandwidth for equal-cost multipaths (ECMP) across the nexthop gateways in percentages.

Sample Output

show route

```
user@host> show route
inet.0: 11 destinations, 12 routes (11 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:65500:1:10.0.0.20/240
                   *[MVPN/70] 19:53:41, metric2 1
                      Indirect
1:65500:1:10.0.0.40/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
                    [BGP/170] 19:53:26, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.24.4 via lt-0/3/0.24, label-switched-path toD
1:65500:1:10.0.0.60/240
                   *[BGP/170] 19:53:29, localpref 100, from 10.0.0.30
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
                    [BGP/170] 19:53:25, localpref 100, from 10.0.0.33
                      AS path: I
                    > to 10.0.28.8 via lt-0/3/0.28, label-switched-path toF
```

Copyright © 2017, Juniper Networks, Inc.

show route (VPN)

The following sample output shows a VPN route with composite next hops enabled. The first **Push** operation corresponds to the outer label. The second **Push** operation corresponds to the inner label.

```
user@host> show route 192.0.2.0

13979:665001.inet.0: 871 destinations, 3556 routes (871 active, 0 holddown, 0
hidden)
+ = Active Route, - = Last Active, * = Both

192.0.2.0/24        [BGP/170] 00:28:32, localpref 100, from 10.9.9.160
                      AS path: 13980 ?, validation-state: unverified
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                    [BGP/170] 00:28:28, localpref 100, from 10.9.9.169
                      AS path: 13980 ?, validation-state: unverified
                    > to 10.100.0.42 via ae2.0, Push 126016, Push 300368(top)
                    #[Multipath/255] 00:28:28, metric2 102
                    > to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
                      to 10.100.0.42 via ae2.0, Push 16, Push 300368(top)
```

show route (with Destination Prefix)

```
user@host> show route 172.16.0.0/12

inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.0.0/12      *[Static/5] 2w4d 12:54:27
                    > to 192.168.167.254 via fxp0.0
```

show route destination-prefix detail

```
user@host> show route 198.51.100.0 detail

inet.0: 15 destinations, 20 routes (15 active, 0 holddown, 0 hidden)
198.51.100.0/24 (2 entries, 2 announced)
        *BGP    Preference: 170/-101
         ...
         BGP-Static Preference: 4294967292
               Next hop type: Discard
                 Address: 0x9041ae4
                 Next-hop reference count: 2
                 State: <NoReadvrt Int Ext AlwaysFlash>
        Inactive reason: Route Preference
        Local AS:   200
        Age: 4d 1:40:40
        Validation State: unverified
        Task: RT
        Announcement bits (1): 2-BGP_RT_Background
        AS path: 4 5 6 I
```

show route extensive

```
user@host> show route extensive
v1.mvpn.0: 5 destinations, 8 routes (5 active, 1 holddown, 0 hidden)
1:65500:1:10.0.0.40/240 (1 entry, 1 announced)
    *BGP   Preference: 170/-101
```


Copyright © 2017, Juniper Networks, Inc.

```
            PMSI: Flags 0x0: Label[0:0:0]: PIM-SM: Sender 10.0.0.40 Group
203.0.113.1
            Next hop type: Indirect
            Address: 0x92455b8
            Next-hop reference count: 2
            Source: 10.0.0.30
            Protocol next hop: 10.0.0.40
            Indirect next hop: 2 no-forward
            State: <Active Int Ext>
                     Local AS: 64510 Peer AS: 64511
            Age: 3  Metric2: 1
           Validation State: unverified
            Task: BGP_64510.10.0.0.30+179
            Announcement bits (2): 0-PIM.v1 1-mvpn global task
            AS path: I (Originator) Cluster list:  10.0.0.30
            AS path:  Originator ID: 10.0.0.40
            Communities: target:64502:100 encapsulation:0L:14           Import
Accepted
            Localpref: 100
            Router ID: 10.0.0.30
            Primary Routing Table bgp.mvpn.0
            Indirect next hops: 1
                 Protocol next hop: 10.0.0.40 Metric: 1
                 Indirect next hop: 2 no-forward
                 Indirect path forwarding next hops: 1
                       Next hop type: Router
                       Next hop: 10.0.24.4 via lt-0/3/0.24 weight 0x1
                 10.0.0.40/32 Originating RIB: inet.3
                       Metric: 1                       Node path count: 1
                       Forwarding nexthops: 1
                          Nexthop: 10.0.24.4 via lt-0/3/0.24
```

show route extensive (ECMP)

```
user@host> show route extensive
*IS-IS Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 1048577
                Address: 0xXXXXXXXXXX
                Next-hop reference count: YY
                Next hop: 172.16.50.2 via ae1.0 balance 43%, selected
                Session Id: 0x141
                Next hop: 192.0.2.2 via ae0.0 balance 57%
```

show route extensive (Multipath Resolution)

```
user@host> show route extensive
inet.0: 37 destinations, 37 routes (36 active, 0 holddown, 1 hidden)
10.1.1.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.1.1.2/32 -> {indirect(1048574)}
        *Static Preference: 5
                Next hop type: Indirect, Next hop index: 0
                Address: 0xb39d1b0
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 581
                Next hop: 12.1.1.2 via ge-2/0/1.0, selected
                Session Id: 0x144
                Next hop: 13.1.1.2 via ge-2/0/2.0, selected
                Session Id: 0x145
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Protocol next hop: 10.1.1.1
                    Indirect next hop: 0xb2b20f0 1048574 INH Session ID: 0x143
                    State: <Active Int Ext>
                    Age: 2:53  Metric2: 0
                    Validation State: unverified
                    Task: RT
                    Announcement bits (2): 0-KRT 2-Resolve tree 1
                    AS path: I
                    Indirect next hops: 1
                            Protocol next hop: 10.1.1.1
                          Indirect next hop: 0xb2b20f0 1048574 INH Session ID: 0x143

                            Indirect path forwarding next hops: 2
                                    Next hop type: Router
                                    Next hop: 12.1.1.2 via ge-2/0/1.0
                                    Session Id: 0x144
                                    Next hop: 13.1.1.2 via ge-2/0/2.0
                                    Session Id: 0x145
                         10.1.1.1/32 Originating RIB: inet.0
                            Node path count: 1
                              Node flags: 1
                            Forwarding nexthops: 2 (Merged)
                            Nexthop: 12.1.1.2 via ge-2/0/1.0
                             Nexthop: 13.1.1.2 via ge-2/0/2.0
```

```
user@host> show route active-path extensive
user@host> show route 141.1.1.1 active-path extensive

inet.0: 1000061 destinations, 1000082 routes (1000061 active, 0 holddown, 0 hidden)
141.1.1.1/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 141.1.1.1/32 -> {indirect(1051215)}
unicast reverse-path: 0
[ae0.0 ae1.0]
Page 0 idx 0, (group Internet-IPv4 type External) Type 1 val 0xbb2e53d8 (adv_entry)
Advertised metrics:
Nexthop: Self
AS path: [500] 410 I
Communities:
Path 141.1.1.1 from 50.0.0.11 Vector len 4. Val: 0
*BGP Preference: 170/-101
Next hop type: Indirect, Next hop index: 0
Address: 0x2e9aacdc
Next-hop reference count: 500000
Source: 50.0.0.11
Next hop type: Router, Next hop index: 0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 25, Push 20(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 20: None;
Label element ptr: 0xb5dc1780
Label parent element ptr: 0x18d48080
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 25, Push 22(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 22: None;
Label element ptr: 0xb5dc1700
Label parent element ptr: 0x18d41000
```


Copyright © 2017, Juniper Networks, Inc.

```
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 24, Push 48(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 48: None;
Label element ptr: 0x18d40800
Label parent element ptr: 0x18d49780
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Label operation: Push 3851, Push 24, Push 49(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 49: None;
Label element ptr: 0xb5dc1680
Label parent element ptr: 0x18d48f00
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Label operation: Push 3851, Push 25, Push 21(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 21: None;
Label element ptr: 0xb5dc1600
Label parent element ptr: 0x18d44d80
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Label operation: Push 3851, Push 25, Push 25(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 25: None; Label 25: None;
Label element ptr: 0xb5dc1580
Label parent element ptr: 0x18d3da80
Label element references: 2
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1, selected
Label operation: Push 3851, Push 24, Push 68(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 68: None;
Label element ptr: 0x18d41500
Label parent element ptr: 0x18d49000
Label element references: 3
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Label operation: Push 3851, Push 24, Push 69(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 3851: None; Label 24: None; Label 69: None;
Label element ptr: 0xb5dc1500
Label parent element ptr: 0x18d48300
Label element references: 2
```

Copyright © 2017, Juniper Networks, Inc.

```
Label element child references: 0
Label element lsp id: 0
Session Id: 0x0
Protocol next hop: 50.0.0.11
Label operation: Push 3851
Label TTL action: prop-ttl
Load balance label: Label 3851: None;
Indirect next hop: 0x1883e200 1051215 INH Session ID: 0xb0d
State:
Local AS: 500 Peer AS: 500
Age: 1:40:03 Metric2: 2
Validation State: unverified
Task: BGP_500.50.0.0.11
Announcement bits (5): 0-KRT 8-KRT 9-BGP_RT_Background 10-Resolve tree 5 11-Resolve
 tree 8
AS path: 410 I
Accepted
Route Label: 3851
Localpref: 100
Router ID: 50.0.0.11
Indirect next hops: 1
Protocol next hop: 50.0.0.11 Metric: 2
Label operation: Push 3851
Label TTL action: prop-ttl
Load balance label: Label 3851: None;
Indirect next hop: 0x1883e200 1051215 INH Session ID: 0xb0d
Indirect path forwarding next hops (Merged): 8
Next hop type: Router
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.12.2 via ae0.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
Next hop: 50.0.13.3 via ae1.0 weight 0x1
Session Id: 0x0
50.0.0.11/32 Originating RIB: inet.3
Metric: 1 Node path count: 4
Node flags: 1
Indirect nexthops: 4
Protocol Nexthop: 50.0.0.4 Metric: 1 Push 24
Indirect nexthop: 0x1880f200 1048597 INH Session ID: 0xb0c
Path forwarding nexthops link: 0x36120400
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
50.0.0.4/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
```


Copyright © 2017, Juniper Networks, Inc.

```
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
Protocol Nexthop: 50.0.0.5 Metric: 1 Push 24
Indirect nexthop: 0x18810000 1048596 INH Session ID: 0xb0b
Path forwarding nexthops link: 0x1545be00
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
50.0.0.5/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
Protocol Nexthop: 50.0.0.6 Metric: 1 Push 25
Indirect nexthop: 0x1880e600 1048588 INH Session ID: 0xb0a
Path forwarding nexthops link: 0x3611f440
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
50.0.0.6/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
Protocol Nexthop: 50.0.0.7 Metric: 1 Push 25
Indirect nexthop: 0x1880dc00 1048586 INH Session ID: 0xb09
Path forwarding nexthops link: 0x15466d80
Path inh link: 0x0
Indirect path forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
50.0.0.7/32 Originating RIB: inet.3
Metric: 1 Node path count: 1
Forwarding nexthops: 2
Nexthop: 50.0.12.2 via ae0.0
Session Id: 0
Nexthop: 50.0.13.3 via ae1.0
Session Id: 0
```

show route (Enhanced Subscriber Management)

```
user@host> show route
inet.0: 41 destinations, 41 routes (40 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

198.51.100.11/24    *[Access-internal/12] 00:00:08
                     > to #0 10.0.0.1.93.65 via demux0.1073741824
```

```
198.51.100.12/24    *[Access-internal/12] 00:00:08
                    Private unicast
```

show route (IPv6 Flow Specification)

```
user@host> show route
inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::10:255:185:19/128
                   *[Direct/0] 05:11:27
                    > via lo0.0
2001:db8::11:11:11:0/120
                   *[BGP/170] 00:28:58, localpref 100
                      AS path: 2000 I, validation-state: unverified
                    > to 2001:db8::13:14:2:2 via ge-1/1/4.0
2001:db8::13:14:2:0/120*[Direct/0] 00:45:07
                    > via ge-1/1/4.0
2001:db8::13:14:2:1/128*[Local/0] 00:45:18
                      Local via ge-1/1/4.0
fe80::2a0:a50f:fc71:71d5/128
                   *[Direct/0] 05:11:27
                    > via lo0.0
fe80::5e5e:abff:feb0:933e/128
                   *[Local/0] 00:45:18
                      Local via ge-1/1/4.0

inet6flow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::11:11:11:10/128,*,proto=6,dstport=80,srcport=65535/term:1
                   *[BGP/170] 00:28:58, localpref 100, from 2001:db8::13:14:2:2
                      AS path: 2000 I, validation-state: unverified
                      Fictitious
2001:db8::11:11:11:30/128,*,icmp6-type=128,len=100,dscp=10/term:2
                   *[BGP/170] 00:20:54, localpref 100, from 2001:db8::13:14:2:2
                      AS path: 2000 I, validation-state: unverified
                      Fictitious
```

show route display-client-data detail

```
user@host> show route 198.51.100.0/24 display-client-data detail
inet.0: 59 destinations, 70 routes (59 active, 0 holddown, 0 hidden)
198.51.100.0/24 (1 entry, 1 announced)
        State: <FlashAll>
        *BGP-Static Preference: 5/-101
                Next hop type: Indirect, Next hop index: 0
                Address: 0xa5c2af8
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 1641
                Next hop: 192.0.2.1 via ge-2/1/1.0, selected
                Session Id: 0x160
                Protocol next hop: 192.0.2.1
                Indirect next hop: 0xa732cb0 1048621 INH Session ID: 0x17e
                State: <Active Int Ext AlwaysFlash NSR-incapable Programmed>
                Age: 3:13       Metric2: 0
                Validation State: unverified
                Announcement bits (3): 0-KRT 5-LDP 6-Resolve tree 3
```


Copyright © 2017, Juniper Networks, Inc.

```
                    AS path: I
                    Client id: 1, Cookie: 1
```

show route te-ipv4-prefix-ip

```
user@host> show route te-ipv4-prefix-ip 10.10.10.10
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both


PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152
                    *[IS-IS/15] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0101.0101.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0202.0202.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0303.0303.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0404.0404.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0505.0505.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0606.0606.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0707.0707.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                    *[IS-IS/18] 00:01:01
                       Fictitious
```

show route te-ipv4-prefix-ip extensive

```
user@host>show route te-ipv4-prefix-ip 10.10.10.10 extensive
lsdist.0: 298 destinations, 298 routes (298 active, 0 holddown, 0 hidden)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State:<Active NotInstall>
                Local AS:   100
                Age: 7:58
```

Copyright © 2017, Juniper Networks, Inc.

```
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0101.0101.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
       State: <Active NotInstall
       Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0>

PREFIX { Node { AS:100 ISO:0100.0202.0202.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
       Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0303.0303.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0404.0404.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
```


Copyright © 2017, Juniper Networks, Inc.

```
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0505.0505.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0606.0606.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0707.0707.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 298
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 7:58
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
```

```
                    Address: 0xa1a2ac4
                    Next-hop reference count: 298
                    Next hop:
                    State: <Active NotInstall>
                    Local AS:   100
                    Age: 7:58
                    Validation State: unverified
                    Task: IS-IS
                    AS path: I
                    Prefix SID: 1000, Flags: 0x40, Algo: 0
```

show route te-ipv4-prefix-node-iso

```
user@host> show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152
                   *[IS-IS/15] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:1.1.1.1/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:2.2.2.2/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:3.3.3.3/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:4.4.4.4/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:5.5.5.5/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:6.6.6.6/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:7.7.7.7/32 } ISIS-L2:0
}/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152
                   *[IS-IS/18] 00:05:20
                      Fictitious
```

show route te-ipv4-prefix-node-iso extensive

```
user@host> show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00 extensive
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
```


Copyright © 2017, Juniper Networks, Inc.

```
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:1.1.1.1/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1001, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:2.2.2.2/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1002, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:3.3.3.3/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
```

Copyright © 2017, Juniper Networks, Inc.

```
                Prefix SID: 1003, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:4.4.4.4/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1004, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:5.5.5.5/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1005, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:6.6.6.6/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1006, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:7.7.7.7/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
```


Copyright © 2017, Juniper Networks, Inc.

```
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1007, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:47
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0
```

show route te-ipv4-prefix-node-iso detail

```
user@host> show route te-ipv4-prefix-node-iso 0100.0a0a.0a0a.00 detail
lsdist.0: 283 destinations, 283 routes (283 active, 0 holddown, 0 hidden)
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L1:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:1.1.1.1/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1001, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:2.2.2.2/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
```

Copyright © 2017, Juniper Networks, Inc.

```
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1002, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:3.3.3.3/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1003, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:4.4.4.4/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1004, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:5.5.5.5/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1005, Flags: 0xe0, Algo: 0
```


Copyright © 2017, Juniper Networks, Inc.

```
PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:6.6.6.6/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1006, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:7.7.7.7/32 } ISIS-L2:0
}/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1007, Flags: 0xe0, Algo: 0

PREFIX { Node { AS:100 ISO:0100.0a0a.0a0a.00 } { IPv4:10.10.10.10/32 } ISIS-L2:0
 }/1152 (1 entry, 0 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Fictitious, Next hop index: 0
                Address: 0xa1a2ac4
                Next-hop reference count: 283
                Next hop:
                State: <Active NotInstall>
                Local AS:   100
                Age: 6:54
                Validation State: unverified
                Task: IS-IS
                AS path: I
                Prefix SID: 1000, Flags: 0x40, Algo: 0
```

Copyright © 2017, Juniper Networks, Inc.

show route active-path

List of Syntax
Syntax on page 500
Syntax (EX Series Switches) on page 500

Syntax
show route active-path
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route active-path
<brief | detail | extensive | terse>

Release Information
Command introduced in Junos OS Release 8.0.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display all active routes for destinations. An active route is a route that is selected as the best path. Inactive routes are not displayed.

Options
none—Display all active routes.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route active-path on page 500
show route active-path brief on page 501
show route active-path detail on page 501
show route active-path extensive on page 502
show route active-path terse on page 504

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route active-path

```
user@host> show route active-path

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.70.19/32    *[Direct/0] 21:33:52
```


Copyright © 2017, Juniper Networks, Inc.

```
                    > via lo0.0
10.255.71.50/32    *[IS-IS/15] 00:18:13, metric 10
                    > to 172.16.100.1 via so-2/1/3.0
172.16.100.1/24      *[Direct/0] 00:18:36
                    > via so-2/1/3.0
172.16.100.1/32      *[Local/0] 00:18:41
                      Local via so-2/1/3.0
192.168.64.0/21    *[Direct/0] 21:33:52
                    > via fxp0.0
192.168.70.19/32   *[Local/0] 21:33:52
                      Local via fxp0.0
```

show route active-path brief

The output for the **show route active-path brief** command is identical to that for the **show route active-path** command. For sample output, see show route active-path on page 281.

show route active-path detail

```
user@host> show route active-path detail

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)

10.255.70.19/32 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via lo0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

10.255.71.50/32 (1 entry, 1 announced)
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 397
                Next-hop reference count: 4
                Next hop: 172.16.100.1 via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:31      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
                AS path: I

172.16.100.0/24 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:54
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3
```

Copyright © 2017, Juniper Networks, Inc.

```
                AS path: I

172.16.100.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: so-2/1/3.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 21:59
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.64.0/21 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via fxp0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.70.19/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: fxp0.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 21:37:10
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I
```

show route active-path extensive

```
user@host> show route active-path extensive

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
10.255.70.19/32 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via lo0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

10.255.71.50/32 (1 entry, 1 announced)
```


Copyright © 2017, Juniper Networks, Inc.

```
TSI:
KRT in-kernel 10.255.71.50/32 -> {172.16.100.1}
IS-IS level 2, LSP fragment 0
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 397
                Next-hop reference count: 4
                Next hop: 172.16.100.1 via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 24:08      Metric: 10
                Task: IS-IS
                Announcement bits (4): 0-KRT 2-IS-IS 5-Resolve tree 2 6-Resolve
tree 3
                AS path: I

172.16.100.1/24 (1 entry, 1 announced)
TSI:
IS-IS level 1, LSP fragment 0
IS-IS level 2, LSP fragment 0
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via so-2/1/3.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 24:31
                Task: IF
                Announcement bits (3): 2-IS-IS 5-Resolve tree 2 6-Resolve tree 3

                AS path: I

172.16.100.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: so-2/1/3.0
                State: ‹Active NoReadvrt Int›
                Local AS:   200
                Age: 24:36
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.64.0/21 (1 entry, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 3
                Next hop: via fxp0.0, selected
                State: ‹Active Int›
                Local AS:   200
                Age: 21:39:47
                Task: IF
                Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
                AS path: I

192.168.70.19/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 11
                Interface: fxp0.0
```

Copyright © 2017, Juniper Networks, Inc.

```
State: ‹Active NoReadvrt Int›
Local AS:   200
Age: 21:39:47
Task: IF
Announcement bits (2): 5-Resolve tree 2 6-Resolve tree 3
AS path: I
```

show route active-path terse

```
user@host> show route active-path terse

inet.0: 7 destinations, 7 routes (6 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.19/32    D   0                      >lo0.0
* 10.255.71.50/32    I  15         10           >172.16.100.1.
* 172.16.100.0/24       D   0                      >so-2/1/3.0
* 172.16.100.2/32       L   0                       Local
* 192.168.64.0/21    D   0                      >fxp0.0
* 192.168.70.19/32   L   0                       Local
```


Copyright © 2017, Juniper Networks, Inc.

show route advertising-protocol

Syntax show route advertising-protocol *protocol neighbor-address*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Release Information Command introduced before Junos OS Release 7.4.

Description Display the routing information as it has been prepared for advertisement to a particular neighbor of a particular dynamic routing protocol.

Options **brief | detail | extensive | terse**—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

neighbor-address—Address of the neighboring router to which the route entry is being transmitted.

protocol—Protocol transmitting the route:

- **bgp**—Border Gateway Protocol
- **dvmrp**—Distance Vector Multicast Routing Protocol
- **msdp**—Multicast Source Discovery Protocol
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation

Additional Information Routes displayed are routes that the routing table has exported into the routing protocol and that have been filtered by the associated protocol's **export** routing policy statements. Starting with Junos OS Release 13.3, you can display the routing instance table **foo** for any address family, on a VPN route reflector, or a VPN AS boundary router that is advertising local VPN routes. However, If you do not specify the **table** in the command, the output displays each VRF prefix twice.

Required Privilege Level view

Related Documentation
- *Example: Configuring the MED Attribute That Determines the Exit Point in an AS*

List of Sample Output
show route advertising-protocol bgp (Layer 3 VPN) on page 508
show route advertising-protocol bgp detail on page 508

Copyright © 2017, Juniper Networks, Inc.

show route advertising-protocol bgp detail (Aggregate Extended Community Bandwidth) on page 508
show route advertising-protocol bgp detail (Labeled Unicast) on page 509
show route advertising-protocol bgp detail (Layer 2 VPN) on page 509
show route advertising-protocol bgp detail (Layer 3 VPN) on page 509
show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address) on page 510

Output Fields Table 16 on page 287 lists the output fields for the **show route advertising-protocol** command. Output fields are listed in the approximate order in which they appear.

Table 37: show route advertising-protocol Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example, inet.0.	All levels
number **destinations**	Number of destinations for which there are routes in the routing table.	All levels
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)	All levels
Prefix	Destination prefix.	**brief** none
destination-prefix **(entry , announced)**	Destination prefix. The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination.	**detail extensive**
BGP group and type	BGP group name and type (**Internal** or **External**).	**detail extensive**
Route Distinguisher	Unique 64-bit prefix augmenting each IP subnet.	**detail extensive**
Advertised Label	Incoming label advertised by the Label Distribution Protocol (LDP). When an IP packet enters a label-switched path (LSP), the ingress router examines the packet and assigns it a label based on its destination, placing the label in the packet's header. The label transforms the packet from one that is forwarded based on its IP routing information to one that is forwarded based on information associated with the label.	**detail extensive**
Label-Base, range	First label in a block of labels and label block size. A remote PE router uses this first label when sending traffic toward the advertising PE router.	**detail extensive**
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routers by advertising VPN labels. VPN labels transit over either a Resource Reservation Protocol (RSVP) or a Label Distribution Protocol (LDP) label-switched path (LSP) tunnel.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 37: show route advertising-protocol Output Fields *(continued)*

Field Name	Field Description	Level of Output
Nexthop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route. If the next-hop advertisement to the peer is **Self**, and the RIB-out next hop is a specific IP address, the RIB-out IP address is included in the extensive output. See show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address) on page 291.	All levels
MED	Multiple exit discriminator value included in the route.	**brief**
Lclpref or **Localpref**	Local preference value included in the route.	All levels
Queued	When BGP route prioritization is enabled and a route is present in a priority queue, this shows which priority queue the route is in.	All levels except **brief**
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if configured on the router, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.	All levels
Route Labels	Stack of labels carried in the BGP route update.	**detail extensive**
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	**detail extensive**
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	**detail extensive**
Communities	Community path attribute for the route. See the output field table for the show route detail command for all possible values for this field.	**detail extensive**
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 37: show route advertising-protocol Output Fields *(continued)*

Field Name	Field Description	Level of Output
Attrset AS	Number, local preference, and path of the autonomous system (AS) that originated the route. These values are stored in the **Attrset** attribute at the originating router.	**detail extensive**
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	**detail extensive**
control flags	Control flags: **none** or **Site Down**.	**detail extensive**
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	**detail extensive**

Sample Output

show route advertising-protocol bgp (Layer 3 VPN)

```
user@host> show route advertising-protocol bgp 10.255.14.171
      VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.172/32   Self                     1        100 I
      VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop                MED    Lclpref AS path
      10.255.14.181/32   Self                     2        100 I
```

show route advertising-protocol bgp detail

```
user@host> show route advertising-protocol bgp 111.222.1.3 detail
bgp20.inet.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
111.222.1.11/32 (1 entry, 1 announced)
 BGP group pe-pe type Internal
     Route Distinguisher: 111.255.14.11:69
     Advertised Label: 100000
     next hop: Self
     Localpref: 100
     AS path: 2 I
     Communities: target:69:20
     AIGP 210
111.8.0.0/16 (1 entry, 1 announced)
 BGP group pe-pe type Internal
     Route Distinguisher: 111.255.14.11:69
     Advertised Label: 100000
     Next hop: Self
     Localpref: 100
     AS path: 2 I
     Communities: target:69:20
     AIGP 210
```

show route advertising-protocol bgp detail (Aggregate Extended Community Bandwidth)

```
user@host> show route advertising-protocol bgp 10.0.4.2 10.0.2.0/30 detail
inet.0: 20 destinations, 26 routes (20 active, 0 holddown, 0 hidden)
* 10.0.2.0/30 (2 entries, 1 announced)
 BGP group external2 type External
     Nexthop: Self
```


Copyright © 2017, Juniper Networks, Inc.

```
    AS path: [65000] 65001 I
    Communities: bandwidth:65000:80000000
```

show route advertising-protocol bgp detail (Labeled Unicast)

```
user@host>show route advertising bgp 1.1.1.3 detail
inet.0: 69 destinations, 70 routes (69 active, 0 holddown, 0 hidden)
* 1.1.1.8/32 (2 entries, 2 announced)
BGP group ibgp type Internal
Route Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: 1.1.1.4
MED: 7
Localpref: 100
AS path: [5] I
Cluster ID: 3.3.3.3
Originator ID: 1.1.1.1
Entropy label capable
inet6.0: 26 destinations, 28 routes (26 active, 0 holddown, 0 hidden)
* 100::1/128 (2 entries, 1 announced)
BGP group ibgp type Internal
Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: ::ffff:1.1.1.4
Localpref: 100
AS path: [5] I
Cluster ID: 3.3.3.3
Originator ID: 1.1.1.1
```

show route advertising-protocol bgp detail (Layer 2 VPN)

```
user@host> show route advertising-protocol bgp 192.168.24.1 detail
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
192.168.16.1:1:1:1/96 (1 entry, 1 announced)
 BGP group int type Internal
     Route Distinguisher: 192.168.16.1:1
     Label-base : 32768, range : 3
     Nexthop: Self
     Localpref: 100
     AS path: I
     Communities: target:65412:100
     AIGP 210
     Layer2-info: encaps:VLAN, control flags:, mtu:
```

show route advertising-protocol bgp detail (Layer 3 VPN)

```
user@host>  show route advertising-protocol bgp 10.255.14.176 detail
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
 BGP group ibgp type Internal
     Route Distinguisher: 10.255.14.174:2
     VPN Label: 101264
     Nexthop: Self
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AIGP 210
     AttrSet AS: 100
         Localpref: 100
         AS path: I
...
```

Copyright © 2017, Juniper Networks, Inc.

show route advertising-protocol bgp extensive all (Next Hop Self with RIB-out IP Address)

```
user@host> show route advertising-protocol bgp 200.0.0.2 170.0.1.0/24 extensive all
inet.0: 13 destinations, 19 routes (13 active, 0 holddown, 6 hidden)
  170.0.1.0/24 (2 entries, 1 announced)
 BGP group eBGP-INTEROP type External
     Nexthop: Self (rib-out 10.100.3.2)
     AS path: [4713] 200 I
...
```


Copyright © 2017, Juniper Networks, Inc.

show route all

List of Syntax
Syntax on page 511
Syntax (EX Series Switches) on page 511

Syntax
show route all
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route all

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display information about all routes in all routing tables, including private, or internal, tables.

Options
none—Display information about all routes in all routing tables, including private, or internal, tables.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show route brief on page 297
- show route detail on page 299

List of Sample Output
show route all on page 511

Output Fields
In Junos OS Release 9.5 and later, only the output fields for the **show route all** command display all routing tables, including private, or hidden, routing tables. The output field table of the show route command does not display entries for private, or hidden, routing tables in Junos OS Release 9.5 and later.

Sample Output

show route all

The following example displays a snippet of output from the **show route** command and then displays the same snippet of output from the **show route all** command:

```
user@host> show route
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:24:39, metric 1
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Receive
1                  *[MPLS/0] 2d 02:24:39, metric 1
                    Receive
2                  *[MPLS/0] 2d 02:24:39, metric 1
                    Receive
800017             *[VPLS/7] 1d 14:00:16
                    > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 14:00:26
                    > via vt-3/2/0.32772, Pop

user@host> show route all
mpls.0: 7 destinations, 7 routes (5 active, 0 holddown, 2 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
0                  *[MPLS/0] 2d 02:19:12, metric 1
                    Receive
1                  *[MPLS/0] 2d 02:19:12, metric 1
                    Receive
2                  *[MPLS/0] 2d 02:19:12, metric 1
                    Receive
800017             *[VPLS/7] 1d 13:54:49
                    > via vt-3/2/0.32769, Pop
800018             *[VPLS/7] 1d 13:54:59
                    > via vt-3/2/0.32772, Pop
vt-3/2/0.32769      [VPLS/7] 1d 13:54:49
                    Unusable
vt-3/2/0.32772      [VPLS/7] 1d 13:54:59
                    Unusable
```


Copyright © 2017, Juniper Networks, Inc.

show route best

List of Syntax
Syntax on page 513
Syntax (EX Series Switches) on page 513

Syntax
show route best *destination-prefix*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route best *destination-prefix*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the route in the routing table that is the best route to the specified address or range of addresses. The best route is the longest matching route.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

destination-prefix—Address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show route brief on page 297
- show route detail on page 299

List of Sample Output
show route best on page 513
show route best detail on page 514
show route best extensive on page 515
show route best terse on page 515

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route best

```
user@host> show route best 10.255.70.103
```

Copyright © 2017, Juniper Networks, Inc.

```
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[OSPF/10] 1d 13:19:20, metric 2
                    > to 10.31.1.6 via ge-3/1/0.0
                      via so-0/3/0.0

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.255.70.103/32   *[RSVP/7] 1d 13:20:13, metric 2
                    > via so-0/3/0.0, label-switched-path green-r1-r3

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.0.0.0/8         *[Direct/0] 2d 01:43:34
                    > via fxp2.0
                    [Direct/0] 2d 01:43:34
                    > via fxp1.0
```

show route best detail

```
user@host> show route best 10.255.70.103 detail
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                Next hop: via so-0/3/0.0
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:06        Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 5
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100016
                State: <Active Int>
                Local AS:    69
                Age: 1d 13:20:59        Metric: 2
                Task: RSVP
                Announcement bits (1): 1-Resolve tree 2
                AS path: I

private1__inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
10.0.0.0/8 (2 entries, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp2.0, selected
                State: <Active Int>
```


Copyright © 2017, Juniper Networks, Inc.

```
            Age: 2d 1:44:20
            Task: IF
            AS path: I
       Direct Preference: 0
            Next hop type: Interface
            Next-hop reference count: 1
            Next hop: via fxp1.0, selected
            State: <NotBest Int>
            Inactive reason: No difference
            Age: 2d 1:44:20
            Task: IF
            AS path: I
```

show route best extensive

The output for the **show route best extensive** command is identical to that for the **show route best detail** command. For sample output, see show route best detail on page 295.

show route best terse

```
user@host> show route best 10.255.70.103 terse
inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   O  10          2            >10.31.1.6
                                                   so-0/3/0.0

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

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.70.103/32   R   7          2            >so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.0.0.0/8         D   0                        >fxp2.0
                     D   0                        >fxp1.0
```

Copyright © 2017, Juniper Networks, Inc.

show route brief

List of Syntax
Syntax on page 516
Syntax (EX Series Switches) on page 516

Syntax
show route brief
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route brief
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display brief information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show route all on page 292
- show route best on page 294

List of Sample Output
show route brief on page 516

Output Fields
For information about output fields, see the Output Field table of the show route command.

Sample Output

show route brief

```
user@host> show route brief
inet.0: 10 destinations, 10 routes (9 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 1w5d 20:30:29
                      Discard
10.255.245.51/32   *[Direct/0] 2w4d 13:11:14
                    > via lo0.0
```


Copyright © 2017, Juniper Networks, Inc.

```
172.16.0.0/12      *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.0.0/18      *[Static/5] 1w5d 20:30:29
                    > to 192.168.167.254 via fxp0.0
192.168.40.0/22    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.64.0/18    *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
192.168.164.0/22   *[Direct/0] 2w4d 13:11:14
                    > via fxp0.0
192.168.164.51/32  *[Local/0] 2w4d 13:11:14
                      Local via fxp0.0
207.17.136.192/32  *[Static/5] 2w4d 13:11:14
                    > to 192.168.167.254 via fxp0.0
green.inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
100.101.0.0/16     *[Direct/0] 1w5d 20:30:28
                    > via fe-0/0/3.0
100.101.2.3/32     *[Local/0] 1w5d 20:30:28
                      Local via fe-0/0/3.0
172.16.233.5/32       *[OSPF/10] 1w5d 20:30:29, metric 1
                      MultiRecv
```

Copyright © 2017, Juniper Networks, Inc.

show route detail

List of Syntax
Syntax on page 518
Syntax (EX Series Switches) on page 518

Syntax
show route detail
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route detail
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 13.2X51-D15 for the QFX Series.
Comand introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Display detailed information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table on all systems.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route detail on page 529
show route detail (with BGP Multipath) on page 535
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 536
show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 536

Output Fields
Table 17 on page 299 describes the output fields for the **show route detail** command. Output fields are listed in the approximate order in which they appear.

Table 38: show route detail Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.

Copyright © 2017, Juniper Networks, Inc.

Table 38: show route detail Output Fields *(continued)*

Field Name	Field Description
number **routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)
route-destination **(entry, announced)**	Route destination (for example:10.0.0.1/24). The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as: • ***MPLS-label*** (for example, 80001). • ***interface-name*** (for example, ge-1/0/2). • ***neighbor-address*****:*****control-word-status*****:*****encapsulation type*****:*****vc-id*****:*****source*** (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96). ◦ ***neighbor-address***—Address of the neighbor. ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**. ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport. ◦ ***vc-id***—Virtual circuit identifier. ◦ ***source***—Source of the advertisement: **Local** or **Remote**. ◦ ***source***—Source of the advertisement: **Local** or **Remote**.
label stacking	**(**Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).

Copyright © 2017, Juniper Networks, Inc.

Table 38: show route detail Output Fields *(continued)*

Field Name	Field Description
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • ***—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value. **Preference2** values are signed integers, that is, **Preference2** values can be either positive or negative values. However, Junos OS evaluates **Preference2** values as unsigned integers that are represented by positive values. Based on the **Preference2** values, Junos OS evaluates a preferred route differently in the following scenarios: • **Both Signed Preference2 values** ◦ Route A = -101 ◦ Route B = -156 Where both the **Preference2** values are signed, Junos OS evaluates only the unsigned value of **Preference2** and Route A, which has a lower **Preference2** value is preferred. • **Unsigned Preference2 values** Now consider both unsigned **Preference2** values: ◦ Route A = 4294967096 ◦ Route B = 200 Here, Junos OS considers the lesser **Preference2** value and Route B with a **Preference2** value of 200 is preferred because it is less than 4294967096. • **Combination of signed and unsigned Preference2 values** When **Preference2** values of two routes are compared, and for one route the **Preference2** is a signed value, and for the other route it is an unsigned value, Junos OS prefers the route with the positive **Preference2** value over the negative **Preference2** value. For example, consider the following signed and unsigned **Preference2** values: ◦ Route A = -200 ◦ Route B = 200 In this case, Route B with a **Preference2** value of 200 is preferred although this value is greater than -200, because Junos OS evaluates only the unsigned value of the **Preference2** value.
Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).

Copyright © 2017, Juniper Networks, Inc.

Table 38: show route detail Output Fields *(continued)*

Field Name	Field Description
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 18 on page 305.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 19 on page 307.
Local AS	AS number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.

Copyright © 2017, Juniper Networks, Inc.

Table 38: show route detail Output Fields *(continued)*

Field Name	Field Description
Metric*n*	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see show route table.
Task	Name of the protocol that has added the route.
Announcement bits	The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the KRT for installing the route into the Packet Forwarding Engine, to a resolve tree, a L2 VC, or even a VPN. For example, ***n*-Resolve inet** indicates that the specified route is used for route resolution for next hops found in the routing table. • ***n***—An index used by Juniper Networks customer support only.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.

Copyright © 2017, Juniper Networks, Inc.

Table 38: show route detail Output Fields *(continued)*

Field Name	Field Description
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
ORR Generation-ID	Displays the optimal route reflection (ORR) generation identifier. ISIS and OSPF interior gateway protocol (IGP) updates filed whenever any of the corresponding ORR route has its metric valued changed, or if the ORR route is added or deleted.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 20 on page 309 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or **Site Down**.
mtu	Maximum transmission unit (MTU) information.

Copyright © 2017, Juniper Networks, Inc.

Table 38: show route detail Output Fields *(continued)*

Field Name	Field Description
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 18 on page 305 describes all possible values for the Next-hop Types output field.

Table 39: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Dynamic List	Dynamic list next hop

Copyright © 2017, Juniper Networks, Inc.

Table 39: Next-hop Types Output Field Values *(continued)*

Next-Hop Type	Description
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.
Router	A specific node or set of nodes to which the routing device forwards packets that match the route prefix. To qualify as next-hop type router, the route must meet the following criteria: • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.

Copyright © 2017, Juniper Networks, Inc.

Table 39: Next-hop Types Output Field Values *(continued)*

Next-Hop Type	Description
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 19 on page 307 describes all possible values for the State output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).

Table 40: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.

Copyright © 2017, Juniper Networks, Inc.

Table 40: State Output Field Values *(continued)*

Value	Description
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.
Pending	Route pending because of a hold-down configured on another route.
Programmed	Route installed programatically by on-box or off-box applications using API.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available
Router ID	Path through a neighbor with lower ID is available.

Copyright © 2017, Juniper Networks, Inc.

Table 40: State Output Field Values *(continued)*

Value	Description
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.
ProtectionCand	Indicates paths requesting protection.
ProtectionPath	Indicates the route entry that can be used as a protection path.

Table 20 on page 309 describes the possible values for the Communities output field.

Table 41: Communities Output Field Values

Value	Description
area-number	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is **0**. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: ***local AS number:link-bandwidth-number***	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
link-bandwidth-number	Link-bandwidth number: from **0** through **4,294,967,295** (bytes per second).
local AS number	Local AS number: from **1** through **65,535**.
options	1 byte. Currently this is only used if the route type is **5** or **7**. Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	1 byte, encoded as **1** or **2** for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); **3** for summary routes; **5** for external routes (area number must be**0**); **7** for NSSA routes; or **129** for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x8000**. The format is ***area-number:ospf-route-type:options***.

Copyright © 2017, Juniper Networks, Inc.

Table 41: Communities Output Field Values *(continued)*

Value	Description
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x0306**. The format is ***area-number:ospf-route-type:options***.
target	Defines which VPN the route participates in; **target** has the format ***32-bit IP address:16-bit number***. For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between **0x1** and **0x7fff**. This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above **0x8000**. This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route detail

```
user@host> show route detail

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:31:43
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

10.31.1.0/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:30:17
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    69
                Age: 1:30:17    Metric: 1
                ORR Generation-ID: 1
      Area: 0.0.0.0
                Task: OSPF
                AS path: I
```

Copyright © 2017, Juniper Networks, Inc.

```
10.31.1.1/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:30:20
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

10.31.2.0/30 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 10.31.1.6 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:29:56    Metric: 2
                Area: 0.0.0.0
                ORR Generation-ID: 1
     Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

172.16.233.2/32 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

172.16.233.22/32 (1 entry, 1 announced)
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:43
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.255.70.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
```


Copyright © 2017, Juniper Networks, Inc.

```
                State: <Active Int>
                Local AS:    69
                Age: 1:25:49    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

10.255.71.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    69
                Age: 1:25:49    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
0   (1 entry, 1 announced)
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    69
                Age: 1:31:45    Metric: 1
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

  299840 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 299840 /52 -> {indirect(1048575)}
          *RSVP   Preference: 7/2
                  Next hop type: Flood
                  Address: 0x9174a30
                  Next-hop reference count: 4
                  Next hop type: Router, Next hop index: 798
                  Address: 0x9174c28
                  Next-hop reference count: 2
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Next hop: 172.16.0.2 via lt-1/2/0.9 weight 0x1
                    Label-switched-path R2-to-R4-2p2mp
                    Label operation: Pop
                    Next hop type: Router, Next hop index: 1048574
                    Address: 0x92544f0
                    Next-hop reference count: 2
                    Next hop: 172.16.0.2 via lt-1/2/0.7 weight 0x1
                    Label-switched-path R2-to-R200-p2mp
                    Label operation: Pop
                    Next hop: 172.16.0.2 via lt-1/2/0.5 weight 0x8001
                    Label operation: Pop
                    State: <Active Int>
                    Age: 1:29       Metric: 1
                    Task: RSVP
                    Announcement bits (1): 0-KRT
                    AS path: I...

800010 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:29:30
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 10.31.1.6 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 10.255.70.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:29:30    Metric2: 2
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

abcd::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
```


Copyright © 2017, Juniper Networks, Inc.

```
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:44
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:45
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    69
                Age: 1:31:43
                Task: MLD
                Announcement bits (1): 0-KRT
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:31:44
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

10.255.70.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.70.103:1
                Next-hop reference count: 7
                Source: 10.255.70.103
                Protocol next hop: 10.255.70.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    69 Peer AS:    69
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Age: 1:25:49    Metric2: 1
                    AIGP 210
                    Task: BGP_69.10.255.70.103+179
                    Announcement bits (1): 0-green-l2vpn
                    AS path: I
                    Communities: target:11111:1 Layer2-info: encaps:VPLS,
                    control flags:, mtu: 0
                    Label-base: 800008, range: 8
                    Localpref: 100
                    Router ID: 10.255.70.103
                    Primary Routing Table bgp.l2vpn.0

10.255.71.52:1:1:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:31:40    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8, status-vector: 0x9F

10.255.71.52:1:5:1/96 (1 entry, 1 announced)
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 10.255.71.52
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:31:40    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    99
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512

inet.0: 45 destinations, 47 routes (44 active, 0 holddown, 1 hidden)
1.1.1.3/32 (1 entry, 1 announced)
        *IS-IS  Preference: 18
                Level: 2
                Next hop type: Router, Next hop index: 580
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Address: 0x9db6ed0
                    Next-hop reference count: 8
                    Next hop: 10.1.1.6 via lt-1/0/10.5, selected
                    Session Id: 0x18a
                    State: <Active Int>
                    Local AS:     2
                    Age: 1:32       Metric: 10
                    Validation State: unverified
                    ORR Generation-ID: 1
                    Task: IS-IS
                    Announcement bits (3): 0-KRT 5-Resolve tree 4 6-Resolve_IGP_FRR
task
                    AS path: I

inet.0: 61 destinations, 77 routes (61 active, 1 holddown, 0 hidden)
1.1.1.1/32 (2 entries, 1 announced)
        *OSPF   Preference: 10
                Next hop type: Router, Next hop index: 673
                Address: 0xc008830
                Next-hop reference count: 3
                Next hop: 10.1.1.1 via ge-0/0/2.0, selected
                Session Id: 0x1b7
                State: <Active Int>
                Local AS:     1
                Age: 3:06:59    Metric: 100
                Validation State: unverified
                ORR Generation-ID: 1
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 1-KRT 9-Resolve tree 2
                AS path: I
```

show route detail (with BGP Multipath)

```
user@host> show route detail

10.1.1.8/30 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 262142
                Address: 0x901a010
                Next-hop reference count: 2
                Source: 10.1.1.2
                Next hop: 10.1.1.2 via ge-0/3/0.1, selected
                Next hop: 10.1.1.6 via ge-0/3/0.5
                State: <Active Ext>
                Local AS:     1 Peer AS:     2
                Age: 5:04:43
                Validation State: unverified
                Task: BGP_2.10.1.1.2+59955
                Announcement bits (1): 0-KRT
                AS path: 2 I
                Accepted Multipath
                Localpref: 100
                Router ID: 172.16.1.2
         BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 678
                Address: 0x8f97520
                Next-hop reference count: 9
                Source: 10.1.1.6
                Next hop: 10.1.1.6 via ge-0/3/0.5, selected
                State: <NotBest Ext>
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Inactive reason: Not Best in its group - Active preferred
                    Local AS:     1 Peer AS:     2
                    Age: 5:04:43
                 Validation State: unverified
                    Task: BGP_2.10.1.1.6+58198
                    AS path: 2 I
                    Accepted MultipathContrib
                    Localpref: 100
                    Router ID: 172.16.1.3
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
                Next hop: via vt-0/1/0.1
                Next-hop index: 661
                Label operation: Pop
                Address: 0x9172130
                Next hop: via so-0/0/3.0
                Next-hop index: 654
                Label operation: Swap 299872
                State: **Active Int>
                Local AS:  1001
                Age: 8:20       Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
               FECs bound to route: P2MP root-addr 10.255.72.166, grp 232.1.1.1,
 src 192.168.142.2
```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
                Next-hop reference count: 3
                Next hop: 1.3.8.2 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 1.3.4.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    10
                Age: 54:05      Metric: 1
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Validation State: unverified
                    Task: LDP
                    Announcement bits (1): 0-KRT
                    AS path: I
                    FECs bound to route: P2MP root-addr 172.16.1.1, grp: 232.1.1.1,
src: 192.168.219.11
                    Primary Upstream : 172.16.1.3:0--172.16.1.2:0
                       RPF Nexthops  :
                            ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
                            ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
                    Backup Upstream : 172.16.1.3:0--172.16.1.6:0
                       RPF Nexthops :
                            ge-1/2/20.0, 1.2.96.1, Label: 301584, weight: 0xfffe
                            ge-1/2/19.0, 1.3.6.1, Label: 301584, weight: 0xfffe
```

Copyright © 2017, Juniper Networks, Inc.

show route exact

List of Syntax
Syntax on page 538
Syntax (EX Series Switches) on page 538

Syntax
show route exact *destination-prefix*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route exact *destination-prefix*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display only the routes that exactly match the specified address or range of addresses.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

destination-prefix—Address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route exact on page 538
show route exact detail on page 539
show route exact extensive on page 539
show route exact terse on page 539

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route exact

```
user@host> show route exact 207.17.136.0/24

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
207.17.136.0/24    *[Static/5] 2d 03:30:22
                     > to 192.168.71.254 via fxp0.0
```


Copyright © 2017, Juniper Networks, Inc.

show route exact detail

```
user@host> show route exact 207.17.136.0/24 detail

inet.0: 24 destinations, 25 routes (23 active, 0 holddown, 1 hidden)
Restart Complete
207.17.136.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2d 3:30:26
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I
```

show route exact extensive

```
user@host> show route exact 207.17.136.0/24 extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 1:25:18
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I
```

show route exact terse

```
user@host> show route exact 207.17.136.0/24 terse

inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 207.17.136.0/24    S   5                      >192.168.71.254
```

Copyright © 2017, Juniper Networks, Inc.

show route export

List of Syntax
Syntax on page 540
Syntax (EX Series Switches) on page 540

Syntax
show route export
<brief | detail>
<instance <*instance-name*> | *routing-table-name*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route export
<brief | detail>
<instance <*instance-name*> | *routing-table-name*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display policy-based route export information. Policy-based export simplifies the process of exchanging route information between routing instances.

Options
none—(Same as **brief**.) Display standard information about policy-based export for all instances and routing tables on all systems.

brief | detail—(Optional) Display the specified level of output.

instance <*instance-name*>—(Optional) Display a particular routing instance for which policy-based export is currently enabled.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—(Optional) Display information about policy-based export for all routing tables whose name begins with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route export inet** command).

Required Privilege Level
view

List of Sample Output
show route export on page 541
show route export detail on page 541
show route export instance detail on page 542

Output Fields
Table 21 on page 322 lists the output fields for the **show route export** command. Output fields are listed in the approximate order in which they appear.

Copyright © 2017, Juniper Networks, Inc.

Table 42: show route export Output Fields

Field Name	Field Description	Level of Output
Table or *table-name*	Name of the routing tables that either import or export routes.	All levels
Routes	Number of routes exported from this table into other tables. If a particular route is exported to different tables, the counter will only increment by one.	**brief** none
Export	Whether the table is currently exporting routes to other tables: **Y** or **N** (Yes or No).	**brief** none
Import	Tables currently importing routes from the originator table. (Not displayed for tables that are not exporting any routes.)	**detail**
Flags	(**instance** keyword only) Flags for this feature on this instance: • **config auto-policy**—The policy was deduced from the configured IGP export policies. • **cleanup**—Configuration information for this instance is no longer valid. • **config**—The instance was explicitly configured.	**detail**
Options	(**instance** keyword only) Configured option displays the type of routing tables the feature handles: • **unicast**—Indicates *instance*.inet.0. • **multicast**—Indicates *instance*.inet.2. • **unicast multicast**—Indicates *instance*.inet.0 and *instance*.inet.2.	**detail**
Import policy	(**instance** keyword only) Policy that **route export** uses to construct the import-export matrix. Not displayed if the instance type is **vrf**.	**detail**
Instance	(**instance** keyword only) Name of the routing instance.	**detail**
Type	(**instance** keyword only) Type of routing instance: **forwarding**, **non-forwarding**, or **vrf**.	**detail**

Sample Output

show route export

```
user@host> show route export
Table                           Export          Routes
inet.0                          N                    0
black.inet.0                    Y                    3
red.inet.0                      Y                    4
```

show route export detail

```
user@host> show route export detail
inet.0                                          Routes:         0
black.inet.0                                    Routes:         3
  Import: [ inet.0 ]
red.inet.0                                      Routes:         4
  Import: [ inet.0 ]
```

Copyright © 2017, Juniper Networks, Inc.

show route export instance detail

```
user@host> show route export instance detail
Instance: master                          Type: forwarding
  Flags: <config auto-policy> Options: <unicast multicast>
  Import policy: [ (ospf-master-from-red || isis-master-from-black) ]
Instance: black                           Type: non-forwarding
Instance: red                             Type: non-forwarding
```


Copyright © 2017, Juniper Networks, Inc.

show route extensive

List of Syntax
Syntax on page 543
Syntax (EX Series Switches) on page 543

Syntax
show route extensive
<*destination-prefix*>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route extensive
<*destination-prefix*>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display extensive information about the active entries in the routing tables.

Options
none—Display all active entries in the routing table.

destination-prefix—(Optional) Display active entries for the specified address or range of addresses.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route extensive on page 550
show route extensive (Access Route) on page 557
show route extensive (BGP PIC Edge) on page 557
show route extensive (FRR and LFA) on page 558
show route extensive (IS-IS) on page 559
show route extensive (Route Reflector) on page 559
show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs) on page 559
show route label detail (Multipoint LDP with Multicast-Only Fast Reroute) on page 560

Output Fields
Table 22 on page 324 describes the output fields for the **show route extensive** command. Output fields are listed in the approximate order in which they appear.

Table 43: show route extensive Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number **destinations**	Number of destinations for which there are routes in the routing table.

Table 43: show route extensive Output Fields *(continued)*

Field Name	Field Description
***number* routes**	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active). • **holddown** (routes that are in the pending state before being declared inactive). • **hidden** (routes that are not used because of a routing policy).

***route-destination* (entry, announced)**	Route destination (for example: 10.0.0.1/24). The **entry** value is the number of route for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.

TSI	Protocol header information.
label stacking	(Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of two or more exits this router with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.

Copyright © 2017, Juniper Networks, Inc.

Table 43: show route extensive Output Fields *(continued)*

Field Name	Field Description
Level	(IS-IS only). In IS-IS, a single autonomous system (AS) can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see the Output Field table in the show route detail command.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Offset	Whether the metric has been increased or decreased by an offset value.
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to recursively derive a forwarding next hop.

Copyright © 2017, Juniper Networks, Inc.

Table 43: show route extensive Output Fields *(continued)*

Field Name	Field Description
label-operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Indirect next hops	When present, a list of nodes that are used to resolve the path to the next-hop destination, in the order that they are resolved. When BGP PIC Edge is enabled, the output lines that contain **Indirect next hop: weight** follow next hops that the software can use to repair paths where a link failure occurs. The next-hop weight has one of the following values: • 0x1 indicates active next hops. • 0x4000 indicates passive next hops.
State	State of the route (a route can be in more than one state). See the Output Field table in the show route detail command.
Session ID	The BFD session ID number that represents the protection using MPLS fast reroute (FRR) and loop-free alternate (LFA).
Weight	Weight for the backup path. If the weight of an indirect next hop is larger than zero, the weight value is shown. For sample output, see show route table.

Copyright © 2017, Juniper Networks, Inc.

Table 43: show route extensive Output Fields *(continued)*

Field Name	Field Description
Inactive reason	If the route is inactive, the reason for its current state is indicated. Typical reasons include: • **Active preferred**—Currently active route was selected over this route. • **Always compare MED**—Path with a lower multiple exit discriminator (MED) is available. • **AS path**—Shorter AS path is available. • **Cisco Non-deterministic MED selection**—Cisco nondeterministic MED is enabled and a path with a lower MED is available. • **Cluster list length**—Path with a shorter cluster list length is available. • **Forwarding use only**—Path is only available for forwarding purposes. • **IGP metric**—Path through the next hop with a lower IGP metric is available. • **IGP metric type**—Path with a lower OSPF link-state advertisement type is available. • **Interior > Exterior > Exterior via Interior**—Direct, static, IGP, or EBGP path is available. • **Local preference**—Path with a higher local preference value is available. • **Next hop address**—Path with a lower metric next hop is available. • **No difference**—Path from a neighbor with a lower IP address is available. • **Not Best in its group**—Occurs when multiple peers of the same external AS advertise the same prefix and are grouped together in the selection process. When this reason is displayed, an additional reason is provided (typically one of the other reasons listed). • **Number of gateways**—Path with a higher number of next hops is available. • **Origin**—Path with a lower origin code is available. • **OSPF version**—Path does not support the indicated OSPF version. • **RIB preference**—Route from a higher-numbered routing table is available. • **Route destinguisher**—64-bit prefix added to IP subnets to make them unique. • **Route metric or MED comparison**—Route with a lower metric or MED is available. • **Route preference**—Route with a lower preference value is available. • **Router ID**—Path through a neighbor with a lower ID is available. • **Unusable path**—Path is not usable because of one of the following conditions: the route is damped, the route is rejected by an import policy, or the route is unresolved. • **Update source**—Last tiebreaker is the lowest IP address value.
Local AS	Autonomous system (AS) number of the local routing device.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances. For sample output, see show route table.

Copyright © 2017, Juniper Networks, Inc.

Table 43: show route extensive Output Fields *(continued)*

Field Name	Field Description
Task	Name of the protocol that has added the route.
Announcement bits	List of protocols that are consumers of the route. Using the following output as an example, **Announcement bits (3): 0-KRT 5-Resolve tree 2 8-BGP RT Background** there are (3) announcement bits to reflect the three clients (protocols) that have state for this route: Kernel (0-KRT), 5 (resolution tree process 2), and 8 (BGP). The notation ***n*-Resolve inet** indicates that the route is used for route resolution for next hops found in the routing table. ***n*** is an index used by Juniper Networks customer support only.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the local AS number associated with the AS path if more than one AS number is configured on the routing device, or if AS path prepending is configured. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
AS path: I <Originator>	(For route reflected output only) Originator ID attribute set by the route reflector.

Copyright © 2017, Juniper Networks, Inc.

Table 43: show route extensive Output Fields *(continued)*

Field Name	Field Description
route status	Indicates the status of a BGP route: • **Accepted**—The specified BGP route is imported by the default BGP policy. • **Import**—The route is imported into a Layer 3 VPN routing instance. • **Import-Protect**—A remote instance egress that is protected. • **Multipath**—A BGP multipath active route. • **MultipathContrib**—The route is not active but contributes to the BGP multipath. • **Protect**—An egress route that is protected. • **Stale**—A route that is marked stale due to graceful restart.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.
Originator ID	(For route reflected output only) Address of router that originally sent the route to the route reflector.
Prefixes bound to route	Forwarding Equivalent Class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or Site Down.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.

Copyright © 2017, Juniper Networks, Inc.

Table 43: show route extensive Output Fields *(continued)*

Field Name	Field Description
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.
Originating RIB	Name of the routing table whose active route was used to determine the forwarding next-hop entry in the resolution database. For example, in the case of inet.0 resolving through inet.0 and inet.3, this field indicates which routing table, inet.0 or inet.3, provided the best path for a particular prefix.
Node path count	Number of nodes in the path.
Forwarding nexthops	Number of forwarding next hops. The forwarding next hop is the network layer address of the directly reachable neighboring system (if applicable) and the interface used to reach it.

Sample Output

show route extensive

```
user@host> show route extensive
inet.0: 22 destinations, 23 routes (21 active, 0 holddown, 1 hidden)
203.0.113.10/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.10/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 29
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    64496
                Age: 1:34:06
                Task: RT
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

203.0.113.30/30 (2 entries, 1 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 2
                Next hop: via so-0/3/0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:32:40
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
```


Copyright © 2017, Juniper Networks, Inc.

```
                State: <Int>
                Inactive reason: Route Preference
                Local AS:    64496
                Age: 1:32:40    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

203.0.113.103/32 (1 entry, 1 announced)
        *Local  Preference: 0
                Next hop type: Local
                Next-hop reference count: 7
                Interface: so-0/3/0.0
                State: <Active NoReadvrt Int>
                Local AS:    644969
                Age: 1:32:43
                Task: IF
                Announcement bits (1): 3-Resolve tree 2
                AS path: I

...

203.0.113.203/30 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.203/30 -> {203.0.113.216}
        *OSPF   Preference: 10
                Next-hop reference count: 9
                Next hop: via so-0/3/0.0
                Next hop: 203.0.113.216 via ge-3/1/0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:32:19    Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.2/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.2/32 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

...

198.51.100.22/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 198.51.100.22/32 -> {}
        *IGMP   Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
```

Copyright © 2017, Juniper Networks, Inc.

```
                Task: IGMP
                Announcement bits (2): 0-KRT 3-Resolve tree 2
                AS path: I

inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

203.0.113.103/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 100096
                State: <Active Int>
                Local AS:    64496
                Age: 1:28:12    Metric: 2
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

203.0.113.238/32 (1 entry, 1 announced)
        State: <FlashAll>
        *RSVP   Preference: 7
                Next-hop reference count: 6
                Next hop: via so-0/3/0.0 weight 0x1, selected
                Label-switched-path green-r1-r2
                State: <Active Int>
                Local AS:    64496
                Age: 1:28:12    Metric: 1
                Task: RSVP
                Announcement bits (2): 1-Resolve tree 1 2-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

...

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

47.0005.80ff.f800.0000.0108.0001.0102.5507.1052/152 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

0   (1 entry, 1 announced)
TSI:
KRT in-kernel 0      /36 -> {}
        *MPLS   Preference: 0
                Next hop type: Receive
                Next-hop reference count: 6
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:08    Metric: 1
```


Copyright © 2017, Juniper Networks, Inc.

```
                Task: MPLS
                Announcement bits (1): 0-KRT
                AS path: I

...

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
  299840 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 299840 /52 -> {indirect(1048575)}
          *RSVP   Preference: 7/2
                  Next hop type: Flood
                  Address: 0x9174a30
                  Next-hop reference count: 4
                  Next hop type: Router, Next hop index: 798
                  Address: 0x9174c28
                  Next-hop reference count: 2
                  Next hop: 198.51.100.2 via lt-1/2/0.9 weight 0x1
                  Label-switched-path R2-to-R4-2p2mp
                  Label operation: Pop
                  Next hop type: Router, Next hop index: 1048574
                  Address: 0x92544f0
                  Next-hop reference count: 2
                  Next hop: 198.51.100.2 via lt-1/2/0.7 weight 0x1
                  Label-switched-path R2-to-R200-p2mp
                  Label operation: Pop
                  Next hop: 198.51.100.2 via lt-1/2/0.5 weight 0x8001
                  Label operation: Pop
                  State: <Active Int>
                  Age: 1:29       Metric: 1
                  Task: RSVP
                  Announcement bits (1): 0-KRT
                  AS path: I...

800010 (1 entry, 1 announced)

TSI:
KRT in-kernel 800010 /36 -> {vt-3/2/0.32769}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: via vt-3/2/0.32769, selected
                Label operation: Pop
                State: <Active Int>
                Age: 1:31:53
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

vt-3/2/0.32769 (1 entry, 1 announced)
TSI:
KRT in-kernel vt-3/2/0.32769.0      /16 -> {indirect(1048574)}
        *VPLS   Preference: 7
                Next-hop reference count: 2
                Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1, selected
                Label-switched-path green-r1-r3
                Label operation: Push 800012, Push 100096(top)
                Protocol next hop: 203.0.113.103
                Push 800012
                Indirect next hop: 87272e4 1048574
                State: <Active Int>
                Age: 1:31:53    Metric2: 2
```

Copyright © 2017, Juniper Networks, Inc.

```
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Indirect next hops: 1
                        Protocol next hop: 203.0.113.103 Metric: 2
                        Push 800012
                        Indirect next hop: 87272e4 1048574
                        Indirect path forwarding next hops: 1
                               Next hop: 203.0.113.216 via ge-3/1/0.0 weight 0x1

                        203.0.113.103/32 Originating RIB: inet.3
                          Metric: 2                     Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 203.0.113.216 via ge-3/1/0.0

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

2001:db8::10:255:71:52/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:07
                Task: IF
                AS path: I

ff02::2/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::2/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::d/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::d/128 -> {}
        *PIM    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:08
```


Copyright © 2017, Juniper Networks, Inc.

```
                Task: PIM Recv6
                Announcement bits (1): 0-KRT
                AS path: I

ff02::16/128 (1 entry, 1 announced)
TSI:
KRT in-kernel ff02::16/128 -> {}
        *MLD    Preference: 0
                Next-hop reference count: 18
                State: <Active NoReadvrt Int>
                Local AS:    64496
                Age: 1:34:06
                Task: MLD
                Announcement bits (1): 0-KRT
                AS path: I

private.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

fe80::280:42ff:fe10:f179/128 (1 entry, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.16385, selected
                State: <Active NoReadvrt Int>
                Age: 1:34:07
                Task: IF
                AS path: I

green.l2vpn.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)

203.0.113.103:1:3:1/96 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 203.0.113.103:1
                Next-hop reference count: 7
                Source: 203.0.113.103
                Protocol next hop: 203.0.113.103
                Indirect next hop: 2 no-forward
                State: <Secondary Active Int Ext>
                Local AS:    64496 Peer AS:    64496
                Age: 1:28:12    Metric2: 1
                Task: BGP_69.203.0.113.103+179
                Announcement bits (1): 0-green-l2vpn
                AS path: I
                Communities: target:11111:1 Layer2-info: encaps:VPLS,
                control flags:, mtu: 0
                Label-base: 800008, range: 8
                Localpref: 100
                Router ID: 203.0.113.103
                Primary Routing Table bgp.l2vpn.0

203.0.113.152:1:1:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699540
        *L2VPN  Preference: 170/-1
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
```

Copyright © 2017, Juniper Networks, Inc.

```
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:Site-Down,
                mtu: 0
                Label-base: 800016, range: 8, status-vector: 0x9F

203.0.113.152:1:5:1/96 (1 entry, 1 announced)
TSI:
Page 0 idx 0 Type 1 val 8699528
        *L2VPN  Preference: 170/-101
                Next-hop reference count: 5
                Protocol next hop: 203.0.113.152
                Indirect next hop: 0 -
                State: <Active Int Ext>
                Age: 1:34:03    Metric2: 1
                Task: green-l2vpn
                Announcement bits (1): 1-BGP.0.0.0.0+179
                AS path: I
                Communities: Layer2-info: encaps:VPLS, control flags:, mtu: 0
                Label-base: 800008, range: 8, status-vector: 0x9F

...

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

TSI:

203.0.113.163:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 203.0.113.163 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    64499
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512


203.0.113.55/24 (1 entry, 1 announced)
TSI:
KRT queued (pending) add
  198.51.100.0/24 -> {Push 300112}
        *BGP    Preference: 170/-101
                Next hop type: Router
                Address: 0x925c208
                Next-hop reference count: 2
                Source: 203.0.113.9
                Next hop: 203.0.113.9 via ge-1/2/0.15, selected
                Label operation: Push 300112
                Label TTL action: prop-ttl
                State: <Active Ext>
                Local AS:  64509 Peer AS: 65539
                Age: 1w0d 23:06:56
                AIGP: 25
                Task: BGP_65539.203.0.113.9+56732
                Announcement bits (1): 0-KRT
                AS path: 65539 64508 I
                Accepted
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Route Label: 300112
                    Localpref: 100
                    Router ID: 213.0.113.99
```

show route extensive (Access Route)

```
user@host> show route 203.0.113.102 extensive
inet.0: 39256 destinations, 39258 routes (39255 active, 0 holddown, 1 hidden)
203.0.113.102/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 203.0.113.102/32 -> {192.0.2.2}
OSPF area : 0.0.0.0, LSA ID : 203.0.113.102, LSA type : Extern
        *Access Preference: 13
                Next-hop reference count: 78472
                Next hop: 192.0.2.2 via fe-0/0/0.0, selected
                State: <Active Int>
       Age: 12
                Task: RPD Unix Domain Server./var/run/rpd_serv.local
                Announcement bits (2): 0-KRT 1-OSPFv2
                AS path: I


user@host> show route 2001:db8:4641:1::/48 extensive

inet6.0: 75 destinations, 81 routes (75 active, 0 holddown, 0 hidden)
  2001:db8:4641:1::/48 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 2001:db8:4641:1::/48 -> {#0 0.13.1.0.0.1}
          *Access Preference: 13
                  Next hop type: Router, Next hop index: 74548
                  Address: 0x1638c1d8
                  Next-hop reference count: 6
                  Next hop: #0 0.13.1.0.0.1 via demux0.1073753267, selected
                  Session Id: 0x0
                  State: <Active Int>
                  Age: 4:17
                  Validation State: unverified
                  Task: RPD Unix Domain Server./var/run/rpd_serv.local
                  Announcement bits (2): 0-KRT 4-Resolve tree 2
                  AS path: I
2001:db8:4641:1::/128 (1 entry, 1 announced)
  TSI:
  KRT in-kernel 2001:db8:4641:1::/128 -> {#0 0.13.1.0.0.1}
          *Access-internal Preference: 12
                  Next hop type: Router, Next hop index: 74548
                  Address: 0x1638c1d8
                  Next-hop reference count: 6
                  Next hop: #0 0.13.1.0.0.1 via demux0.1073753267, selected
                  Session Id: 0x0
                  State: <Active Int>
                  Age: 4:17
                  Validation State: unverified
                  Task: RPD Unix Domain Server./var/run/rpd_serv.local
                  Announcement bits (2): 0-KRT 4-Resolve tree 2
                  AS path: I
```

show route extensive (BGP PIC Edge)

```
user@host> show route 198.51.100.6 extensive
ed.inet.0: 6 destinations, 9 routes (6 active, 0 holddown, 0 hidden)
  198.51.100.6/32 (3 entries, 2 announced)
```

```
          State: <CalcForwarding>
  TSI:
  KRT in-kernel 198.51.100.6/32 -> {indirect(1048574), indirect(1048577)}
  Page 0 idx 0 Type 1 val 9219e30
      Nexthop: Self
      AS path: [2] 3 I
      Communities: target:2:1
  Path 198.51.100.6 from 198.51.100.4 Vector len 4.  Val: 0
..
          #Multipath Preference: 255
                  Next hop type: Indirect
                  Address: 0x93f4010
                  Next-hop reference count: 2
..
                  Protocol next hop: 198.51.1001.4
                  Push 299824
                  Indirect next hop: 944c000 1048574 INH Session ID: 0x3
                  Indirect next hop: weight 0x1
                  Protocol next hop: 198.51.100.5
                  Push 299824
                  Indirect next hop: 944c1d8 1048577 INH Session ID: 0x4
                  Indirect next hop: weight 0x4000
                  State: <ForwardingOnly Int Ext)>
                  Inactive reason: Forwarding use only
                  Age: 25         Metric2: 15
                  Validation State: unverified
                  Task: RT
                  Announcement bits (1): 0-KRT
                  AS path: 3 I
                  Communities: target:2:1
```

show route extensive (FRR and LFA)

```
user@host> show route 203.0.113.20 extensive
inet.0: 46 destinations, 49 routes (45 active, 0 holddown, 1 hidden)
203.0.113.20/24 (2 entries, 1 announced)
        State: FlashAll
TSI:
KRT in-kernel 203.0.113.20/24 -> {Push 299776, Push 299792}
        *RSVP   Preference: 7/1
                Next hop type: Router, Next hop index: 1048574
                Address: 0xbbbc010
                Next-hop reference count: 5
                Next hop: 203.0.113.112 via ge-2/1/8.0 weight 0x1, selected
                Label-switched-path europa-d-to-europa-e
                Label operation: Push 299776
                Label TTL action: prop-ttl
                Session Id: 0x201
                Next hop: 203.0.113.122 via ge-2/1/4.0 weight 0x4001
                Label-switched-path europa-d-to-europa-e
                Label operation: Push 299792
                Label TTL action: prop-ttl
                Session Id: 0x202
                State: Active Int
                Local AS:   64500
                Age: 5:31  Metric: 2
                Task: RSVP
                Announcement bits (1): 0-KRT
                AS path: I
         OSPF   Preference: 10
                Next hop type: Router, Next hop index: 615
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Address: 0xb9d78c4
                    Next-hop reference count: 7
                    Next hop: 203.0.113.112 via ge-2/1/8.0, selected
                    Session Id: 0x201
                    State: Int
                    Inactive reason: Route Preference
                    Local AS:   64500
                    Age: 5:35  Metric: 3
                    Area: 0.0.0.0
                    Task: OSPF
                    AS path: I
```

show route extensive (IS-IS)

```
user@host> show route extensive
IS-IS Preference: 15
                Level: 1
                Next hop type: Router, Next hop index: 1048577
                Address: 0xXXXXXXXXXX
                Next-hop reference count: YY
                Next hop: 203.0.113.22 via ae1.0 balance 43%, selected
                Session Id: 0x141
                Next hop: 203.0.113.22 via ae0.0 balance 57%
```

show route extensive (Route Reflector)

```
user@host> show route extensive
203.0.113.0/8 (1 entry, 1 announced)

TSI:
KRT in-kernel 203.0.113.0/8 -> {indirect(40)}
        *BGP    Preference: 170/-101
                Source: 192.168.4.214
                Protocol next hop: 198.51.100.192 Indirect next hop: 84ac908 40
                State: <Active Int Ext>
                Local AS: 65548 Peer AS: 65548
                Age: 3:09       Metric: 0       Metric2: 0
                Task: BGP_65548.192.168.4.214+1033
                Announcement bits (2): 0-KRT 4-Resolve inet.0
                AS path: 65544 64507 I <Originator>
                Cluster list:  198.51.100.1
                Originator ID: 203.0.113.88
                Communities: 7777:7777
                Localpref: 100
                Router ID: 203.0.113.4
                Indirect next hops: 1
                        Protocol next hop: 203.0.113.192 Metric: 0
                        Indirect next hop: 84ac908 40
                        Indirect path forwarding next hops: 0
                                Next hop type: Discard
```

show route label detail (Multipoint LDP Inband Signaling for Point-to-Multipoint LSPs)

```
user@host> show route label 299872 detail
mpls.0: 13 destinations, 13 routes (13 active, 0 holddown, 0 hidden)
299872 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Next-hop reference count: 3
                Address: 0x9097d90
```

Copyright © 2017, Juniper Networks, Inc.

```
                    Next hop: via vt-0/1/0.1
                    Next-hop index: 661
                    Label operation: Pop
                    Address: 0x9172130
                    Next hop: via so-0/0/3.0
                    Next-hop index: 654
                    Label operation: Swap 299872
                    State: **Active Int>
                    Local AS:  64511
                    Age: 8:20       Metric: 1
                    Task: LDP
                    Announcement bits (1): 0-KRT
                    AS path: I
                   FECs bound to route: P2MP root-addr 203.0.113.166, grp 203.0.1.1,
 src 192.168.142.2
```

show route label detail (Multipoint LDP with Multicast-Only Fast Reroute)

```
user@host> show route label 301568 detail

mpls.0: 18 destinations, 18 routes (18 active, 0 holddown, 0 hidden)
301568 (1 entry, 1 announced)
        *LDP    Preference: 9
                Next hop type: Flood
                Address: 0x2735208
                Next-hop reference count: 3
                Next hop type: Router, Next hop index: 1397
                Address: 0x2735d2c
                Next-hop reference count: 3
                Next hop: 203.0.113.82 via ge-1/2/22.0
                Label operation: Pop
                Load balance label:  None;
                Next hop type: Router, Next hop index: 1395
                Address: 0x2736290
                Next-hop reference count: 3
                Next hop: 203.0.113.2 via ge-1/2/18.0
                Label operation: Pop
                Load balance label:  None;
                State: <Active Int AckRequest MulticastRPF>
                Local AS:    64500
                Age: 54:05      Metric: 1
                Validation State: unverified
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                FECs bound to route: P2MP root-addr 198.51.100.1, grp: 232.1.1.1,
 src: 192.168.219.11
                Primary Upstream : 198.51.100.3:0--198.51.100.2:0
                   RPF Nexthops  :
                        ge-1/2/15.0, 1.2.94.1, Label: 301568, weight: 0x1
                        ge-1/2/14.0, 1.2.3.1, Label: 301568, weight: 0x1
                Backup Upstream : 198.51.100.3:0--198.51.100.6:0
                   RPF Nexthops :
                       ge-1/2/20.0, 198.51.100.96, Label: 301584, weight: 0xfffe

                       ge-1/2/19.0, 198.51.100.36, Label: 301584, weight: 0xfffe
```


Copyright © 2017, Juniper Networks, Inc.

show route forwarding-table

List of Syntax

Syntax on page 561
Syntax (MX Series Routers) on page 561
Syntax (TX Matrix and TX Matrix Plus Routers) on page 561

Syntax

show route forwarding-table
<detail | extensive | summary>
<all>
<ccc *interface-name*>
<destination *destination-prefix*>
<family *family* | matching *matching*>
<interface-name *interface-name*>
<label *name*>
<matching *matching*>
<multicast>
<table (default | *logical-system-name/routing-instance-name* | *routing-instance-name*)>
<vlan (all | *vlan-name*)>
<vpn *vpn*>

Syntax (MX Series Routers)

show route forwarding-table
<detail | extensive | summary>
<all>
<bridge-domain (all | *domain-name*)>
<ccc *interface-name*>
<destination *destination-prefix*>
<family *family* | matching *matching*>
<interface-name *interface-name*>
<label *name*>
<learning-vlan-id *learning-vlan-id*>
<matching *matching*>
<multicast>
<table (default | *logical-system-name/routing-instance-name* | *routing-instance-name*)>
<vlan (all | *vlan-name*)>
<vpn *vpn*>

Syntax (TX Matrix and TX Matrix Plus Routers)

show route forwarding-table
<detail | extensive | summary>
<all>
<ccc *interface-name*>
<destination *destination-prefix*>
<family *family* | matching *matching*>
<interface-name *interface-name*>
<matching *matching*>
<label *name*>
<lcc *number*>
<multicast>
<table *routing-instance-name*>
<vpn *vpn*>

Release Information

Command introduced before Junos OS Release 7.4.
Option **bridge-domain** introduced in Junos OS Release 7.5

Copyright © 2017, Juniper Networks, Inc.

Option **learning-vlan-id** introduced in Junos OS Release 8.4
Options **all** and **vlan** introduced in Junos OS Release 9.6.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description Display the Routing Engine's forwarding table, including the network-layer prefixes and their next hops. This command is used to help verify that the routing protocol process has relayed the correction information to the forwarding table. The Routing Engine constructs and maintains one or more routing tables. From the routing tables, the Routing Engine derives a table of active routes, called the forwarding table.

NOTE: The Routing Engine copies the forwarding table to the Packet Forwarding Engine, the part of the router that is responsible for forwarding packets. To display the entries in the Packet Forwarding Engine's forwarding table, use the **show pfe route** command.

Options

none—Display the routes in the forwarding tables. By default, the **show route forwarding-table** command does not display information about private, or internal, forwarding tables.

detail | extensive | summary—(Optional) Display the specified level of output.

all—(Optional) Display routing table entries for all forwarding tables, including private, or internal, tables.

bridge-domain (all | *bridge-domain-name*)—(MX Series routers only) (Optional) Display route entries for all bridge domains or the specified bridge domain.

ccc *interface-name*—(Optional) Display route entries for the specified circuit cross-connect interface.

destination *destination-prefix*—(Optional) Destination prefix.

family *family*—(Optional) Display routing table entries for the specified family: **bridge (ccc | destination | detail | extensive | interface-name | label | learning-vlan-id | matching | multicast | summary | table | vlan | vpn)**, **ethernet-switching**, **evpn**, **fibre-channel**, **fmembers**, **inet**, **inet6**, **iso**, **mcsnoop-inet**, **mcsnoop-inet6**, **mpls**, **satellite-inet**, **satellite-inet6**, **satellite-vpls**, **tnp**, **unix**, **vpls**, or **vlan-classification**.

interface-name *interface-name*—(Optional) Display routing table entries for the specified interface.

label *name*—(Optional) Display route entries for the specified label.

lcc *number*—(TX Matrix and TX matrix Plus routers only) (Optional) On a routing matrix composed of a TX Matrix router and T640 routers, display information for the specified T640 router (or line-card chassis) connected to the TX Matrix router. On a routing matrix composed of the TX Matrix Plus router and T1600 or T4000 routers,

Copyright © 2017, Juniper Networks, Inc.

display information for the specified router (line-card chassis) connected to the TX Matrix Plus router.

Replace *number* with the following values depending on the LCC configuration:

- 0 through 3, when T640 routers are connected to a TX Matrix router in a routing matrix.
- 0 through 3, when T1600 routers are connected to a TX Matrix Plus router in a routing matrix.
- 0 through 7, when T1600 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.
- 0, 2, 4, or 6, when T4000 routers are connected to a TX Matrix Plus router with 3D SIBs in a routing matrix.

learning-vlan-id *learning-vlan-id*—(MX Series routers only) (Optional) Display learned information for all VLANs or for the specified VLAN.

matching *matching*—(Optional) Display routing table entries matching the specified prefix or prefix length.

multicast—(Optional) Display routing table entries for multicast routes.

table —(Optional) Display route entries for all the routing tables in the main routing instance or for the specified routing instance. If your device supports logical systems, you can also display route entries for the specified logical system and routing instance. To view the routing instances on your device, use the show route instance command.

vlan (all | *vlan-name*)—(Optional) Display information for all VLANs or for the specified VLAN.

vpn *vpn*—(Optional) Display routing table entries for a specified VPN.

Required Privilege Level view

List of Sample Output

show route forwarding-table on page 568
show route forwarding-table detail on page 569
show route forwarding-table destination extensive (Weights and Balances) on page 569
show route forwarding-table extensive on page 570
show route forwarding-table extensive (RPF) on page 571
show route forwarding-table (dynamic list next hop) on page 572
show route forwarding-table family mpls on page 573
show route forwarding-table family mpls ccc ge-0/0/1.1004 on page 573
show route forwarding-table family vpls on page 573
show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled) on page 573
show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics) on page 574
show route forwarding-table family vpls extensive on page 574

Copyright © 2017, Juniper Networks, Inc.

show route forwarding-table table default on page 575
show route forwarding-table table
logical-system-name/routing-instance-name on page 576
show route forwarding-table vpn on page 577

Output Fields Table 23 on page 345 lists the output fields for the **show route forwarding-table** command. Output fields are listed in the approximate order in which they appear. Field names might be abbreviated (as shown in parentheses) when no level of output is specified, or when the **detail** keyword is used instead of the **extensive** keyword.

Table 44: show route forwarding-table Output Fields

Field Name	Field Description	Level of Output
Logical system	Name of the logical system. This field is displayed if you specify the **table** ***logical-system-name/routing-instance-name*** option on a device that is configured for and supports logical systems.	All levels
Routing table	Name of the routing table (for example, inet, inet6, mpls).	All levels

Copyright © 2017, Juniper Networks, Inc.

Table 44: show route forwarding-table Output Fields *(continued)*

Field Name	Field Description	Level of Output
Enabled protocols	The features and protocols that have been enabled for a given routing table. This field can contain the following values: • BUM hashing—BUM hashing is enabled. • MAC Stats—Mac Statistics is enabled. • Bridging—Routing instance is a normal layer 2 bridge. • No VLAN—No VLANs are associated with the bridge domain. • All VLANs—The **vlan-id all** statement has been enabled for this bridge domain. • Single VLAN—Single VLAN ID is associated with the bridge domain. • MAC action drop—New MACs will be dropped when the MAC address limit is reached. • Dual VLAN—Dual VLAN tags are associated with the bridge domain • No local switching—No local switching is enabled for this routing instance.. • Learning disabled—Layer 2 learning is disabled for this routing instance. • MAC limit reached—The maximum number of MAC addresses that was configured for this routing instance has been reached. • VPLS—The VPLS protocol is enabled. • No IRB l2-copy—The no-irb-layer-2-copy feature is enabled for this routing instance. • ACKed by all peers—All peers have acknowledged this routing instance. • BUM Pruning—BUM pruning is enabled on the VPLS instance. • Def BD VXLAN—VXLAN is enabled for the default bridge domain. • EVPN—EVPN protocol is enabled for this routing instance. • Def BD OVSDB—Open vSwitch Database (OVSDB) is enabled on the default bridge domain. • Def BD Ingress replication—VXLAN ingress node replication is enabled on the default bridge domain. • L2 backhaul—Layer 2 backhaul is enabled. • FRR optimize—Fast reroute optimization • MAC pinning—MAC pinning is enabled for this bridge domain. • MAC Aging Timer—The MAC table aging time is set per routing instance. • EVPN VXLAN—This routing instance supports EVPN with VXLAN encapsulation. • PBBN—This routing instance is configured as a provider backbone bridged network. • PBN—This routing instance is configured as a provider bridge network. • ETREE—The ETREE protocol is enabled on this EVPN routing instance. • ARP/NDP suppression—EVPN ARP NDP suppression is enabled in this routing instance. • Def BD EVPN VXLAN—EVPN VXLAN is enabled for the default bridge domain. • MPLS control word—Control word is enabled for this MPLS routing instance.	All levels
Address family	Address family (for example, **IP**, **IPv6**, **ISO**, **MPLS**, and **VPLS**).	All levels
Destination	Destination of the route.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 44: show route forwarding-table Output Fields *(continued)*

Field Name	Field Description	Level of Output
Route Type (Type)	How the route was placed into the forwarding table. When the **detail** keyword is used, the route type might be abbreviated (as shown in parentheses): • **cloned** (**clon**)—(TCP or multicast only) Cloned route. • **destination** (**dest**)—Remote addresses directly reachable through an interface. • **destination down** (**iddn**)—Destination route for which the interface is unreachable. • **interface cloned** (**ifcl**)—Cloned route for which the interface is unreachable. • **route down** (**ifdn**)—Interface route for which the interface is unreachable. • **ignore** (**ignr**)—Ignore this route. • **interface** (**intf**)—Installed as a result of configuring an interface. • **permanent** (**perm**)—Routes installed by the kernel when the routing table is initialized. • **user**—Routes installed by the routing protocol process or as a result of the configuration.	All levels
Route Reference (RtRef)	Number of routes to reference.	**detail extensive**
Flags	Route type flags: • **none**—No flags are enabled. • **accounting**—Route has accounting enabled. • **cached**—Cache route. • **incoming-iface** ***interface-number***—Check against incoming interface. • **prefix load balance**—Load balancing is enabled for this prefix. • **rt nh decoupled**—Route has been decoupled from the next hop to the destination. • **sent to PFE**—Route has been sent to the Packet Forwarding Engine. • **static**—Static route.	**extensive**
Next hop	IP address of the next hop to the destination.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 44: show route forwarding-table Output Fields *(continued)*

Field Name	Field Description	Level of Output
Next hop Type (Type)	Next-hop type. When the **detail** keyword is used, the next-hop type might be abbreviated (as indicated in parentheses): • **broadcast (bcst)**—Broadcast. • **deny**—Deny. • **discard (dscd)** —Discard. • **hold**—Next hop is waiting to be resolved into a unicast or multicast type. • **indexed (idxd)**—Indexed next hop. • **indirect (indr)**—Indirect next hop. • **local (locl)**—Local address on an interface. • **routed multicast (mcrt)**—Regular multicast next hop. • **multicast (mcst)**—Wire multicast next hop (limited to the LAN). • **multicast discard (mdsc)**—Multicast discard. • **multicast group (mgrp)**—Multicast group member. • **receive (recv)**—Receive. • **reject (rjct)**—Discard. An ICMP unreachable message was sent. • **resolve (rslv)**—Resolving the next hop. • **unicast (ucst)**—Unicast. • **unilist (ulst)**—List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.	**detail extensive**
Index	Software index of the next hop that is used to route the traffic for a given prefix.	**detail extensive** none
Route interface-index	Logical interface index from which the route is learned. For example, for interface routes, this is the logical interface index of the route itself. For static routes, this field is zero. For routes learned through routing protocols, this is the logical interface index from which the route is learned.	**extensive**
Reference (NhRef)	Number of routes that refer to this next hop.	**detail extensive** none
Next-hop interface (Netif)	Interface used to reach the next hop.	**detail extensive** none
Weight	Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible (see the **Balance** field description).	**extensive**
Balance	Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a router is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.	**extensive**
RPF interface	List of interfaces from which the prefix can be accepted. Reverse path forwarding (RPF) information is displayed only when **rpf-check** is configured on the interface.	**extensive**

Copyright © 2017, Juniper Networks, Inc.

Sample Output

show route forwarding-table

```
user@host> show route forwarding-table
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    46     4
0.0.0.0/32         perm     0                    dscd    44     1
172.16.1.0/24         ifdn     0                    rslv   608     1 ge-2/0/1.0
172.16.1.0/32         iddn     0 172.16.1.0           recv   606     1 ge-2/0/1.0
172.16.1.1/32         user     0                    rjct    46     4
172.16.1.1/32         intf     0 172.16.1.1           locl   607     2
172.16.1.1/32         iddn     0 172.16.1.1           locl   607     2
172.16.1.255/32       iddn     0 ff:ff:ff:ff:ff:ff  bcst   605     1 ge-2/0/1.0
10.0.0.0/24        intf     0                    rslv   616     1 ge-2/0/0.0
10.0.0.0/32        dest     0 10.0.0.0           recv   614     1 ge-2/0/0.0
10.0.0.1/32        intf     0 10.0.0.1           locl   615     2
10.0.0.1/32        dest     0 10.0.0.1           locl   615     2
10.0.0.255/32      dest     0 10.0.0.255         bcst   613     1 ge-2/0/0.0
10.1.1.0/24        ifdn     0                    rslv   612     1 ge-2/0/1.0
10.1.1.0/32        iddn     0 10.1.1.0           recv   610     1 ge-2/0/1.0
10.1.1.1/32        user     0                    rjct    46     4
10.1.1.1/32        intf     0 10.1.1.1           locl   611     2
10.1.1.1/32        iddn     0 10.1.1.1           locl   611     2
10.1.1.255/32      iddn     0 ff:ff:ff:ff:ff:ff  bcst   609     1 ge-2/0/1.0
10.206.0.0/16      user     0 10.209.63.254      ucst   419    20 fxp0.0
10.209.0.0/16      user     1 0:12:1e:ca:98:0    ucst   419    20 fxp0.0
10.209.0.0/18      intf     0                    rslv   418     1 fxp0.0
10.209.0.0/32      dest     0 10.209.0.0         recv   416     1 fxp0.0
10.209.2.131/32    intf     0 10.209.2.131       locl   417     2
10.209.2.131/32    dest     0 10.209.2.131       locl   417     2
10.209.17.55/32    dest     0 0:30:48:5b:78:d2   ucst   435     1 fxp0.0
10.209.63.42/32    dest     0 0:23:7d:58:92:ca   ucst   434     1 fxp0.0
10.209.63.254/32   dest     0 0:12:1e:ca:98:0    ucst   419    20 fxp0.0
10.209.63.255/32   dest     0 10.209.63.255      bcst   415     1 fxp0.0
10.227.0.0/16      user     0 10.209.63.254      ucst   419    20 fxp0.0

...

Routing table: iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                   rjct    27     1
47.0005.80ff.f800.0000.0108.0003.0102.5524.5220.00
intf     0                   locl    28     1

Routing table: inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                   rjct     6     1
ff00::/8           perm     0                   mdsc     4     1
ff02::1/128        perm     0 ff02::1           mcst     3     1

Routing table: ccc
MPLS:
Interface.Label    Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 16     1
100004(top)fe-0/0/1.0
```


Copyright © 2017, Juniper Networks, Inc.

show route forwarding-table detail

```
user@host> show route forwarding-table detail
Routing table: inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            user     2 0:90:69:8e:b1:1b   ucst   132     4 fxp0.0
default            perm     0                    rjct    14     1
10.1.1.0/24        intf     0 ff.3.0.21          ucst   322     1 so-5/3/0.0
10.1.1.0/32        dest     0 10.1.1.0           recv   324     1 so-5/3/0.0
10.1.1.1/32        intf     0 10.1.1.1           locl   321     1
10.1.1.255/32      dest     0 10.1.1.255         bcst   323     1 so-5/3/0.0
10.21.21.0/24      intf     0 ff.3.0.21          ucst   326     1 so-5/3/0.0
10.21.21.0/32      dest     0 10.21.21.0         recv   328     1 so-5/3/0.0
10.21.21.1/32      intf     0 10.21.21.1         locl   325     1
10.21.21.255/32    dest     0 10.21.21.255       bcst   327     1 so-5/3/0.0
127.0.0.1/32       intf     0 127.0.0.1          locl   320     1
172.17.28.19/32    clon     1 192.168.4.254      ucst   132     4 fxp0.0
172.17.28.44/32    clon     1 192.168.4.254      ucst   132     4 fxp0.0

...

Routing table: private1__.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    46     1
10.0.0.0/8         intf     0                    rslv   136     1 fxp1.0
10.0.0.0/32        dest     0 10.0.0.0           recv   134     1 fxp1.0
10.0.0.4/32        intf     0 10.0.0.4           locl   135     2
10.0.0.4/32        dest     0 10.0.0.4           locl   135     2

...

Routing table: iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    38     1

Routing table: inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    22     1
ff00::/8           perm     0                    mdsc    21     1
ff02::1/128        perm     0 ff02::1            mcst    17     1

...

Routing table: mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct 28     1
```

show route forwarding-table destination extensive (Weights and Balances)

```
user@host> show route forwarding-table destination 3.4.2.1 extensive
Routing table: inet [Index 0]
Internet:

Destination:  3.4.2.1/32
  Route type: user
```

```
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: unilist                Index: 262143   Reference: 1
  Nexthop: 172.16.4.4
  Next-hop type: unicast                Index: 335      Reference: 2
  Next-hop interface: so-1/1/0.0        Weight: 22      Balance: 3
  Nexthop: 145.12.1.2
  Next-hop type: unicast                Index: 337      Reference: 2
  Next-hop interface: so-0/1/2.0        Weight: 33      Balance: 33
```

show route forwarding-table extensive

```
user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:

Destination:  default
  Route type: user
  Route reference: 2                    Route interface-index: 0
  Flags: sent to PFE
  Nexthop: 00:00:5E:00:53:1b
  Next-hop type: unicast                Index: 132      Reference: 4
  Next-hop interface: fxp0.0

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: none
  Next-hop type: reject                 Index: 14       Reference: 1

Destination:  127.0.0.1/32
  Route type: interface
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Nexthop: 127.0.0.1
  Next-hop type: local                  Index: 320      Reference: 1

...

Routing table: private1__.inet [Index 1]
Internet:

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                 Index: 46       Reference: 1

Destination:  10.0.0.0/8
  Route type: interface
  Route reference: 0                    Route interface-index: 3
  Flags: sent to PFE
  Next-hop type: resolve                Index: 136      Reference: 1
  Next-hop interface: fxp1.0

...

Routing table: iso [Index 0]
ISO:

Destination:  default
```


Copyright © 2017, Juniper Networks, Inc.

```
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                 Index: 38       Reference: 1

Routing table: inet6 [Index 0]
Internet6:

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                 Index: 22       Reference: 1

Destination:  ff00::/8
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: multicast discard      Index: 21       Reference: 1

...

Routing table: private1__.inet6 [Index 1]
Internet6:

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Next-hop type: reject                 Index: 54       Reference: 1

Destination:  fe80::2a0:a5ff:fe3d:375/128
  Route type: interface
  Route reference: 0                    Route interface-index: 0
  Flags: sent to PFE
  Nexthop: fe80::2a0:a5ff:fe3d:375
  Next-hop type: local                  Index: 75       Reference: 1

...
```

show route forwarding-table extensive (RPF)

The next example is based on the following configuration, which enables an RPF check on all routes that are learned from this interface, including the interface route:

```
so-1/1/0 {
    unit 0 {
        family inet {
            rpf-check;
            address 192.0.2.2/30;
        }
    }
}
```

```
user@host> show route forwarding-table extensive
Routing table: inet [Index 0]
Internet:
...
...
Destination:  192.0.2.3/32
```

```
  Route type: destination
  Route reference: 0                   Route interface-index: 67
  Flags: sent to PFE
  Nexthop: 192.0.2.3
  Next-hop type: broadcast             Index: 328      Reference: 1
  Next-hop interface: so-1/1/0.0
  RPF interface: so-1/1/0.0
```

show route forwarding-table (dynamic list next hop)

The **show route forwarding table** output shows the two next hop elements for a multihomed EVPN destination.

```
user@host> show route forwarding-table label 299952 extensive
MPLS:

Destination:  299952
  Route type: user
  Route reference: 0                   Route interface-index: 0
  Multicast RPF nh index: 0
  P2mpidx: 0
  Flags: sent to PFE, rt nh decoupled
  Next-hop type: indirect              Index: 1048575  Reference: 2
  Nexthop:
  Next-hop type: composite             Index: 601      Reference: 2
  Next-hop type: indirect              Index: 1048574  Reference: 3
  Nexthop: 1.0.0.4
  Next-hop type: Push 301632, Push 299776(top) Index: 600 Reference: 2
  Load Balance Label: None
  Next-hop interface: ge-0/0/1.0
  Next-hop type: indirect              Index: 1048577  Reference: 3
  Nexthop: 1.0.0.4
  Next-hop type: Push 301344, Push 299792(top) Index: 619 Reference: 2
  Load Balance Label: None
  Next-hop interface: ge-0/0/1.0
```

After one of the PE router has been disabled in the EVPN multihomed network, the same **show route forwarding table** output command shows one next hop element and one empty next hop element.

```
user@host> show route forwarding-table label 299952 extensive
Routing table: default.mpls [Index 0]
MPLS:

Destination:  299952
  Route type: user
  Route reference: 0                   Route interface-index: 0
  Multicast RPF nh index: 0
  P2mpidx: 0
  Flags: sent to PFE, rt nh decoupled
  Next-hop type: indirect              Index: 1048575  Reference: 2
  Nexthop:
  Next-hop type: composite             Index: 601      Reference: 2
  Next-hop type: indirect              Index: 1048577  Reference: 3
  Nexthop: 1.0.0.4
  Next-hop type: Push 301344, Push 299792(top) Index: 619 Reference: 2
  Load Balance Label: None
  Next-hop interface: ge-0/0/1.0
```


Copyright © 2017, Juniper Networks, Inc.

show route forwarding-table family mpls

```
user@host> show route forwarding-table family mpls
Routing table: mpls
MPLS:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            perm     0                   rjct    19     1
0                  user     0                   recv    18     3
1                  user     0                   recv    18     3
2                  user     0                   recv    18     3
100000             user     0 10.31.1.6         swap  100001     fe-1/1/0.0
800002             user     0                   Pop              vt-0/3/0.32770

vt-0/3/0.32770 (VPLS)
                   user     0                   indr   351     4
                                                Push 800000, Push 100002(top)
so-0/0/0.0
```

show route forwarding-table family mpls ccc ge-0/0/1.1004

```
user@host>show route forwarding-table mpls ccc ge-0/0/1.1004
Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop          Type  Index    NhRef Netif
ge-0/0/1.1004  (CCC) user     0                    ulst   1048577     2
                                                   comp      754        3
                                                  comp      755        3
                                                comp      756        3

Routing table: __mpls-oam__.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index    NhRef Netif
default            perm     0                    dscd      556     1
```

show route forwarding-table family vpls

```
user@host> show route forwarding-table family vpls
Routing table: green.vpls
VPLS:
Destination        Type RtRef Next hop          Type Index NhRef Netif
default            dynm     0                   flood   353     1
default            perm     0                   rjct    298     1
fe-0/1/0.0         dynm     0                   flood   355     1
00:00:5E:00:53:1f/48             <<<<<Remote CE

                   dynm     0                   indr   351     4
                                                Push 800000, Push 100002(top)
so-0/0/0.0
00:00:5E:00:53:1f/48             <<<<<<Local CE

                   dynm     0                   ucst   354     2 fe-0/1/0.0
```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled)

```
user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing
Destination        Type RtRef Next hop           Type Index    NhRef Netif
```

Copyright © 2017, Juniper Networks, Inc.

```
default             perm     0                   dscd      519     1
lsi.1048832         intf     0                   indr  1048574     4
                               172.16.3.2          Push 262145      621     2
ge-3/0/0.0
00:00:5E:00:53:01/48 user     0                   ucst      590     5 ge-2/3/9.0
0x30003/51          user     0                   comp      627     2
ge-2/3/9.0          intf     0                   ucst      590     5 ge-2/3/9.0
ge-3/1/3.0          intf     0                   ucst      619     4 ge-3/1/3.0
0x30002/51          user     0                   comp      600     2
0x30001/51          user     0                   comp      597     2
```

show route forwarding-table vpls (Broadcast, unknown unicast, and multicast (BUM) hashing is enabled with MAC Statistics)

```
user@host> show route forwarding-table vpls
Routing table: green.vpls
VPLS:
Enabled protocols: BUM hashing, MAC Stats
Destination         Type RtRef Next hop           Type Index    NhRef Netif
default             perm     0                   dscd      519     1
lsi.1048834         intf     0                   indr  1048574     4
                               172.16.3.2          Push 262145      592     2
ge-3/0/0.0
00:19:e2:25:d0:01/48 user     0                   ucst      590     5 ge-2/3/9.0
0x30003/51          user     0                   comp      630     2
ge-2/3/9.0          intf     0                   ucst      590     5 ge-2/3/9.0
ge-3/1/3.0          intf     0                   ucst      591     4 ge-3/1/3.0
0x30002/51          user     0                   comp      627     2
0x30001/51          user     0                   comp      624     2
```

show route forwarding-table family vpls extensive

```
user@host> show route forwarding-table family vpls extensive
Routing table: green.vpls [Index 2]
VPLS:

Destination:  default
  Route type: dynamic
  Route reference: 0                    Route interface-index: 72
  Flags: sent to PFE
  Next-hop type: flood                  Index: 289      Reference: 1
  Next-hop type: unicast                Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0
  Next-hop type: unicast                Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0

Destination:  default
  Route type: permanent
  Route reference: 0                    Route interface-index: 0
  Flags: none
  Next-hop type: discard                Index: 341      Reference: 1

Destination:  fe-0/1/2.0
  Route type: dynamic
  Route reference: 0                    Route interface-index: 69
  Flags: sent to PFE
  Next-hop type: flood                  Index: 293      Reference: 1
  Next-hop type: indirect               Index: 363      Reference: 4
  Next-hop type: Push 800016
  Next-hop interface: at-1/0/1.0
```


Copyright © 2017, Juniper Networks, Inc.

```
  Next-hop type: indirect              Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
  Next-hop type: unicast               Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0

Destination:  fe-0/1/3.0
  Route type: dynamic
  Route reference: 0                   Route interface-index: 70
  Flags: sent to PFE
  Next-hop type: flood                 Index: 292      Reference: 1
  Next-hop type: indirect              Index: 363      Reference: 4
  Next-hop type: Push 800016
  Next-hop interface: at-1/0/1.0
  Next-hop type: indirect              Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
  Next-hop type: unicast               Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0

Destination:  00:00:5E:00:53:01/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 70
  Flags: sent to PFE, prefix load balance
  Next-hop type: unicast               Index: 291      Reference: 3
  Next-hop interface: fe-0/1/3.0
  Route used as destination:
    Packet count:         6640    Byte count:      675786
  Route used as source
    Packet count:         6894    Byte count:      696424

Destination:  00:00:5E:00:53:04/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 69
  Flags: sent to PFE, prefix load balance
  Next-hop type: unicast               Index: 290      Reference: 3
  Next-hop interface: fe-0/1/2.0
  Route used as destination:
    Packet count:           96    Byte count:        8079
  Route used as source:
    Packet count:          296    Byte count:       24955

Destination:  00:00:5E:00:53:05/48
  Route type: dynamic
  Route reference: 0                   Route interface-index: 74
  Flags: sent to PFE, prefix load balance
  Next-hop type: indirect              Index: 301      Reference: 5
  Next hop: 10.31.3.2
  Next-hop type: Push 800000
  Next-hop interface: fe-0/1/1.0
```

show route forwarding-table table default

```
user@host> show route forwarding-table table default
Routing table: default.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    36     2
0.0.0.0/32         perm     0                    dscd    34     1
```

```
10.0.60.0/30       user     0 10.0.60.13         ucst   713     5 fe-0/1/3.0
10.0.60.12/30      intf     0                    rslv   688     1 fe-0/1/3.0
10.0.60.12/32      dest     0 10.0.60.12         recv   686     1 fe-0/1/3.0
10.0.60.13/32      dest     0 0:5:85:8b:bc:22    ucst   713     5 fe-0/1/3.0
10.0.60.14/32      intf     0 10.0.60.14         locl   687     2
10.0.60.14/32      dest     0 10.0.60.14         locl   687     2
10.0.60.15/32      dest     0 10.0.60.15         bcst   685     1 fe-0/1/3.0
10.0.67.12/30      user     0 10.0.60.13         ucst   713     5 fe-0/1/3.0
10.0.80.0/30       ifdn     0 ff.3.0.21          ucst   676     1 so-0/0/1.0
10.0.80.0/32       dest     0 10.0.80.0          recv   678     1 so-0/0/1.0
10.0.80.2/32       user     0                    rjct    36     2
10.0.80.2/32       intf     0 10.0.80.2          locl   675     1
10.0.80.3/32       dest     0 10.0.80.3          bcst   677     1 so-0/0/1.0
10.0.90.12/30      intf     0                    rslv   684     1 fe-0/1/0.0
10.0.90.12/32      dest     0 10.0.90.12         recv   682     1 fe-0/1/0.0
10.0.90.14/32      intf     0 10.0.90.14         locl   683     2
10.0.90.14/32      dest     0 10.0.90.14         locl   683     2
10.0.90.15/32      dest     0 10.0.90.15         bcst   681     1 fe-0/1/0.0
10.5.0.0/16        user     0 192.168.187.126    ucst   324    15 fxp0.0
10.10.0.0/16       user     0 192.168.187.126    ucst   324    15 fxp0.0
10.13.10.0/23      user     0 192.168.187.126    ucst   324    15 fxp0.0
10.84.0.0/16       user     0 192.168.187.126    ucst   324    15 fxp0.0
10.150.0.0/16      user     0 192.168.187.126    ucst   324    15 fxp0.0
10.157.64.0/19     user     0 192.168.187.126    ucst   324    15 fxp0.0
10.209.0.0/16      user     0 192.168.187.126    ucst   324    15 fxp0.0

...

Routing table: default.iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    60     1

Routing table: default.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct    44     1
::/128             perm     0                    dscd    42     1
ff00::/8           perm     0                    mdsc    43     1
ff02::1/128        perm     0 ff02::1            mcst    39     1

Routing table: default.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd    50     1
```

show route forwarding-table table logical-system-name/routing-instance-name

```
user@host> show route forwarding-table table R4/vpn-red
Logical system: R4
Routing table: vpn-red.inet
Internet:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   563     1
0.0.0.0/32         perm     0                    dscd   561     2
172.16.0.1/32         user     0                    dscd   561     2
172.16.2.0/24         intf     0                    rslv   771     1 ge-1/2/0.3
172.16.2.0/32        dest     0 172.16.2.0           recv   769     1 ge-1/2/0.3
172.16.2.1/32         intf     0 172.16.2.1           locl   770     2
172.16.2.1/32         dest     0 172.16.2.1           locl   770     2
172.16.2.2/32         dest     0 0.4.80.3.0.1b.c0.d5.e4.bd.0.1b.c0.d5.e4.bc.8.0
```


Copyright © 2017, Juniper Networks, Inc.

```
                                                ucst   789     1 ge-1/2/0.3
172.16.2.255/32       dest     0 172.16.2.255          bcst   768     1 ge-1/2/0.3
172.16.233.0/4         perm     1                     mdsc   562     1
172.16.233.1/32        perm     0 172.16.233.1           mcst   558     1
255.255.255.255/32 perm     0                     bcst   559     1

Logical system: R4
Routing table: vpn-red.iso
ISO:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   608     1

Logical system: R4
Routing table: vpn-red.inet6
Internet6:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    rjct   708     1
::/128             perm     0                    dscd   706     1
ff00::/8           perm     0                    mdsc   707     1
ff02::1/128        perm     0 ff02::1            mcst   704     1

Logical system: R4
Routing table: vpn-red.mpls
MPLS:
Destination        Type RtRef Next hop           Type Index NhRef Netif
default            perm     0                    dscd   638
```

show route forwarding-table vpn

```
user@host> show route forwarding-table vpn VPN-A
      Routing table:: VPN-A.inet
Internet:
      Destination        Type RtRef Nexthop           Type Index NhRef Netif
      default            perm     0                   rjct     4     4
      10.39.10.20/30     intf     0 ff.3.0.21         ucst    40     1
      so-0/0/0.0
      10.39.10.21/32     intf     0 10.39.10.21       locl    36     1
      10.255.14.172/32   user     0                   ucst    69     2
      so-0/0/0.0
      10.255.14.175/32   user     0                   indr    81     3
                                                      Push 100004, Push
      100004(top) so-1/0/0.0
      172.16.233.0/4         perm     2                   mdsc     5     3
      172.16.233.1/32        perm     0 172.16.233.1         mcst     1     8
      172.16.233.5/32        user     1 172.16.233.5         mcst     1     8
      255.255.255.255/32 perm     0                   bcst     2     3
```

On QFX5200, the results for this command look like this:

```
show route forwarding-table family mpls
Routing table: default.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 65 1
0 user 0 recv 64 4
1 user 0 recv 64 4
2 user 0 recv 64 4
13 user 0 recv 64 4
300384 user 0 9.1.1.1 Pop 1711 2 xe-0/0/34.0
300384(S=0) user 0 9.1.1.1 Pop 1712 2 xe-0/0/34.0
```

Copyright © 2017, Juniper Networks, Inc.

```
300400 user 0 ulst 131071 2
                              10.1.1.2 Pop 1713 1 xe-0/0/38.0
                              172.16.11.2 Pop 1714 1 xe-0/0/40.0
300400(S=0) user 0 ulst 131072 2
                              10.1.1.2 Pop 1715 1 xe-0/0/38.0
                              172.16.11.2 Pop 1716 1 xe-0/0/40.0

Routing table: __mpls-oam__.mpls
MPLS:
Destination Type RtRef Next hop Type Index NhRef Netif
default perm 0 dscd 1681 1
```


Copyright © 2017, Juniper Networks, Inc.

show route hidden

Syntax
show route hidden
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Release Information
Command introduced before Junos OS Release 7.4.

Description
Display only hidden route information. A hidden route is unusable, even if it is the best path.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- *Understanding Hidden Routes*

List of Sample Output
show route hidden on page 579
show route hidden detail on page 580
show route hidden extensive on page 580
show route hidden terse on page 580

Output Fields
For information about output fields, see the output field table for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route hidden

```
user@host> show route hidden
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
127.0.0.1/32        [Direct/0] 04:26:38
                    > via lo0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
10.5.5.5/32         [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                      AS path: 100 I
```

Copyright © 2017, Juniper Networks, Inc.

```
                      Unusable
10.12.1.0/24        [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                      AS path: 100 I
                      Unusable
10.12.80.4/30       [BGP/170] 03:44:10, localpref 100, from 10.4.4.4
                      AS path: I
                      Unusable
...
```

show route hidden detail

```
user@host> show route hidden detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
127.0.0.1/32 (1 entry, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via lo0.0, selected
                State: <Hidden Martian Int>
                Local AS:     1
                Age: 4:27:37
                Task: IF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.5.5.5/32 (1 entry, 0 announced)
         BGP    Preference: 170/-101
                Route Distinguisher: 10.4.4.4:4
                Next hop type: Unusable
                Next-hop reference count: 6
                State: <Secondary Hidden Int Ext>
                Local AS:     1 Peer AS:     1
                Age: 3:45:09
                Task: BGP_1.10.4.4.4+2493
                AS path: 100 I
                Communities: target:1:999
                VPN Label: 100064
                Localpref: 100
                Router ID: 10.4.4.4
                Primary Routing Table bgp.l3vpn.0

...
```

show route hidden extensive

The output for the **show route hidden extensive** command is identical to that of the **show route hidden detail** command. For sample output, see show route hidden detail on page 361.

show route hidden terse

```
user@host> show route hidden terse
```

 Copyright © 2017, Juniper Networks, Inc.

```
inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  127.0.0.1/32       D   0                     >lo0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.5.5.5/32        B 170        100             Unusable        100 I
  10.12.1.0/24       B 170        100             Unusable        100 I
  10.12.80.4/30      B 170        100             Unusable        I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.4.4.4:4:10.5.5.5/32
                     B 170        100             Unusable        100 I
  10.4.4.4:4:10.12.1.0/24
                     B 170        100             Unusable        100 I
  10.4.4.4:4:10.12.80.4/30
                     B 170        100             Unusable        I

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

show route inactive-path

List of Syntax
Syntax on page 582
Syntax (EX Series Switches) on page 582

Syntax
show route inactive-path
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route inactive-path
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display routes for destinations that have no active route. An inactive route is a route that was not selected as the best path.

Options
none—Display all inactive routes.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

Related Documentation
- show route active-path on page 281

List of Sample Output
show route inactive-path on page 582
show route inactive-path detail on page 583
show route inactive-path extensive on page 584
show route inactive-path terse on page 584

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route inactive-path

```
user@host> show route inactive-path

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

```
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.100.12/30     [OSPF/10] 03:57:28, metric 1
                    > via so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.0/8          [Direct/0] 04:39:56
                    > via fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.12.80.0/30       [BGP/170] 04:38:17, localpref 100
                      AS path: 100 I
                    > to 10.12.80.1 via ge-6/3/2.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route inactive-path detail

```
user@host> show route inactive-path detail

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete

10.12.100.12/30 (2 entries, 1 announced)
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/3/0.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Local AS:     1
                Age: 3:58:24    Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

10.0.0.0/8 (2 entries, 0 announced)
         Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via fxp1.0, selected
                State: <NotBest Int>
```

```
                    Inactive reason: No difference
                    Age: 4:40:52
                    Task: IF
                    AS path: I

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete

10.12.80.0/30 (2 entries, 1 announced)
         BGP    Preference: 170/-101
                Next-hop reference count: 6
                Source: 10.12.80.1
                Next hop: 10.12.80.1 via ge-6/3/2.0, selected
                State: <Ext>
                Inactive reason: Route Preference
                Peer AS:   100
                Age: 4:39:13
                Task: BGP_100.10.12.80.1+179
                AS path: 100 I
                Localpref: 100
                Router ID: 10.0.0.0
```

show route inactive-path extensive

The output for the **show route inactive-path extensive** command is identical to that of the **show route inactive-path detail** command. For sample output, see show route inactive-path detail on page 364.

show route inactive-path terse

```
user@host> show route inactive-path terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.100.12/30    O  10          1            >so-0/3/0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.0.0.0/8         D   0                     >fxp1.0

red.inet.0: 6 destinations, 8 routes (4 active, 0 holddown, 3 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  10.12.80.0/30      B 170        100            >10.12.80.1      100 I

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
```


Copyright © 2017, Juniper Networks, Inc.

```
bgp.l3vpn.0: 3 destinations, 3 routes (0 active, 0 holddown, 3 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

show route instance

List of Syntax
Syntax on page 586
Syntax (EX Series Switches and QFX Series) on page 586

Syntax
show route instance
<brief | detail | summary>
<*instance-name*>
<logical-system (all | *logical-system-name*)>
<operational>

Syntax (EX Series Switches and QFX Series)
show route instance
<brief | detail | summary>
<*instance-name*>
<operational>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Command introduced in Junos OS Release 11.3 for the QFX Series.
Command introduced in Junos OS Release 14.1X53-D20 for the OCX Series.

Description
Display routing instance information.

Options
none—(Same as **brief**) Display standard information about all routing instances.

brief | detail | summary—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**. (These options are not available with the **operational** keyword.)

instance-name—(Optional) Display information for all routing instances whose name begins with this string (for example, **cust1**, **cust11**, and **cust111** are all displayed when you run the **show route instance cust1** command).

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

operational—(Optional) Display operational routing instances.

Required Privilege Level
view

Related Documentation
- *Example: Transporting IPv6 Traffic Across IPv4 Using Filter-Based Tunneling*
- *Example: Configuring the Helper Capability Mode for OSPFv3 Graceful Restart*

List of Sample Output
show route instance on page 588
show route instance detail (Graceful Restart Complete) on page 588
show route instance detail (Graceful Restart Incomplete) on page 590

Copyright © 2017, Juniper Networks, Inc.

show route instance detail (VPLS Routing Instance) on page 592
show route instance operational on page 592
show route instance summary on page 592

Output Fields Table 24 on page 368 lists the output fields for the **show route instance** command. Output fields are listed in the approximate order in which they appear.

Table 45: show route instance Output Fields

Field Name	Field Description	Level of Output
Instance or *instance-name*	Name of the routing instance.	All levels
Operational Routing Instances	(**operational** keyword only) Names of all operational routing instances.	—
Type	Type of routing instance: **forwarding**, **l2vpn**, **no-forwarding**, **vpls**, **virtual-router**, or **vrf**.	All levels
State	State of the routing instance: **active** or **inactive**.	**brief detail** none
Interfaces	Name of interfaces belonging to this routing instance.	**brief detail** none
Restart State	Status of graceful restart for this instance: **Pending** or **Complete**.	**detail**
Path selection timeout	Maximum amount of time, in seconds, remaining until graceful restart is declared complete. The default is **300**.	**detail**
Tables	Tables (and number of routes) associated with this routing instance.	**brief detail** none
Route-distinguisher	Unique route distinguisher associated with this routing instance.	**detail**
Vrf-import	VPN routing and forwarding instance import policy name.	**detail**
Vrf-export	VPN routing and forwarding instance export policy name.	**detail**
Vrf-import-target	VPN routing and forwarding instance import target community name.	**detail**
Vrf-export-target	VPN routing and forwarding instance export target community name.	**detail**
Vrf-edge-protection-id	Context identifier configured for edge-protection.	**detail**
Fast-reroute-priority	Fast reroute priority setting for a VPLS routing instance: **high**, **medium**, or **low**. The default is **low**.	**detail**
Restart State	Restart state: • **Pending:*protocol-name***—List of protocols that have not yet completed graceful restart for this routing table. • **Complete**—All protocols have restarted for this routing table.	**detail**

Copyright © 2017, Juniper Networks, Inc.

Table 45: show route instance Output Fields *(continued)*

Field Name	Field Description	Level of Output
Primary rib	Primary table for this routing instance.	**brief** none **summary**
Active/holddown/hidden	Number of active, hold-down, and hidden routes.	All levels

Sample Output

show route instance

```
user@host> show route instance
Instance             Type
         Primary RIB                                    Active/holddown/hidden
master               forwarding
         inet.0                                         16/0/1
         iso.0                                          1/0/0
         mpls.0                                         0/0/0
         inet6.0                                        2/0/0
         l2circuit.0                                    0/0/0
__juniper_private1__ forwarding
         __juniper_private1__.inet.0                    12/0/0
         __juniper_private1__.inet6.0                   1/0/0
```

show route instance detail (Graceful Restart Complete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding        State: Active
  Restart State: Complete Path selection timeout: 300
  Tables:
    inet.0                : 17 routes (15 active, 0 holddown, 1 hidden)
    Restart Complete
    inet.3                : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    iso.0                 : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                : 19 routes (19 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l3vpn.0           : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Complete
    inet6.0               : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0           : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0       : 4 routes (4 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

```
    Restart Complete
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0           : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
    BGP-L.mpls.0           : 3 routes (3 active, 0 holddown, 0 hidden)
    Restart Complete
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn             State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
LDP:
  Router ID: 10.69.105.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0             : 5 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
OSPF:
  Router ID: 10.69.101.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Vrf-import-target: [ target:11111
  Tables:
    OSPF.inet.0            : 8 routes (7 active, 0 holddown, 0 hidden)
    Restart Complete
RIP:
  Router ID: 10.69.102.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
```

Copyright © 2017, Juniper Networks, Inc.

```
  Tables:
    RIP.inet.0              : 6 routes (6 active, 0 holddown, 0 hidden)
    Restart Complete
STATIC:
  Router ID: 10.69.100.1
  Type: vrf               State: Active
  Restart State: Complete Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
    STATIC.inet.0           : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Complete
```

show route instance detail (Graceful Restart Incomplete)

```
user@host> show route instance detail
master:
  Router ID: 10.255.14.176
  Type: forwarding        State: Active
  Restart State: Pending  Path selection timeout: 300
  Tables:
    inet.0                  : 17 routes (15 active, 1 holddown, 1 hidden)
    Restart Pending: OSPF LDP
    inet.3                  : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: OSPF LDP
    iso.0                   : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Complete
    mpls.0                  : 23 routes (23 active, 0 holddown, 0 hidden)
    Restart Pending: LDP VPN
    bgp.l3vpn.0             : 10 routes (10 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
    inet6.0                 : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Complete
    bgp.l2vpn.0             : 1 routes (1 active, 0 holddown, 0 hidden)
    Restart Pending: BGP VPN
BGP-INET:
  Router ID: 10.69.103.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.103
  Route-distinguisher: 10.255.14.176:103
  Vrf-import: [ BGP-INET-import ]
  Vrf-export: [ BGP-INET-export ]
  Tables:
    BGP-INET.inet.0         : 6 routes (5 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
BGP-L:
  Router ID: 10.69.104.1
  Type: vrf               State: Active
  Restart State: Pending  Path selection timeout: 300
  Interfaces:
    t3-0/0/0.104
  Route-distinguisher: 10.255.14.176:104
  Vrf-import: [ BGP-L-import ]
  Vrf-export: [ BGP-L-export ]
  Tables:
    BGP-L.inet.0            : 6 routes (5 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

```
    Restart Pending: VPN
    BGP-L.mpls.0           : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
L2VPN:
  Router ID: 0.0.0.0
  Type: l2vpn              State: Active
  Restart State: Pending   Path selection timeout: 300
  Interfaces:
    t3-0/0/0.512
  Route-distinguisher: 10.255.14.176:512
  Vrf-import: [ L2VPN-import ]
  Vrf-export: [ L2VPN-export ]
  Tables:
    L2VPN.l2vpn.0          : 2 routes (2 active, 0 holddown, 0 hidden)
    Restart Pending: VPN L2VPN
LDP:
  Router ID: 10.69.105.1
  Type: vrf                State: Active
  Restart State: Pending   Path selection timeout: 300
  Interfaces:
    t3-0/0/0.105
  Route-distinguisher: 10.255.14.176:105
  Vrf-import: [ LDP-import ]
  Vrf-export: [ LDP-export ]
  Tables:
    LDP.inet.0             : 5 routes (4 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF LDP VPN
OSPF:
  Router ID: 10.69.101.1
  Type: vrf                State: Active
  Restart State: Pending   Path selection timeout: 300
  Interfaces:
    t3-0/0/0.101
  Route-distinguisher: 10.255.14.176:101
  Vrf-import: [ OSPF-import ]
  Vrf-export: [ OSPF-export ]
  Tables:
    OSPF.inet.0            : 8 routes (7 active, 1 holddown, 0 hidden)
    Restart Pending: OSPF VPN
RIP:
  Router ID: 10.69.102.1
  Type: vrf                State: Active
  Restart State: Pending   Path selection timeout: 300
  Interfaces:
    t3-0/0/0.102
  Route-distinguisher: 10.255.14.176:102
  Vrf-import: [ RIP-import ]
  Vrf-export: [ RIP-export ]
  Tables:
    RIP.inet.0             : 8 routes (6 active, 2 holddown, 0 hidden)
    Restart Pending: RIP VPN
STATIC:
  Router ID: 10.69.100.1
  Type: vrf                State: Active
  Restart State: Pending   Path selection timeout: 300
  Interfaces:
    t3-0/0/0.100
  Route-distinguisher: 10.255.14.176:100
  Vrf-import: [ STATIC-import ]
  Vrf-export: [ STATIC-export ]
  Tables:
```

Copyright © 2017, Juniper Networks, Inc.

```
    STATIC.inet.0          : 4 routes (4 active, 0 holddown, 0 hidden)
    Restart Pending: VPN
```

show route instance detail (VPLS Routing Instance)

```
user@host> show route instance detail test-vpls
test-vpls:
  Router ID: 0.0.0.0
  Type: vpls              State: Active
  Interfaces:
    lsi.1048833
    lsi.1048832
    fe-0/1/0.513
  Route-distinguisher: 10.255.37.65:1
  Vrf-import: [ __vrf-import-test-vpls-internal__ ]
  Vrf-export: [ __vrf-export-test-vpls-internal__ ]
  Vrf-import-target: [ target:300:1 ]
  Vrf-export-target: [ target:300:1 ]
  Vrf-edge-protection-id:  166.1.3.1  Fast-reroute-priority: high
  Tables:
    test-vpls.l2vpn.0          : 3 routes (3 active, 0 holddown, 0 hidden)
```

show route instance operational

```
user@host> show route instance operational
Operational Routing Instances:

master
default
```

show route instance summary

```
user@host> show route instance summary
Instance           Type           Primary rib            Active/holddown/hidden
master             forwarding
                                  inet.0                 15/0/1
                                  iso.0                  1/0/0
                                  mpls.0                 35/0/0
                                  l3vpn.0                0/0/0
                                  inet6.0                2/0/0
                                  l2vpn.0                0/0/0
                                  l2circuit.0            0/0/0
BGP-INET           vrf
                                  BGP-INET.inet.0        5/0/0
                                  BGP-INET.iso.0         0/0/0
                                  BGP-INET.inet6.0       0/0/0
BGP-L              vrf
                                  BGP-L.inet.0           5/0/0
                                  BGP-L.iso.0            0/0/0
                                  BGP-L.mpls.0           4/0/0
                                  BGP-L.inet6.0          0/0/0
L2VPN              l2vpn
                                  L2VPN.inet.0           0/0/0
                                  L2VPN.iso.0            0/0/0
                                  L2VPN.inet6.0          0/0/0
                                  L2VPN.l2vpn.0          2/0/0
LDP                vrf
                                  LDP.inet.0             4/0/0
                                  LDP.iso.0              0/0/0
                                  LDP.mpls.0             0/0/0
```


Copyright © 2017, Juniper Networks, Inc.

```
                                    LDP.inet6.0             0/0/0
                                    LDP.l2circuit.0         0/0/0
OSPF                vrf
                                    OSPF.inet.0             7/0/0
                                    OSPF.iso.0              0/0/0
                                    OSPF.inet6.0            0/0/0
RIP                 vrf
                                    RIP.inet.0              6/0/0
                                    RIP.iso.0               0/0/0
                                    RIP.inet6.0             0/0/0
STATIC              vrf
                                    STATIC.inet.0           4/0/0
                                    STATIC.iso.0            0/0/0
                                    STATIC.inet6.0          0/0/0
```

Copyright © 2017, Juniper Networks, Inc.

show route next-hop

List of Syntax
Syntax on page 594
Syntax (EX Series Switches) on page 594

Syntax
show route next-hop *next-hop*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route next-hop *next-hop*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the entries in the routing table that are being sent to the specified next-hop address.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of ouput.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

next-hop—Next-hop address.

Required Privilege Level
view

List of Sample Output
show route next-hop on page 594
show route next-hop detail on page 595
show route next-hop extensive on page 597
show route next-hop terse on page 598

Output Fields
For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route next-hop

```
user@host> show route next-hop 192.168.71.254

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both

10.10.0.0/16       *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
10.209.0.0/16      *[Static/5] 06:26:25
```


Copyright © 2017, Juniper Networks, Inc.

```
                    > to 192.168.71.254 via fxp0.0
172.16.0.0/12       *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
192.168.0.0/16      *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
192.168.102.0/23    *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
207.17.136.0/24     *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0
207.17.136.192/32   *[Static/5] 06:26:25
                    > to 192.168.71.254 via fxp0.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop detail

```
user@host> show route next-hop 192.168.71.254 detail

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
Restart Complete
10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
```

Copyright © 2017, Juniper Networks, Inc.

```
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

192.168.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

192.168.102.0/23 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

207.17.136.0/24 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

207.17.136.192/32 (1 entry, 1 announced)
        *Static Preference: 5
                Next-hop reference count: 36
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:     1
                Age: 6:27:41
                Task: RT
                Announcement bits (3): 0-KRT 3-Resolve tree 1 5-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete
```


Copyright © 2017, Juniper Networks, Inc.

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop extensive

```
user@host> show route next-hop 192.168.71.254 extensive

inet.0: 18 destinations, 18 routes (17 active, 0 holddown, 1 hidden)
10.10.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.10.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.209.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 10.209.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

172.16.0.0/12 (1 entry, 1 announced)
TSI:
KRT in-kernel 172.16.0.0/12 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

192.168.0.0/16 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.0.0/16 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```

Copyright © 2017, Juniper Networks, Inc.

```
192.168.102.0/23 (1 entry, 1 announced)
TSI:
KRT in-kernel 192.168.102.0/23 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

207.17.136.0/24 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.0/24 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

207.17.136.192/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 207.17.136.192/32 -> {192.168.71.254}
        *Static Preference: 5
                Next-hop reference count: 22
                Next hop: 192.168.71.254 via fxp0.0, selected
                State: <Active NoReadvrt Int Ext>
                Local AS:    69
                Age: 2:02:28
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

green.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

red.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route next-hop terse

```
user@host> show route next-hop 192.168.71.254 terse

inet.0: 25 destinations, 26 routes (24 active, 0 holddown, 1 hidden)
Restart Complete
+ = Active Route, - = Last Active, * = Both
```


Copyright © 2017, Juniper Networks, Inc.

```
A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.10.0.0/16       S   5                       >192.168.71.254
* 10.209.0.0/16      S   5                       >192.168.71.254
* 172.16.0.0/12      S   5                       >192.168.71.254
* 192.168.0.0/16     S   5                       >192.168.71.254
* 192.168.102.0/23   S   5                       >192.168.71.254
* 207.17.136.0/24    S   5                       >192.168.71.254
* 207.17.136.192/32  S   5                       >192.168.71.254

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

red.inet.0: 4 destinations, 5 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Restart Complete

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Restart Complete
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

show route output

List of Syntax
Syntax on page 600
Syntax (EX Series Switches) on page 600

Syntax
show route output (address *ip-address* | interface *interface-name*)
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route output (address *ip-address* | interface *interface-name*)
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the entries in the routing table learned through static routes and interior gateway protocols that are to be sent out the interface with either the specified IP address or specified name.

To view routes advertised to a neighbor or received from a neighbor for the BGP protocol, use the **show route advertising-protocol bgp** and **show route receive-protocol bgp** commands instead.

Options
address *ip-address*—Display entries in the routing table that are to be sent out the interface with the specified IP address.

brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

interface *interface-name*—Display entries in the routing table that are to be sent out the interface with the specified name.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output
show route output address on page 601
show route output address detail on page 601
show route output address extensive on page 602
show route output address terse on page 602
show route output interface on page 602
show route output interface detail on page 603
show route output interface extensive on page 603
show route output interface terse on page 603

Copyright © 2017, Juniper Networks, Inc.

Output Fields For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route output address

```
user@host> show route output address 172.16.36.1/24

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.36.0/24         *[Direct/0] 00:19:56
                     > via so-0/1/2.0
                     [OSPF/10] 00:19:55, metric 1
                     > via so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output address detail

```
user@host> show route output address 172.16.36.1 detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
172.16.36.0/24 (2 entries, 0 announced)
        *Direct Preference: 0
                Next hop type: Interface
                Next-hop reference count: 1
                Next hop: via so-0/1/2.0, selected
                State: <Active Int>
                Age: 23:00
                Task: IF
                AS path: I
         OSPF   Preference: 10
                Next-hop reference count: 1
                Next hop: via so-0/1/2.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Age: 22:59      Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                AS path: I

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

```
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output address extensive

The output for the **show route output address extensive** command is identical to that of the **show route output address detail** command. For sample output, see show route output address detail on page 382.

show route output address terse

```
user@host> show route output address 172.16.36.1 terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 172.16.36.0/24        D   0                      >so-0/1/2.0
                     O  10          1            >so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output interface

```
user@host> show route output interface so-0/1/2.0

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.255.71.240/32   *[OSPF/10] 00:13:00, metric 2
                      via so-0/1/2.0
                    > via so-0/3/2.0
10.255.71.241/32   *[OSPF/10] 00:13:10, metric 1
                    > via so-0/1/2.0
172.16.14.0/24        *[OSPF/10] 00:05:11, metric 3
                      to 35.1.1.2 via ge-3/1/0.0
                    > via so-0/1/2.0
                      via so-0/3/2.0
172.16.16.0/24        *[OSPF/10] 00:13:10, metric 2
                    > via so-0/1/2.0
172.16.36.0/24        *[Direct/0] 00:13:21
                    > via so-0/1/2.0
                    [OSPF/10] 00:13:20, metric 1
                    > via so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

```
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```

show route output interface detail

```
user@host> show route output interface so-0/1/2.0 detail

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
10.255.71.240/32 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 2
                Next hop: via so-0/1/2.0
                Next hop: via so-0/3/2.0, selected
                State: <Active Int>
                Age: 14:52      Metric: 2
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (1): 0-KRT
                AS path: I

10.255.71.241/32 (1 entry, 1 announced)
        *OSPF   Preference: 10
                Next-hop reference count: 4
                Next hop: via so-0/1/2.0, selected
                State: <Active Int>
                Age: 15:02      Metric: 1
                Area: 0.0.0.0
                Task: OSPF
                Announcement bits (1): 0-KRT
                AS path: I
...
```

show route output interface extensive

The output for the **show route output interface extensive** command is identical to that of the **show route output interface detail** command. For sample output, see show route output interface detail on page 384.

show route output interface terse

```
user@host> show route output interface so-0/1/2.0 terse

inet.0: 28 destinations, 30 routes (27 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* 10.255.71.240/32   O  10          2             so-0/1/2.0
                                                 >so-0/3/2.0
* 10.255.71.241/32   O  10          1            >so-0/1/2.0
* 172.16.14.0/24        O  10          3             35.1.1.2
                                                 >so-0/1/2.0
                                                  so-0/3/2.0
* 172.16.16.0/24        O  10          2            >so-0/1/2.0
* 172.16.36.0/24        D   0                       >so-0/1/2.0
                      O  10          1            >so-0/1/2.0

private1__.inet.0: 2 destinations, 3 routes (2 active, 0 holddown, 0 hidden)
```

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
private1__.inet6.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
```


Copyright © 2017, Juniper Networks, Inc.

show route protocol

List of Syntax
Syntax on page 605
Syntax (EX Series Switches) on page 605

Syntax
show route protocol *protocol*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route protocol *protocol*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
ospf2 and **ospf3** options introduced in Junos OS Release 9.2.
ospf2 and **ospf3** options introduced in Junos OS Release 9.2 for EX Series switches.
flow option introduced in Junos OS Release 10.0.
flow option introduced in Junos OS Release 10.0 for EX Series switches.

Description
Display the route entries in the routing table that were learned from a particular protocol.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output. If you do not specify a level of output, the system defaults to **brief**.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

protocol—Protocol from which the route was learned:

- **access**—Access route for use by DHCP application
- **access-internal**—Access-internal route for use by DHCP application
- **aggregate**—Locally generated aggregate route
- **arp**—Route learned through the Address Resolution Protocol
- **atmvpn**—Asynchronous Transfer Mode virtual private network
- **bgp**—Border Gateway Protocol
- **ccc**—Circuit cross-connect
- **direct**—Directly connected route
- **dvmrp**—Distance Vector Multicast Routing Protocol
- **esis**—End System-to-Intermediate System
- **flow**—Locally defined flow-specification route
- **frr**—Precomputed protection route or backup route used when a link goes down
- **isis**—Intermediate System-to-Intermediate System

Copyright © 2017, Juniper Networks, Inc.

- **ldp**—Label Distribution Protocol
- **l2circuit**—Layer 2 circuit
- **l2vpn**—Layer 2 virtual private network
- **local**—Local address
- **mpls**—Multiprotocol Label Switching
- **msdp**—Multicast Source Discovery Protocol
- **ospf**—Open Shortest Path First versions 2 and 3
- **ospf2**—Open Shortest Path First versions 2 only
- **ospf3**—Open Shortest Path First version 3 only
- **pim**—Protocol Independent Multicast
- **rip**—Routing Information Protocol
- **ripng**—Routing Information Protocol next generation
- **rsvp**—Resource Reservation Protocol
- **rtarget**—Local route target virtual private network
- **static**—Statically defined route
- **tunnel**—Dynamic tunnel
- **vpn**—Virtual private network

NOTE: EX Series switches run a subset of these protocols. See the switch CLI for details.

Required Privilege Level view

List of Sample Output

show route protocol access on page 607
show route protocol access-internal extensive on page 607
show route protocol arp on page 607
show route protocol bgp on page 608
show route protocol bgp detail on page 608
show route protocol bgp detail (Labeled Unicast) on page 609
show route protocol bgp detail (Aggregate Extended Community Bandwidth) on page 609
show route protocol bgp extensive on page 610
show route protocol bgp terse on page 611
show route protocol direct on page 611
show route protocol frr on page 611
show route protocol l2circuit detail on page 612

Copyright © 2017, Juniper Networks, Inc.

show route protocol l2vpn extensive on page 613
show route protocol ldp on page 613
show route protocol ldp extensive on page 614
show route protocol ospf (Layer 3 VPN) on page 615
show route protocol ospf detail on page 616
show route protocol rip on page 616
show route protocol rip detail on page 616
show route protocol ripng table inet6 on page 616
show route protocol static detail on page 617

Output Fields For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

show route protocol access

```
user@host> show route protocol access
inet.0: 30380 destinations, 30382 routes (30379 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

13.160.0.3/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
13.160.0.4/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
13.160.0.5/32      *[Access/13] 00:00:09
                    > to 13.160.0.2 via fe-0/0/0.0
```

show route protocol access-internal extensive

```
user@host> show route protocol access-internal 13.160.0.19 extensive
inet.0: 100020 destinations, 100022 routes (100019 active, 0 holddown, 1 hidden)
13.160.0.19/32 (1 entry, 1 announced)
TSI:
KRT in-kernel 13.160.0.19/32 -> {13.160.0.2}
        *Access-internal Preference: 12
                Next-hop reference count: 200000
                Next hop: 13.160.0.2 via fe-0/0/0.0, selected
                State: <Active Int>
       Age: 36
                Task: RPD Unix Domain Server./var/run/rpd_serv.local
                Announcement bits (1): 0-KRT
                AS path: I
```

show route protocol arp

```
user@host>  show route protocol arp
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
```

Copyright © 2017, Juniper Networks, Inc.

```
                      Unusable
20.20.1.4/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.5/32        [ARP/4294967293] 00:04:32, from 20.20.1.1
                      Unusable
20.20.1.6/32        [ARP/4294967293] 00:04:34, from 20.20.1.1
                      Unusable
20.20.1.7/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.8/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.9/32        [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.10/32       [ARP/4294967293] 00:04:35, from 20.20.1.1
                      Unusable
20.20.1.11/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
20.20.1.12/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
20.20.1.13/32       [ARP/4294967293] 00:04:33, from 20.20.1.1
                      Unusable
...
```

show route protocol bgp

```
user@host> show route protocol bgp 192.168.64.0/21
inet.0: 335832 destinations, 335833 routes (335383 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.64.0/21       *[BGP/170] 6d 10:41:16, localpref 100, from 192.168.69.71
                      AS path: 10458 14203 2914 4788 4788 I
                    > to 192.168.167.254 via fxp0.0
```

show route protocol bgp detail

```
user@host> show route protocol bgp 66.117.63.0/24 detail
inet.0: 335805 destinations, 335806 routes (335356 active, 0 holddown, 450 hidden)
66.117.63.0/24  (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006436
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                State: <Active Ext>
                Local AS:    69 Peer AS: 10458
                Age: 6d 10:42:42        Metric2: 0
                Task: BGP_10458.192.168.69.71+179
                Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
                AS path: 10458 14203 2914 4788 4788 I
                Communities: 2914:410 2914:2403 2914:3400
                Accepted
                Localpref: 100
                Router ID: 207.17.136.192
```


Copyright © 2017, Juniper Networks, Inc.

show route protocol bgp detail (Labeled Unicast)

```
user@host> show route protocol bgp 1.1.1.8/32 detail
inet.0: 45 destinations, 46 routes (45 active, 0 holddown, 0 hidden)
1.1.1.8/32 (2 entries, 2 announced)
State:
*BGP Preference: 1/-101
Next hop type: Indirect, Next hop index: 0
Address: 0xc007f30
Next-hop reference count: 2
Source: 1.1.1.1
Next hop type: Router, Next hop index: 614
Next hop: 20.1.1.2 via ge-0/0/1.0, selected
Label-switched-path lsp1
Label operation: Push 1000126, Push 1000125, Push 1000124, Push 1000123, Push
299872(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl, prop-ttl, prop-ttl(top)
Load balance label: Label 1000126: None; Label 1000125: None; Label 1000124: None;
 Label 1000123: None; Label 299872: None;
Label element ptr: 0xc007860
Label parent element ptr: 0xc0089a0
Label element references: 1
Label element child references: 0
Label element lsp id: 0
Session Id: 0x140
Protocol next hop: 1.1.1.4
Label operation: Push 1000126, Push 1000125, Push 1000124, Push 1000123(top)
Label TTL action: prop-ttl, prop-ttl, prop-ttl, prop-ttl
Load balance label: Label 1000126: None; Label 1000125: None; Label 1000124: None;
 Label 1000123: None;
Indirect next hop: 0xae8d300 1048576 INH Session ID: 0x142
State:
Local AS: 5 Peer AS: 5
Age: 22:43 Metric2: 2
Validation State: unverified
Task: BGP_5.1.1.1.1
Announcement bits (2): 0-KRT 7-Resolve tree 2
AS path: I
Accepted
Route Labels: 1000123(top) 1000124 1000125 1000126
Localpref: 100
Router ID: 1.1.1.1
```

show route protocol bgp detail (Aggregate Extended Community Bandwidth)

```
user@host> show route 10.0.2.0 protocol bgp detail
inet.0: 20 destinations, 26 routes (20 active, 0 holddown, 0 hidden)
10.0.2.0/30 (2 entries, 1 announced)
        *BGP    Preference: 170/-101
                Next hop type: Router, Next hop index: 0
                Address: 0xb618990
                Next-hop reference count: 3
                Source: 10.0.1.1
                Next hop: 10.0.0.2 via ge-0/0/0.0 balance 40%
                Session Id: 0x0
                Next hop: 10.0.1.1 via ge-0/0/1.0 balance 60%, selected
                Session Id: 0x0
                State: <Active Ext>
                Local AS: 65000 Peer AS: 65001
                Age: 20:33
```

Copyright © 2017, Juniper Networks, Inc.

```
                Validation State: unverified
                Task: BGP_65001.10.0.1.1
                Announcement bits (3): 0-KRT 2-BGP_Listen.0.0.0.0+179
3-BGP_RT_Background
                AS path: 65001 I
                Communities: bandwidth:65000:60000000
         Accepted Multipath
                Localpref: 100
                Router ID: 128.49.121.137
        BGP     Preference: 170/-101
                Next hop type: Router, Next hop index: 595
                Address: 0xb7a1330
                Next-hop reference count: 9
                Source: 10.0.0.2
                Next hop: 10.0.0.2 via ge-0/0/0.0, selected
                Session Id: 0x141
                State: <NotBest Ext>
                Inactive reason: Not Best in its group - Active preferred
                Local AS: 65000 Peer AS: 65001
                Age: 20:33
                Validation State: unverified
                Task: BGP_65001.10.0.0.2
                AS path: 65001 I
                Communities: bandwidth:65000:40000000
                Accepted MultipathContrib
                Localpref: 100
                Router ID: 128.49.121.132
```

show route protocol bgp extensive

```
user@host> show route protocol bgp 192.168.64.0/21 extensive

inet.0: 335827 destinations, 335828 routes (335378 active, 0 holddown, 450 hidden)
192.168.64.0/21 (1 entry, 1 announced)
TSI:
KRT in-kernel 1.9.0.0/16 -> {indirect(342)}
Page 0 idx 1 Type 1 val db31a80
    Nexthop: Self
    AS path: [69] 10458 14203 2914 4788 4788 I
    Communities: 2914:410 2914:2403 2914:3400
Path 1.9.0.0 from 192.168.69.71 Vector len 4.  Val: 1
        *BGP    Preference: 170/-101
                Next hop type: Indirect
                Next-hop reference count: 1006502
                Source: 192.168.69.71
                Next hop type: Router, Next hop index: 324
                Next hop: 192.168.167.254 via fxp0.0, selected
                Protocol next hop: 192.168.69.71
                Indirect next hop: 8e166c0 342
                State: <Active Ext>
                Local AS:    69 Peer AS: 10458
                Age: 6d 10:44:45        Metric2: 0
                Task: BGP_10458.192.168.69.71+179
                Announcement bits (3): 0-KRT 2-BGP RT Background 3-Resolve tree
1
                AS path: 10458 14203 2914 4788 4788 I
                Communities: 2914:410 2914:2403 2914:3400
                Accepted
                Localpref: 100
                Router ID: 207.17.136.192
                Indirect next hops: 1
```


Copyright © 2017, Juniper Networks, Inc.

```
                                    Protocol next hop: 192.168.69.71
                                    Indirect next hop: 8e166c0 342
                                    Indirect path forwarding next hops: 1
                                            Next hop type: Router
                                            Next hop: 192.168.167.254 via fxp0.0
                                    192.168.0.0/16 Originating RIB: inet.0
                                      Node path count: 1
                                      Forwarding nexthops: 1
                                            Nexthop: 192.168.167.254 via fxp0.0
```

show route protocol bgp terse

```
user@host> show route protocol bgp 192.168.64.0/21 terse

inet.0: 24 destinations, 32 routes (23 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

A Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
  192.168.64.0/21    B 170        100            >172.16.100.1    10023 21 I
```

show route protocol direct

```
user@host> show route protocol direct

inet.0: 335843 destinations, 335844 routes (335394 active, 0 holddown, 450 hidden)
+ = Active Route, - = Last Active, * = Both

172.16.8.0/24          *[Direct/0] 17w0d 10:31:49
                        > via fe-1/3/1.0
10.255.165.1/32     *[Direct/0] 25w4d 04:13:18
                        > via lo0.0
172.16.30.0/24       *[Direct/0] 17w0d 23:06:26
                        > via fe-1/3/2.0
192.168.164.0/22    *[Direct/0] 25w4d 04:13:20
                        > via fxp0.0

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

47.0005.80ff.f800.0000.0108.0001.0102.5516.5001/152
                    *[Direct/0] 25w4d 04:13:21
                        > via lo0.0

inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2001:db8::10:255:165:1/128
                    *[Direct/0] 25w4d 04:13:21
                        > via lo0.0
fe80::2a0:a5ff:fe12:ad7/128
                    *[Direct/0] 25w4d 04:13:21
                        > via lo0.0
```

show route protocol frr

```
user@host> show route protocol frr
inet.0: 43 destinations, 43 routes (42 active, 0 holddown, 1 hidden)

inet.3: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

```
cust1.inet.0: 1033 destinations, 2043 routes (1033 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

20.20.1.3/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.3 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.4/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.4 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.5/32       *[FRR/200] 00:05:35, from 20.20.1.1
                    > to 20.20.1.5 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.6/32       *[FRR/200] 00:05:37, from 20.20.1.1
                    > to 20.20.1.6 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.7/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.7 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.8/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.8 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.9/32       *[FRR/200] 00:05:38, from 20.20.1.1
                    > to 20.20.1.9 via ge-4/1/0.0
                      to 10.10.15.1 via ge-0/2/4.0, Push 16, Push 299792(top)
20.20.1.10/32      *[FRR/200] 00:05:38, from 20.20.1.1
...
```

show route protocol l2circuit detail

```
user@host> show route protocol l2circuit detail

mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
100000 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via ge-2/0/0.0, selected
                Label operation: Pop       Offset: 4
                State: <Active Int>
                Local AS:    99
                Age: 9:52
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

ge-2/0/0.0 (1 entry, 1 announced)
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000, Push 100000(top)[0] Offset: -4
                Protocol next hop: 10.245.255.63
                Push 100000 Offset: -4
                 Indirect next hop: 86af0c0 298
                State: <Active Int>
                Local AS:    99
                Age: 9:52
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I

l2circuit.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

10.245.255.63:CtrlWord:4:3:Local/96 (1 entry, 1 announced)
```


Copyright © 2017, Juniper Networks, Inc.

```
        *L2CKT  Preference: 7
                Next hop: via so-1/1/2.0 weight 1, selected
                Label-switched-path my-lsp
                Label operation: Push 100000[0]
                Protocol next hop: 10.245.255.63 Indirect next hop: 86af000 296
                State: <Active Int>
                Local AS:    99
                Age: 10:21
                Task: l2 circuit
                Announcement bits (1): 0-LDP
                AS path: I
                VC Label 100000, MTU 1500, VLAN ID 512
```

show route protocol l2vpn extensive

```
user@host> show route protocol l2vpn extensive

inet.0: 14 destinations, 15 routes (13 active, 0 holddown, 1 hidden)

inet.3: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden)
800001 (1 entry, 1 announced)
TSI:
KRT in-kernel 800001 /36 -> {so-0/0/0.0}
        *L2VPN  Preference: 7
                Next hop: via so-0/0/0.0 weight 49087 balance 97%, selected
                Label operation: Pop       Offset: 4
                State: <Active Int>
                Local AS:    69
                Age: 7:48
                Task: Common L2 VC
                Announcement bits (1): 0-KRT
                AS path: I

so-0/0/0.0 (1 entry, 1 announced)
TSI:
KRT in-kernel so-0/0/0.0.0      /16 -> {indirect(288)}
        *L2VPN  Preference: 7
                Next hop: via so-0/0/1.0, selected
                Label operation: Push 800000 Offset: -4
                Protocol next hop: 10.255.14.220
                Push 800000 Offset: -4
                 Indirect next hop: 85142a0 288
                State: <Active Int>
                Local AS:    69
                Age: 7:48
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 1-Common L2 VC
                AS path: I
                Communities: target:69:1 Layer2-info: encaps:PPP,
                control flags:2, mtu: 0
```

show route protocol ldp

```
user@host> show route protocol ldp
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

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

192.168.16.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Push 100000
192.168.17.1/32    *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0

private1__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100064             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100064(S=0)        *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Pop
100080             *[LDP/9] 1d 23:03:35, metric 1
                    > via t1-4/0/0.0, Swap 100000
```

show route protocol ldp extensive

```
user@host> show route protocol ldp extensive
192.168.16.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *LDP    Preference: 9
                Next-hop reference count: 3
                Next hop: via t1-4/0/0.0, selected
                Label operation: Push 100000
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
                AS path: I

192.168.17.1/32 (1 entry, 1 announced)
        State: <FlashAll>
        *LDP    Preference: 9
                Next-hop reference count: 3
                Next hop: via t1-4/0/0.0, selected
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (2): 0-Resolve tree 1 2-Resolve tree 2
                AS path: I

private1__.inet.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

mpls.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)

100064 (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /36 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
```


Copyright © 2017, Juniper Networks, Inc.

```
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 192.168.17.1/32

100064(S=0) (1 entry, 1 announced)
TSI:
KRT in-kernel 100064 /40 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                Label operation: Pop
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I

100080 (1 entry, 1 announced)
TSI:
KRT in-kernel 100080 /36 -> {t1-4/0/0.0}
        *LDP    Preference: 9
                Next-hop reference count: 2
                Next hop: via t1-4/0/0.0, selected
                Label operation: Swap 100000
                State: <Active Int>
                Local AS: 64500
                Age: 1d 23:03:58        Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 192.168.16.1/32
```

show route protocol ospf (Layer 3 VPN)

```
user@host> show route protocol ospf
inet.0: 40 destinations, 40 routes (39 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.4/30       *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.39.1.8/30        [OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
10.255.14.171/32   *[OSPF/10] 00:05:18, metric 4
                    > via t3-3/2/0.0
10.255.14.179/32   *[OSPF/10] 00:05:18, metric 2
                    > via t3-3/2/0.0
172.16.233.5/32       *[OSPF/10] 20:25:55, metric 1

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30       [OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
10.255.14.173/32   *[OSPF/10] 00:05:43, metric 1
                    > via so-0/2/2.0
172.16.233.5/32       *[OSPF/10] 20:26:20, metric 1
```

Copyright © 2017, Juniper Networks, Inc.

show route protocol ospf detail

```
user@host> show route protocol ospf detail
VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.39.1.16/30 (2 entries, 0 announced)
         OSPF   Preference:  10
                Nexthop: via so-0/2/2.0, selected
                State: <Int>
                Inactive reason: Route Preference
                Age: 6:25       Metric: 1
                Area: 0.0.0.0
                Task: VPN-AB-OSPF
                AS path: I
                Communities: Route-Type:0.0.0.0:1:0

...
```

show route protocol rip

```
user@host> show route protocol rip
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32   *[RIP/100] 20:24:34, metric 2
                    > to 10.39.1.22 via t3-0/2/2.0
172.16.233.9/32      *[RIP/100] 00:03:59, metric 1
```

show route protocol rip detail

```
user@host> show route protocol rip detail
inet.0: 26 destinations, 27 routes (25 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

VPN-AB.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both
10.255.14.177/32 (1 entry, 1 announced)
        *RIP    Preference: 100
                Nexthop: 10.39.1.22 via t3-0/2/2.0, selected
                State: <Active Int>
                Age: 20:25:02   Metric: 2
                Task: VPN-AB-RIPv2
                Announcement bits (2): 0-KRT 2-BGP.0.0.0.0+179
                AS path: I
                Route learned from 10.39.1.22 expires in 96 seconds
```

show route protocol ripng table inet6

```
user@host> show route protocol ripng table inet6
inet6.0: 4215 destinations, 4215 routes (4214 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

1111::1/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::2/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::3/128        *[RIPng/100] 02:13:33, metric 2
```


Copyright © 2017, Juniper Networks, Inc.

```
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::4/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::5/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
1111::6/128        *[RIPng/100] 02:13:33, metric 2
                    > to fe80::2a0:a5ff:fe3d:56 via t3-0/2/0.0
```

show route protocol static detail

```
user@host> show route protocol static detail
inet.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
10.5.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.10.0.0/16 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I

10.13.10.0/23 (1 entry, 1 announced)
        *Static Preference: 5
                Next hop type: Router, Next hop index: 324
                Address: 0x9274010
                Next-hop reference count: 27
                Next hop: 192.168.187.126 via fxp0.0, selected
                Session Id: 0x0
                State: <Active NoReadvrt Int Ext>
                Age: 7w3d 21:24:25
                Validation State: unverified
                Task: RT
                Announcement bits (1): 0-KRT
                AS path: I
```

Copyright © 2017, Juniper Networks, Inc.

show route receive-protocol

List of Syntax
Syntax on page 618
Syntax (EX Series Switches) on page 618

Syntax
show route receive-protocol *protocol neighbor-address*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)

Syntax (EX Series Switches)
show route receive-protocol *protocol neighbor-address*
<brief | detail | extensive | terse>

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description
Display the routing information as it was received through a particular neighbor using a particular dynamic routing protocol.

Options
brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

protocol neighbor-address—Protocol transmitting the route (**bgp**, **dvmrp**, **msdp**, **pim**, **rip**, or **ripng**) and address of the neighboring router from which the route entry was received.

Additional Information
The output displays the selected routes and the attributes with which they were received, but does not show the effects of import policy on the routing attributes.

Required Privilege Level
view

List of Sample Output
show route receive-protocol bgp on page 621
show route receive-protocol bgp extensive on page 621
show route receive-protocol bgp table extensive on page 622
show route receive-protocol bgp logical-system extensive on page 622
show route receive-protocol bgp detail (Layer 2 VPN) on page 623
show route receive-protocol bgp extensive (Layer 2 VPN) on page 623
show route receive-protocol bgp (Layer 3 VPN) on page 624
show route receive-protocol bgp detail (Layer 3 VPN) on page 624
show route receive-protocol bgp detail (Long-Lived Graceful Restart) on page 625
show route receive-protocol bgp detail (Labeled Unicast) on page 625
show route receive-protocol bgp extensive (Layer 3 VPN) on page 626
Show route receive protocol (Segment Routing Traffic Engineering on page 626

Copyright © 2017, Juniper Networks, Inc.

Output Fields Table 25 on page 400 describes the output fields for the **show route receive-protocol** command. Output fields are listed in the approximate order in which they appear.

Table 46: show route receive-protocol Output Fields

Field Name	Field Description	Level of Output
routing-table-name	Name of the routing table—for example, inet.0.	All levels
number destinations	Number of destinations for which there are routes in the routing table.	All levels
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** • **holddown** (routes that are in pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)	All levels
Prefix	Destination prefix.	none **brief**
MED	Multiple exit discriminator value included in the route.	none **brief**
destination-prefix (entry, announced)	Destination prefix. The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination.	**detail extensive**
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag.	**detail extensive**
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag may be displayed for a route. Neither of these flags are displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table	**detail extensive**
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.	**detail extensive**
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.	**detail extensive**
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 46: show route receive-protocol Output Fields *(continued)*

Field Name	Field Description	Level of Output
VPN Label	Virtual private network (VPN) label. Packets are sent between CE and PE routing devices by advertising VPN labels. VPN labels transit over either an RSVP or an LDP label-switched path (LSP) tunnel.	**detail extensive**
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.	All levels
Localpref or Lclpref	Local preference value included in the route.	All levels
AS path	Autonomous system (AS) path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the router, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.	All levels
Route Labels	Stack of labels carried in the BGP route update.	**detail extensive**
Cluster list	(For route reflected output only) Cluster ID sent by the route reflector.	**detail extensive**
Originator ID	(For route reflected output only) Address of routing device that originally sent the route to the route reflector.	**detail extensive**
Communities	Community path attribute for the route. See the Output Field table in the show route detail command for all possible values for this field.	**detail extensive**
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.	**detail extensive**
Attrset AS	Number, local preference, and path of the AS that originated the route. These values are stored in the **Attrset** attribute at the originating routing device.	**detail extensive**

Copyright © 2017, Juniper Networks, Inc.

Table 46: show route receive-protocol Output Fields *(continued)*

Field Name	Field Description	Level of Output
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).	**detail extensive**
control flags	Control flags: **none** or **Site Down**.	**detail extensive**
mtu	Maximum transmission unit (MTU) of the Layer 2 circuit.	**detail extensive**

Sample Output

show route receive-protocol bgp

```
user@host> show route receive-protocol bgp 10.255.245.215

inet.0: 28 destinations, 33 routes (27 active, 0 holddown, 1 hidden)
Prefix                  Next hop              MED     Lclpref    AS path
10.22.1.0/24            10.255.245.215        0       100        I
10.22.2.0/24            10.255.245.215        0       100        I
```

show route receive-protocol bgp extensive

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
Prefix             Next hop               MED    Lclpref AS path
172.16.1.0/24 (1 entry, 1 announced)
     Next hop: 10.0.50.3
     Localpref: 100
     AS path: I <Orginator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.45
172.16.163.0/16 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.68
172.16.164.0/16 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.45
172.16.195.0/24 (1 entry, 1 announced)
     Next hop: 111.222.5.254
     Localpref: 100
     AS path: I <Originator>
     Cluster list:  10.2.3.1
     Originator ID: 10.255.245.68
inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
Prefix             Next hop               MED    Lclpref AS path
inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Next hop               MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Next hop               MED    Lclpref AS path
mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

show route receive-protocol bgp table extensive

```
user@host> show route receive-protocol bgp 207.17.136.192 table inet.0 66.117.68.0/24 extensive
inet.0: 227315 destinations, 227316 routes (227302 active, 0 holddown, 13 hidden)
* 66.117.63.0/24 (1 entry, 1 announced)
     Nexthop: 207.17.136.29
     Localpref: 100
     AS path: AS2 PA[6]: 14203 2914 3356 29748 33437 AS_TRANS
     AS path: AS4 PA[2]: 33437 393219
     AS path: Merged[6]: 14203 2914 3356 29748 33437 393219 I
     Communities: 2914:420
```

show route receive-protocol bgp logical-system extensive

```
user@host> show route receive-protocol bgp 10.0.0.9 logical-system PE4 extensive
inet.0: 12 destinations, 13 routes (12 active, 0 holddown, 0 hidden)
* 10.0.0.0/30 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 10.0.0.4/30 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

  10.0.0.8/30 (2 entries, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 10.9.9.1/32 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 10.100.1.1/32 (1 entry, 1 announced)
     Accepted
     Route Label: 3
     Nexthop: 10.0.0.9
     AS path: 13979 I

* 172.16.44.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300096
     Nexthop: 10.0.0.9
     AS path: 13979 I
     AIGP: 203

* 172.16.55.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300112
     Nexthop: 10.0.0.9
     AS path: 13979 7018 I
     AIGP: 25
```


Copyright © 2017, Juniper Networks, Inc.

```
* 172.16.66.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300144
     Nexthop: 10.0.0.9
     AS path: 13979 7018 I

* 172.16.99.0/24 (1 entry, 1 announced)
     Accepted
     Route Label: 300160
     Nexthop: 10.0.0.9
     AS path: 13979 7018 I
```

show route receive-protocol bgp detail (Layer 2 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171 detail
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix             Nexthop               MED    Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0
hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags: 0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
```

show route receive-protocol bgp extensive (Layer 2 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171 extensive
inet.0: 68 destinations, 68 routes (67 active, 0 holddown, 1 hidden)
Prefix             Nexthop               MED    Lclpref AS path
inet.3: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
mpls.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
frame-vpn.l2vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 1 announced)
     Route Distinguisher: 10.255.245.35:1
```

Copyright © 2017, Juniper Networks, Inc.

```
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Prefix             Nexthop               MED    Lclpref AS path
10.255.245.35:1:5:1/96 (1 entry, 0 announced)
     Route Distinguisher: 10.255.245.35:1
     Label-base : 800000, range : 4, status-vector : 0x0
     Nexthop: 10.255.245.35
     Localpref: 100
     AS path: I
     Communities: target:65299:100 Layer2-info: encaps:FRAME RELAY,
     control flags:0, mtu: 0
```

show route receive-protocol bgp (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.171
      inet.0: 33 destinations, 33 routes (32 active, 0 holddown, 1 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      VPN-A.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      10.255.14.175/32   10.255.14.171                    100 2 I
      10.255.14.179/32   10.255.14.171           2        100 I
      VPN-B.inet.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      10.255.14.175/32   10.255.14.171                    100 2 I
      10.255.14.177/32   10.255.14.171                    100 I
      iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      mpls.0: 9 destinations, 9 routes (9 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      bgp.l3vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
      Prefix             Nexthop               MED    Lclpref AS path
      10.255.14.171:300:10.255.14.177/32
                         10.255.14.171                    100 I
      10.255.14.171:100:10.255.14.179/32
                         10.255.14.171           2        100 I
      10.255.14.171:200:10.255.14.175/32
                         10.255.14.171                    100 2 I
```

show route receive-protocol bgp detail (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.14.174 detail
inet.0: 16 destinations, 17 routes (15 active, 0 holddown, 1 hidden)
inet.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
vpna.inet.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
* 10.49.0.0/30 (1 entry, 1 announced)
     Route Distinguisher: 10.255.14.176:2
     VPN Label: 101264
     Nexthop: 10.255.14.174
     Localpref: 100
     AS path: I
     Communities: target:200:100
     AttrSet AS: 100
         Localpref: 100
```


Copyright © 2017, Juniper Networks, Inc.

```
        AS path: I
* 10.255.14.172/32 (1 entry, 1 announced)
    Route Distinguisher: 10.255.14.176:2
    VPN Label: 101280
    Nexthop: 10.255.14.174
    Localpref: 100
    AS path: I
    Communities: target:200:100
    AttrSet AS: 100
        Localpref: 100
        AS path: I
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
mpls.0: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
* 10.255.14.174:2:10.49.0.0/30 (1 entry, 0 announced)
    Route Distinguisher: 10.255.14.174:2
    VPN Label: 101264
    Nexthop: 10.255.14.174
    Localpref: 100
    AS path: I
    Communities: target:200:100
    AttrSet AS: 100
        Localpref: 100
        AS path: I
* 10.255.14.174:2:10.255.14.172/32 (1 entry, 0 announced)
    Route Distinguisher: 10.255.14.174:2
    VPN Label: 101280
    Nexthop: 10.255.14.174
    Localpref: 100
    AS path: I
    Communities: target:200:100
    AttrSet AS: 100
        Localpref: 100
        AS path: I
inet6.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp detail (Long-Lived Graceful Restart)

```
user@host> show route receive-protocol bgp 10.4.12.11 detail

  bgp.l2vpn.0: 38 destinations, 39 routes (37 active, 0 holddown, 1 hidden)
  * 172.16.1.4:100:172.16.1.4/96 AD (1 entry, 1 announced)
    Accepted LongLivedStale LongLivedStaleImport
    Nexthop: 10.4.12.11
    Localpref: 100
    AS path: I
```

show route receive-protocol bgp detail (Labeled Unicast)

```
user@host> show route receive-protocol bgp 1.1.1.1 detail
inet.0: 45 destinations, 46 routes (45 active, 0 holddown, 0 hidden)
* 1.1.1.8/32 (2 entries, 2 announced)
Accepted
Route Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: 1.1.1.4
Localpref: 100
AS path: I
Entropy label capable, next hop field matches route next hop

inet.3: 15 destinations, 21 routes (6 active, 0 holddown, 14 hidden)
```

Copyright © 2017, Juniper Networks, Inc.

```
iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)

mpls.0: 11 destinations, 11 routes (11 active, 0 holddown, 0 hidden)

inet6.0: 26 destinations, 28 routes (26 active, 0 holddown, 0 hidden)

* 100::1/128 (2 entries, 2 announced)
Accepted
Route Labels: 1000123(top) 1000124 1000125 1000126
Nexthop: ::ffff:1.1.1.4
Localpref: 100
AS path: I

inet6.3: 22 destinations, 23 routes (22 active, 0 holddown, 0 hidden)
```

show route receive-protocol bgp extensive (Layer 3 VPN)

```
user@host> show route receive-protocol bgp 10.255.245.63 extensive
inet.0: 244 destinations, 244 routes (243 active, 0 holddown, 1 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     172.16.1.0/24 (1 entry, 1 announced)
          Nexthop: 10.0.50.3
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     172.16.163.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68
     172.16.164.0/16 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.45
     172.16.195.0/24 (1 entry, 1 announced)
          Nexthop: 111.222.5.254
          Localpref: 100
          AS path: I <Originator>
          Cluster list:  10.2.3.1
          Originator ID: 10.255.245.68
     inet.2: 63 destinations, 63 routes (63 active, 0 holddown, 0 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     inet.3: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     iso.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
     Prefix             Nexthop               MED    Lclpref AS path
     mpls.0: 48 destinations, 48 routes (48 active, 0 holddown, 0 hidden)
```

Show route receive protocol (Segment Routing Traffic Engineering

```
show route receive protocol bgp 10.1.1.4
bgp.inetcolor.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

* 50-4.4.4.4-1234(sr6)/96 (1 entry, 0 announced)
     Import Accepted
```


Copyright © 2017, Juniper Networks, Inc.

```
     Distinguisher: 50
     Color: 1234
     Nexthop: 10.1.1.4
     Localpref: 100
     AS path: 3 I
     Communities: target:1.1.1.1:1


inetcolor.0: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
* 4.4.4.4-1234(c6)/64 (1 entry, 1 announced)
     Import Accepted
     Color: 1234
     Nexthop: 10.1.1.4
     Localpref: 100
     AS path: 3 I
     Communities: target:1.1.1.1:1


user@host# run show route receive-protocol bgp 5001:1::4
bgp.inet6color.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)

* 50-2001:1::4-1234(sr6)/192 (1 entry, 0 announced)
     Import Accepted
     Distinguisher: 50
     Color: 1234
     Nexthop: ::ffff:1.1.1.4
     Localpref: 100
     AS path: 3 I
     Communities: target:1.1.1.1:1


inet6color.0: 6 destinations, 7 routes (6 active, 0 holddown, 0 hidden)
* 2001::5-1234(c6)/160 (1 entry, 1 announced)
     Import Accepted
     Color: 1234
     Nexthop: ::ffff:1.1.1.5
     Localpref: 100
     AS path: 3 I
     Communities: target:2:1
```

Copyright © 2017, Juniper Networks, Inc.

show route table

List of Syntax

Syntax on page 628
Syntax (EX Series Switches, QFX Series Switches) on page 628

Syntax

show route table *routing-table-name*
<brief | detail | extensive | terse>
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches, QFX Series Switches)

show route table *routing-table-name*
<brief | detail | extensive | terse>

Release Information

Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.
Statement introduced in Junos OS Release 14.1X53-D15 for QFX Series switches.
Show route table evpn statement introduced in Junos OS Release 15.1X53-D30 for QFX Series switches.

Description

Display the route entries in a particular routing table.

Options

brief | detail | extensive | terse—(Optional) Display the specified level of output.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

routing-table-name—Display route entries for all routing tables whose names begin with this string (for example, inet.0 and inet6.0 are both displayed when you run the **show route table inet** command).

Required Privilege Level

view

Related Documentation

- *show route summary*

List of Sample Output

show route table bgp.l2.vpn on page 639
show route table bgp.l3vpn.0 on page 639
show route table bgp.l3vpn.0 detail on page 640
show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured) on page 641
show route table bgp.evpn.0 on page 641
show route table evpna.evpn.0 on page 642
show route table inet.0 on page 642
show route table inet.3 on page 643
show route table inet.3 protocol ospf on page 643
show route table inet6.0 on page 643
show route table inet6.3 on page 643

Copyright © 2017, Juniper Networks, Inc.

show route table inetflow detail on page 644
show route table lsdist.0 extensive on page 644
show route table l2circuit.0 on page 646
show route table mpls on page 646
show route table mpls extensive on page 646
show route table mpls.0 on page 647
show route table mpls.0 detail (PTX Series) on page 648
show route table mpls.0 ccc ge-0/0/1.1004 detail on page 648
show route table mpls.0 protocol evpn on page 649
show route table mpls.0 protocol ospf on page 655
show route table mpls.0 extensive (PTX Series) on page 656
show route table mpls.0 (RSVP Route—Transit LSP) on page 656
show route table vpls_1 detail on page 657
show route table vpn-a on page 657
show route table vpn-a.mdt.0 on page 657
show route table VPN-A detail on page 658
show route table VPN-AB.inet.0 on page 658
show route table VPN_blue.mvpn-inet6.0 on page 659
show route table vrf1.mvpn.0 extensive on page 659
show route table inetflow detail on page 659
show route table bgp.evpn.0 extensive |no-more (EVPN) on page 662

Output Fields Table 15 on page 261 describes the output fields for the **show route table** command. Output fields are listed in the approximate order in which they appear.

Table 47: show route table Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).

Restart complete	All protocols have restarted for this routing table.
	
	Restart state:
	
	• **Pending:*protocol-name***—List of protocols that have not yet completed graceful restart for this routing table.
	• **Complete**—All protocols have restarted for this routing table.
	
	For example, if the output shows-
	
	• `LDP.inet.0 : 5 routes (4 active, 1 holddown, 0 hidden)`
	`Restart Pending: OSPF LDP VPN`
	
	This indicates that **OSPF**, **LDP**, and **VPN** protocols did not restart for the **LDP.inet.0** routing table.
	• `vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)`
	`Restart Complete`
	
	This indicates that all protocols have restarted for the **vpls_1.l2vpn.0** routing table.

number destinations	Number of destinations for which there are routes in the routing table.

Copyright © 2017, Juniper Networks, Inc.

Table 47: show route table Output Fields *(continued)*

Field Name	Field Description
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)

route-destination (entry, announced)	Route destination (for example:10.0.0.1/24). The **entry** value is the number of routes for this destination, and the **announced** value is the number of routes being announced for this destination. Sometimes the route destination is presented in another format, such as:
	• ***MPLS-label*** (for example, 80001).
	• ***interface-name*** (for example, ge-1/0/2).
	• ***neighbor-address*:*control-word-status*:*encapsulation type*:*vc-id*:*source*** (Layer 2 circuit only; for example, 10.1.1.195:NoCtrlWord:1:1:Local/96).
	  ◦ ***neighbor-address***—Address of the neighbor.
	  ◦ ***control-word-status***—Whether the use of the control word has been negotiated for this virtual circuit: **NoCtrlWord** or **CtrlWord**.
	  ◦ ***encapsulation type***—Type of encapsulation, represented by a number: (1) Frame Relay DLCI, (2) ATM AAL5 VCC transport, (3) ATM transparent cell transport, (4) Ethernet, (5) VLAN Ethernet, (6) HDLC, (7) PPP, (8) ATM VCC cell transport, (10) ATM VPC cell transport.
	  ◦ ***vc-id***—Virtual circuit identifier.
	  ◦ ***source***—Source of the advertisement: **Local** or **Remote**.
	• ***inclusive multicast Ethernet tag route***—Type of route destination represented by (for example, 3:100.100.100.10:100::0::10::100.100.100.10/384):
	  ◦ *route distinguisher*—(8 octets) Route distinguisher (RD) must be the RD of the EVPN instance (EVI) that is advertising the NLRI.
	  ◦ *Ethernet tag ID*—(4 octets) Identifier of the Ethernet tag. Can set to 0 or to a valid Ethernet tag value.
	  ◦ *IP address length*—(1 octet) Length of IP address in bits.
	  ◦ *originating router's IP address*—(4 or 16 octets) Must set to the provider edge (PE) device's IP address. This address should be common for all EVIs on the PE device, and may be the PE device's loopback address.

label stacking	**(**Next-to-the-last-hop routing device for MPLS only) Depth of the MPLS label stack, where the label-popping operation is needed to remove one or more labels from the top of the stack. A pair of routes is displayed, because the pop operation is performed only when the stack depth is two or more labels. • **S=0 route** indicates that a packet with an incoming label stack depth of 2 or more exits this routing device with one fewer label (the label-popping operation is performed). • If there is no **S=** information, the route is a normal MPLS route, which has a stack depth of 1 (the label-popping operation is not performed).

Copyright © 2017, Juniper Networks, Inc.

Table 47: show route table Output Fields *(continued)*

Field Name	Field Description
[*protocol*, *preference*]	Protocol from which the route was learned and the preference value for the route. • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • ***—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.
Level	(IS-IS only). In IS-IS, a single AS can be divided into smaller groups called areas. Routing between areas is organized hierarchically, allowing a domain to be administratively divided into smaller areas. This organization is accomplished by configuring Level 1 and Level 2 intermediate systems. Level 1 systems route within an area. When the destination is outside an area, they route toward a Level 2 system. Level 2 intermediate systems route between areas and toward other ASs.
Route Distinguisher	IP subnet augmented with a 64-bit prefix.
PMSI	Provider multicast service interface (MVPN routing table).
Next-hop type	Type of next hop. For a description of possible values for this field, see Table 18 on page 305.
Next-hop reference count	Number of references made to the next hop.
Flood nexthop branches exceed maximum message	Indicates that the number of flood next-hop branches exceeded the system limit of 32 branches, and only a subset of the flood next-hop branches were installed in the kernel.
Source	IP address of the route source.
Next hop	Network layer address of the directly reachable neighboring system.
via	Interface used to reach the next hop. If there is more than one interface available to the next hop, the name of the interface that is actually used is followed by the word **Selected**. This field can also contain the following information: • Weight—Value used to distinguish primary, secondary, and fast reroute backup routes. Weight information is available when MPLS label-switched path (LSP) link protection, node-link protection, or fast reroute is enabled, or when the standby state is enabled for secondary paths. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible. • Balance—Balance coefficient indicating how traffic of unequal cost is distributed among next hops when a routing device is performing unequal-cost load balancing. This information is available when you enable BGP multipath load balancing.
Label-switched-path *lsp-path-name*	Name of the LSP used to reach the next hop.

Copyright © 2017, Juniper Networks, Inc.

Table 47: show route table Output Fields *(continued)*

Field Name	Field Description
Label operation	MPLS label and operation occurring at this routing device. The operation can be **pop** (where a label is removed from the top of the stack), **push** (where another label is added to the label stack), or **swap** (where a label is replaced by another label).
Interface	(Local only) Local interface name.
Protocol next hop	Network layer address of the remote routing device that advertised the prefix. This address is used to derive a forwarding next hop.
Indirect next hop	Index designation used to specify the mapping between protocol next hops, tags, kernel export policy, and the forwarding next hops.
State	State of the route (a route can be in more than one state). See Table 19 on page 307.
Local AS	AS number of the local routing devices.
Age	How long the route has been known.
AIGP	Accumulated interior gateway protocol (AIGP) BGP attribute.
Metric*n*	Cost value of the indicated route. For routes within an AS, the cost is determined by IGP and the individual protocol metrics. For external routes, destinations, or routing domains, the cost is determined by a preference value.
MED-plus-IGP	Metric value for BGP path selection to which the IGP cost to the next-hop destination has been added.
TTL-Action	For MPLS LSPs, state of the TTL propagation attribute. Can be enabled or disabled for all RSVP-signaled and LDP-signaled LSPs or for specific VRF routing instances.
Task	Name of the protocol that has added the route.
Announcement bits	The number of BGP peers or protocols to which Junos OS has announced this route, followed by the list of the recipients of the announcement. Junos OS can also announce the route to the kernel routing table (KRT) for installing the route into the Packet Forwarding Engine, to a resolve tree, a Layer 2 VC, or even a VPN. For example, ***n*-Resolve inet** indicates that the specified route is used for route resolution for next hops found in the routing table. • ***n***—An index used by Juniper Networks customer support only.

Copyright © 2017, Juniper Networks, Inc.

Table 47: show route table Output Fields *(continued)*

Field Name	Field Description
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **Recorded**—The AS path is recorded by the sample process (sampled). • **?**—Incomplete; typically, the AS path was aggregated. When AS path numbers are included in the route, the format is as follows: • **[ ]**—Brackets enclose the number that precedes the AS path. This number represents the number of ASs present in the AS path, when calculated as defined in RFC 4271. This value is used in the AS-path merge process, as defined in RFC 4893. • **[ ]**—If more than one AS number is configured on the routing device, or if AS path prepending is configured, brackets enclose the local AS number associated with the AS path. • **{ }**—Braces enclose AS sets, which are groups of AS numbers in which the order does not matter. A set commonly results from route aggregation. The numbers in each AS set are displayed in ascending order. • **( )**—Parentheses enclose a confederation. • **([ ])**—Parentheses and brackets enclose a confederation set. NOTE: In Junos OS Release 10.3 and later, the AS path field displays an unrecognized attribute and associated hexadecimal value if BGP receives attribute 128 (attribute set) and you have not configured an independent domain in any routing instance.
validation-state	(BGP-learned routes) Validation status of the route: • **Invalid**—Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **Unknown**—Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **Unverified**—Indicates that the origin of the prefix is not verified against the database. This is because the database got populated and the validation is not called for in the BGP import policy, although origin validation is enabled, or the origin validation is not enabled for the BGP peers. • **Valid**—Indicates that the prefix and autonomous system pair are found in the database.
FECs bound to route	Indicates point-to-multipoint root address, multicast source address, and multicast group address when multipoint LDP (M-LDP) inband signaling is configured.
Primary Upstream	When multipoint LDP with multicast-only fast reroute (MoFRR) is configured, indicates the primary upstream path. MoFRR transmits a multicast join message from a receiver toward a source on a primary path, while also transmitting a secondary multicast join message from the receiver toward the source on a backup path.
RPF Nexthops	When multipoint LDP with MoFRR is configured, indicates the reverse-path forwarding (RPF) next-hop information. Data packets are received from both the primary path and the secondary paths. The redundant packets are discarded at topology merge points due to the RPF checks.
Label	Multiple MPLS labels are used to control MoFRR stream selection. Each label represents a separate route, but each references the same interface list check. Only the primary label is forwarded while all others are dropped. Multiple interfaces can receive packets using the same label.

Copyright © 2017, Juniper Networks, Inc.

Table 47: show route table Output Fields *(continued)*

Field Name	Field Description
weight	Value used to distinguish MoFRR primary and backup routes. A lower weight value is preferred. Among routes with the same weight value, load balancing is possible.
VC Label	MPLS label assigned to the Layer 2 circuit virtual connection.
MTU	Maximum transmission unit (MTU) of the Layer 2 circuit.
VLAN ID	VLAN identifier of the Layer 2 circuit.
Prefixes bound to route	Forwarding equivalent class (FEC) bound to this route. Applicable only to routes installed by LDP.
Communities	Community path attribute for the route. See Table 20 on page 309 for all possible values for this field.
Layer2-info: encaps	Layer 2 encapsulation (for example, VPLS).
control flags	Control flags: **none** or **Site Down**.
mtu	Maximum transmission unit (MTU) information.
Label-Base, range	First label in a block of labels and label block size. A remote PE routing device uses this first label when sending traffic toward the advertising PE routing device.
status vector	Layer 2 VPN and VPLS network layer reachability information (NLRI).
Accepted Multipath	Current active path when BGP multipath is configured.
Accepted LongLivedStale	The LongLivedStale flag indicates that the route was marked LLGR-stale by this router, as part of the operation of LLGR receiver mode. Either this flag or the LongLivedStaleImport flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag.
Accepted LongLivedStaleImport	The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy. Either this flag or the LongLivedStale flag might be displayed for a route. Neither of these flags is displayed at the same time as the Stale (ordinary GR stale) flag. Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and import into the inet.0 routing table
ImportAccepted LongLivedStaleImport	Accept all received BGP long-lived graceful restart (LLGR) and LLGR stale routes learned from configured neighbors and imported into the inet.0 routing table The LongLivedStaleImport flag indicates that the route was marked LLGR-stale when it was received from a peer, or by import policy.
Accepted MultipathContrib	Path currently contributing to BGP multipath.
Localpref	Local preference value included in the route.
Router ID	BGP router ID as advertised by the neighbor in the open message.

Copyright © 2017, Juniper Networks, Inc.

Table 47: show route table Output Fields *(continued)*

Field Name	Field Description
Primary Routing Table	In a routing table group, the name of the primary routing table in which the route resides.
Secondary Tables	In a routing table group, the name of one or more secondary tables in which the route resides.

Table 18 on page 305 describes all possible values for the Next-hop Types output field.

Table 48: Next-hop Types Output Field Values

Next-Hop Type	Description
Broadcast (bcast)	Broadcast next hop.
Deny	Deny next hop.
Discard	Discard next hop.
Flood	Flood next hop. Consists of components called branches, up to a maximum of 32 branches. Each flood next-hop branch sends a copy of the traffic to the forwarding interface. Used by point-to-multipoint RSVP, point-to-multipoint LDP, point-to-multipoint CCC, and multicast.
Hold	Next hop is waiting to be resolved into a unicast or multicast type.
Indexed (idxd)	Indexed next hop.
Indirect (indr)	Used with applications that have a protocol next hop address that is remote. You are likely to see this next-hop type for internal BGP (IBGP) routes when the BGP next hop is a BGP neighbor that is not directly connected.
Interface	Used for a network address assigned to an interface. Unlike the router next hop, the interface next hop does not reference any specific node on the network.
Local (locl)	Local address on an interface. This next-hop type causes packets with this destination address to be received locally.
Multicast (mcst)	Wire multicast next hop (limited to the LAN).
Multicast discard (mdsc)	Multicast discard.
Multicast group (mgrp)	Multicast group member.
Receive (recv)	Receive.
Reject (rjct)	Discard. An ICMP unreachable message was sent.

Copyright © 2017, Juniper Networks, Inc.

Table 48: Next-hop Types Output Field Values *(continued)*

Next-Hop Type	Description
Resolve (rslv)	Resolving next hop.
Routed multicast (mcrt)	Regular multicast next hop.
Router	A specific node or set of nodes to which the routing device forwards packets that match the route prefix. To qualify as a next-hop type router, the route must meet the following criteria: • Must not be a direct or local subnet for the routing device. • Must have a next hop that is directly connected to the routing device.
Table	Routing table next hop.
Unicast (ucst)	Unicast.
Unilist (ulst)	List of unicast next hops. A packet sent to this next hop goes to any next hop in the list.

Table 19 on page 307 describes all possible values for the State output field. A route can be in more than one state (for example, **<Active NoReadvrt Int Ext>**).

Table 49: State Output Field Values

Value	Description
Accounting	Route needs accounting.
Active	Route is active.
Always Compare MED	Path with a lower multiple exit discriminator (MED) is available.
AS path	Shorter AS path is available.
Cisco Non-deterministic MED selection	Cisco nondeterministic MED is enabled, and a path with a lower MED is available.
Clone	Route is a clone.
Cluster list length	Length of cluster list sent by the route reflector.
Delete	Route has been deleted.
Ex	Exterior route.

Copyright © 2017, Juniper Networks, Inc.

Table 49: State Output Field Values *(continued)*

Value	Description
Ext	BGP route received from an external BGP neighbor.
FlashAll	Forces all protocols to be notified of a change to any route, active or inactive, for a prefix. When not set, protocols are informed of a prefix only when the active route changes.
Hidden	Route not used because of routing policy.
IfCheck	Route needs forwarding RPF check.
IGP metric	Path through next hop with lower IGP metric is available.
Inactive reason	Flags for this route, which was not selected as best for a particular destination.
Initial	Route being added.
Int	Interior route.
Int Ext	BGP route received from an internal BGP peer or a BGP confederation peer.
Interior > Exterior > Exterior via Interior	Direct, static, IGP, or EBGP path is available.
Local Preference	Path with a higher local preference value is available.
Martian	Route is a martian (ignored because it is obviously invalid).
MartianOK	Route exempt from martian filtering.
Next hop address	Path with lower metric next hop is available.
No difference	Path from neighbor with lower IP address is available.
NoReadvrt	Route not to be advertised.
NotBest	Route not chosen because it does not have the lowest MED.
Not Best in its group	Incoming BGP AS is not the best of a group (only one AS can be the best).
NotInstall	Route not to be installed in the forwarding table.
Number of gateways	Path with a greater number of next hops is available.
Origin	Path with a lower origin code is available.

Copyright © 2017, Juniper Networks, Inc.

Table 49: State Output Field Values *(continued)*

Value	Description
Pending	Route pending because of a hold-down configured on another route.
Release	Route scheduled for release.
RIB preference	Route from a higher-numbered routing table is available.
Route Distinguisher	64-bit prefix added to IP subnets to make them unique.
Route Metric or MED comparison	Route with a lower metric or MED is available.
Route Preference	Route with lower preference value is available.
Router ID	Path through a neighbor with lower ID is available.
Secondary	Route not a primary route.
Unusable path	Path is not usable because of one of the following conditions: • The route is damped. • The route is rejected by an import policy. • The route is unresolved.
Update source	Last tiebreaker is the lowest IP address value.

Table 20 on page 309 describes the possible values for the Communities output field.

Table 50: Communities Output Field Values

Value	Description
area-number	4 bytes, encoding a 32-bit area number. For AS-external routes, the value is **0**. A nonzero value identifies the route as internal to the OSPF domain, and as within the identified area. Area numbers are relative to a particular OSPF domain.
bandwidth: ***local AS number:link-bandwidth-number***	Link-bandwidth community value used for unequal-cost load balancing. When BGP has several candidate paths available for multipath purposes, it does not perform unequal-cost load balancing according to the link-bandwidth community unless all candidate paths have this attribute.
domain-id	Unique configurable number that identifies the OSPF domain.
domain-id-vendor	Unique configurable number that further identifies the OSPF domain.
link-bandwidth-number	Link-bandwidth number: from **0** through **4,294,967,295** (bytes per second).
local AS number	Local AS number: from **1** through **65,535**.

Copyright © 2017, Juniper Networks, Inc.

Table 50: Communities Output Field Values *(continued)*

Value	Description
options	1 byte. Currently this is only used if the route type is **5** or **7**. Setting the least significant bit in the field indicates that the route carries a type 2 metric.
origin	(Used with VPNs) Identifies where the route came from.
ospf-route-type	1 byte, encoded as **1** or **2** for intra-area routes (depending on whether the route came from a type 1 or a type 2 LSA); **3** for summary routes; **5** for external routes (area number must be**0**); **7** for NSSA routes; or **129** for sham link endpoint addresses.
route-type-vendor	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x8000**. The format is *area-number:ospf-route-type:options*.
rte-type	Displays the area number, OSPF route type, and option of the route. This is configured using the BGP extended community attribute **0x0306**. The format is *area-number:ospf-route-type:options*.
target	Defines which VPN the route participates in; **target** has the format ***32-bit IP address:16-bit number***. For example, 10.19.0.0:100.
unknown IANA	Incoming IANA codes with a value between **0x1** and **0x7fff**. This code of the BGP extended community attribute is accepted, but it is not recognized.
unknown OSPF vendor community	Incoming IANA codes with a value above **0x8000**. This code of the BGP extended community attribute is accepted, but it is not recognized.

Sample Output

show route table bgp.l2.vpn

```
user@host> show route table bgp.l2.vpn
bgp.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

192.168.24.1:1:4:1/96
                   *[BGP/170] 01:08:58, localpref 100, from 192.168.24.1
                      AS path: I
                    > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table bgp.l3vpn.0

```
user@host> show route table bgp.l3vpn.0
      bgp.l3vpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.255.71.15:100:10.255.71.17/32
                         *[BGP/170] 00:03:59, MED 1, localpref 100, from
      10.255.71.15
                            AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.15:200:10.255.71.18/32
```

Copyright © 2017, Juniper Networks, Inc.

```
                                *[BGP/170] 00:03:59, MED 1, localpref 100, from
      10.255.71.15
                                  AS path: I
                                > via so-2/1/0.0, Push 100021, Push 100011(top)
```

show route table bgp.l3vpn.0 detail

```
user@host> show route table bgp.l3vpn.0 detail
bgp.l3vpn.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)

10.255.245.12:1:172.16.4.0/8 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
         Route Distinguisher: 10.255.245.12:1
         Source: 10.255.245.12
         Next hop: 192.168.208.66 via fe-0/0/0.0, selected
         Label operation: Push 182449
         Protocol next hop: 10.255.245.12
         Push 182449
         Indirect next hop: 863a630 297
         State: <Active Int Ext>
         Local AS:    35 Peer AS:    35
         Age: 12:19      Metric2: 1
         Task: BGP_35.10.255.245.12+179
         Announcement bits (1): 0-BGP.0.0.0.0+179
         AS path: 30 10458 14203 2914 3356 I (Atomic) Aggregator: 3356 4.68.0.11

         Communities: 2914:420 target:11111:1 origin:56:78
         VPN Label: 182449
         Localpref: 100
          Router ID: 10.255.245.12

10.255.245.12:1:4.17.225.0/24 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
         Route Distinguisher: 10.255.245.12:1
         Source: 10.255.245.12
         Next hop: 192.168.208.66 via fe-0/0/0.0, selected
         Label operation: Push 182465
         Protocol next hop: 10.255.245.12
         Push 182465
         Indirect next hop: 863a8f0 305
         State: <Active Int Ext>
         Local AS:    35 Peer AS:    35
         Age: 12:19      Metric2: 1
         Task: BGP_35.10.255.245.12+179
         Announcement bits (1): 0-BGP.0.0.0.0+179
AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496 6496 I
                Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                VPN Label: 182465
                Localpref: 100
                Router ID: 10.255.245.12

10.255.245.12:1:4.17.226.0/23 (1 entry, 1 announced)
  *BGP   Preference: 170/-101
         Route Distinguisher: 10.255.245.12:1
         Source: 10.255.245.12
         Next hop: 192.168.208.66 via fe-0/0/0.0, selected
         Label operation: Push 182465
         Protocol next hop: 10.255.245.12
         Push 182465
         Indirect next hop: 86bd210 330
         State: <Active Int Ext>
```


Copyright © 2017, Juniper Networks, Inc.

```
                Local AS:    35 Peer AS:    35
                Age: 12:19      Metric2: 1
                Task: BGP_35.10.255.245.12+179
                Announcement bits (1): 0-BGP.0.0.0.0+179
                AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

                6496 I
                Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                VPN Label: 182465
                Localpref: 100
                Router ID: 10.255.245.12

        10.255.245.12:1:4.17.251.0/24 (1 entry, 1 announced)
          *BGP   Preference: 170/-101
                Route Distinguisher: 10.255.245.12:1
                Source: 10.255.245.12
                Next hop: 192.168.208.66 via fe-0/0/0.0, selected
                Label operation: Push 182465
                Protocol next hop: 10.255.245.12
                Push 182465
                Indirect next hop: 86bd210 330
                State: <Active Int Ext>
                Local AS:    35 Peer AS:    35
                Age: 12:19      Metric2: 1
                Task: BGP_35.10.255.245.12+179
                Announcement bits (1): 0-BGP.0.0.0.0+179
                AS path: 30 10458 14203 2914 11853 11853 11853 6496 6496 6496 6496 6496

                6496 I
                Communities: 2914:410 target:12:34 target:11111:1 origin:12:34
                VPN Label: 182465
                Localpref: 100
```

show route table bgp.rtarget.0 (When Proxy BGP Route Target Filtering Is Configured)

```
user@host> show route table bgp.rtarget.o
bgp.rtarget.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

100:100:100/96
                   *[RTarget/5] 00:03:14
                      Type Proxy
                        for 10.255.165.103
                        for 10.255.166.124
                      Local
```

show route table bgp.evpn.0

```
user@host> show route table bgp.evpn.0
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

2:100.100.100.2:100::0::00:26:88:5f:67:b0/304
                   *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
2:100.100.100.2:100::0::00:51:51:51:51:51/304
                   *[BGP/170] 11:00:05, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
```

Copyright © 2017, Juniper Networks, Inc.

```
2:100.100.100.3:100::0::00:52:52:52:52:52/304
                   *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                      AS path: I, validation-state: unverified
                    > to 100.64.13.3 via ge-2/0/8.0, label-switched-path R0toR2
2:100.100.100.3:100::0::a8:d0:e5:5b:01:c8/304
                   *[BGP/170] 10:59:58, localpref 100, from 100.100.100.3
                      AS path: I, validation-state: unverified
                    > to 100.64.13.3 via ge-2/0/8.0, label-switched-path R0toR2
3:100.100.100.2:100::1000::100.100.100.2/304
                   *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
3:100.100.100.2:100::2000::100.100.100.2/304
                   *[BGP/170] 11:00:16, localpref 100, from 100.100.100.2
                      AS path: I, validation-state: unverified
                    > to 100.64.12.2 via xe-2/2/0.0, label-switched-path R0toR1
```

show route table evpna.evpn.0

```
user@host> show route table evpna.evpn.0
evpna.evpn.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

3:100.100.100.10:100::0::10::100.100.100.10/384
                   *[EVPN/170] 01:37:09
                      Indirect
3:100.100.100.2:100::2000::100.100.100.2/304
                   *[EVPN/170] 01:37:12
                      Indirect
```

show route table inet.0

```
user@host> show route table inet.0
inet.0: 12 destinations, 12 routes (11 active, 0 holddown, 1 hidden)
+ = Active Route, - = Last Active, * = Both

0.0.0.0/0          *[Static/5] 00:51:57
                    > to 172.16.5.254 via fxp0.0
10.0.0.1/32         *[Direct/0] 00:51:58
                    > via at-5/3/0.0
10.0.0.2/32         *[Local/0] 00:51:58
                      Local
10.12.12.21/32     *[Local/0] 00:51:57
                      Reject
10.13.13.13/32     *[Direct/0] 00:51:58
                    > via t3-5/2/1.0
10.13.13.14/32     *[Local/0] 00:51:58
                      Local
10.13.13.21/32     *[Local/0] 00:51:58
                      Local
10.13.13.22/32     *[Direct/0] 00:33:59
                    > via t3-5/2/0.0
127.0.0.1/32        [Direct/0] 00:51:58
                    > via lo0.0
10.222.5.0/24     *[Direct/0] 00:51:58
                    > via fxp0.0
10.222.5.81/32    *[Local/0] 00:51:58
                      Local
```


Copyright © 2017, Juniper Networks, Inc.

show route table inet.3

```
user@host> show route table inet.3
inet.3: 5 destinations, 5 routes (5 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.0.0.5/32        *[LDP/9] 00:25:43, metric 10, tag 200
                      to 10.2.94.2 via lt-1/2/0.49
                    > to 10.2.3.2 via lt-1/2/0.23
```

show route table inet.3 protocol ospf

```
user@host> show route table inet.3 protocol ospf
inet.3: 9 destinations, 18 routes (9 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1.1.1.20/32         [L-OSPF/10] 1d 00:00:56, metric 2
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800020
                     to 10.0.6.60 via lt-1/2/0.12, Push 800020, Push 800030(top)
1.1.1.30/32         [L-OSPF/10] 1d 00:01:01, metric 3
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800030
                      to 10.0.6.60 via lt-1/2/0.12, Push 800030
1.1.1.40/32         [L-OSPF/10] 1d 00:01:01, metric 4
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800040
                      to 10.0.6.60 via lt-1/2/0.12, Push 800040
1.1.1.50/32         [L-OSPF/10] 1d 00:01:01, metric 5
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800050
                      to 10.0.6.60 via lt-1/2/0.12, Push 800050
1.1.1.60/32         [L-OSPF/10] 1d 00:01:01, metric 6
                    > to 10.0.10.70 via lt-1/2/0.14, Push 800060
                      to 10.0.6.60 via lt-1/2/0.12, Pop
```

show route table inet6.0

```
user@host> show route table inet6.0
inet6.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Route, * = Both

fec0:0:0:3::/64 *[Direct/0] 00:01:34
>via fe-0/1/0.0

fec0:0:0:3::/128 *[Local/0] 00:01:34
>Local

fec0:0:0:4::/64 *[Static/5] 00:01:34
>to fec0:0:0:3::ffff via fe-0/1/0.0
```

show route table inet6.3

```
user@router> show route table inet6.3
inet6.3: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

::10.255.245.195/128
                   *[LDP/9] 00:00:22, metric 1
                    > via so-1/0/0.0
::10.255.245.196/128
                   *[LDP/9] 00:00:08, metric 1
                    > via so-1/0/0.0, Push 100008
```

Copyright © 2017, Juniper Networks, Inc.

show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 2
                State: <Active Ext>
                Local AS: 64502 Peer AS: 64500
                Age: 4
                Task: BGP_64500.10.12.99.5+3792
                Announcement bits (1): 0-Flow
                AS path: 64500 I
                Communities: traffic-rate:0:0
                Validation state: Accept, Originator: 10.12.99.5
                Via: 10.12.44.0/24, Active
                Localpref: 100
                Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
        *Flow   Preference: 5
                Next-hop reference count: 2
                State: <Active>
                Local AS: 64502
                Age: 6:30
                Task: RT Flow
                Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
                AS path: I
                Communities: 1:1
```

show route table lsdist.0 extensive

```
user@host> show route table lsdist.0 extensive
lsdist.0: 10 destinations, 10 routes (10 active, 0 holddown, 0 hidden)
NODE { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 ISIS-L1:0 }/1152
 (1 entry, 1 announced)
TSI:
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xa62f378 (adv_entry)
   Advertised metrics:
     Nexthop: Self
     Localpref: 100
     AS path: [4170512532] I
     Communities:
Path NODE { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 ISIS-L1:0 }
 Vector len 4.  Val: 0
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0x95dfc64
                Next-hop reference count: 9
                State: <Active NotInstall>
                Local AS: 4170512532
                Age: 6:05
                Validation State: unverified
                Task: IS-IS
                Announcement bits (1): 0-BGP_RT_Background
                AS path: I
                IPv4 Router-ids:
                   128.220.11.197
                Area membership:
```


Copyright © 2017, Juniper Networks, Inc.

```
                47 00 05 80 ff f8 00 00 00 01 08 00 01
             SPRING-Capabilities:                 - SRGB block [Start: 800000,
 Range: 256, Flags: 0xc0]
              SPRING-Algorithms:
                 - Algo: 0
 LINK { Local { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 }.{
IPv4:8.65.1.105 } Remote { AS:4170512532 BGP-LS ID:4170512532 ISO:4284.3300.5067)
TSI:
Page 0 idx 0, (group ibgp type Internal) Type 1 val 0xa62f3cc (adv_entry)
    Advertised metrics:
      Nexthop: Self
      Localpref: 100
      AS path: [4170512532] I
      Communities:
Path LINK { Local { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 }.{
 IPv4:8.65.1.105 } Remote { AS:4170512532 BGP-LS ID:4170512532 ISO:4284.33000
      *IS-IS  Preference: 15
               Level: 1
               Next hop type: Fictitious, Next hop index: 0
               Address: 0x95dfc64
               Next-hop reference count: 9
               State: <Active NotInstall>
               Local AS: 4170512532
               Age: 6:05
               Validation State: unverified
               Task: IS-IS
               Announcement bits (1): 0-BGP_RT_Background
               AS path: I
               Color: 32768
               Maximum bandwidth: 1000Mbps
               Reservable bandwidth: 1000Mbps
               Unreserved bandwidth by priority:
                  0   1000Mbps
                  1   1000Mbps
                  2   1000Mbps
                  3   1000Mbps
                  4   1000Mbps
                  5   1000Mbps
                  6   1000Mbps
                  7   1000Mbps
                Metric: 10
                TE Metric: 10
                LAN IPV4 Adj-SID -  Label: 299776, Flags: 0x30,
Weight: 0, Nbr: 10.220.1.83

PREFIX { Node { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 } {
IPv4:128.220.11.197/32 } ISIS-L1:0 }/1152 (1 entry, 1 announced) TSI: Page 0 idx
 0, (group ibgp type Internal) Type 1 val 0xa62f43c (adv_entry)
    Advertised metrics:
      Nexthop: Self
      Localpref: 100
      AS path: [4170512532] I
      Communities:
Path PREFIX { Node { AS:4170512532 BGP-LS ID:4170512532 ISO:3245.3412.3456.00 }
{ IPv4:128.220.11.197/32 } ISIS-L1:0 } Vector len 4.  Val: 0
        *IS-IS  Preference: 15
                Level: 1
                Next hop type: Fictitious, Next hop index: 0
                Address: 0x95dfc64
                Next-hop reference count: 9
                State:<Active NotInstall>
```

```
                    Local AS: 4170512532
                    Age: 6:05
                    Validation State: unverified
                    Task: IS-IS
                    Announcement bits (1): 0-BGP_RT_Background
                    AS path: I
                                 Prefix SID: 67, Flags: 0x40, Algo: 0
```

show route table l2circuit.0

```
user@host> show route table l2circuit.0
l2circuit.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.195:NoCtrlWord:1:1:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:NoCtrlWord:1:1:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
10.1.1.195:CtrlWord:1:2:Local/96
                   *[L2CKT/7] 00:50:47
                    > via so-0/1/2.0, Push 100049
                      via so-0/1/3.0, Push 100049
10.1.1.195:CtrlWord:1:2:Remote/96
                   *[LDP/9] 00:50:14
                      Discard
```

show route table mpls

```
user@host> show route table mpls
mpls.0: 4 destinations, 4 routes (4 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1                  *[MPLS/0] 00:13:55, metric 1
                      Receive
2                  *[MPLS/0] 00:13:55, metric 1
                      Receive
1024               *[VPN/0] 00:04:18
                      to table red.inet.0, Pop
```

show route table mpls extensive

```
user@host> show route table mpls extensive
100000 (1 entry, 1 announced)
TSI:
KRT in-kernel 100000 /36 -> {so-1/0/0.0}
        *LDP    Preference:   9
                Next hop:  via so-1/0/0.0, selected
                Pop
                State: <Active Int>
                Age: 29:50      Metric: 1
                Task: LDP
                Announcement bits (1): 0-KRT
                AS path: I
                Prefixes bound to route: 10.0.0.194/32
```


Copyright © 2017, Juniper Networks, Inc.

show route table mpls.0

```
user@host> show route table mpls.0
mpls.0: 18 destinations, 19 routes (18 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 11:39:56, metric 1
                      to table inet.0
0(S=0)             *[MPLS/0] 11:39:56, metric 1
                      to table mpls.0
1                  *[MPLS/0] 11:39:56, metric 1
                      Receive
2                  *[MPLS/0] 11:39:56, metric 1
                      to table inet6.0
2(S=0)             *[MPLS/0] 11:39:56, metric 1
                      to table mpls.0
13                 *[MPLS/0] 11:39:56, metric 1
                      Receive
303168             *[EVPN/7] 11:00:49, routing-instance pbbn10, route-type
Ingress-MAC, ISID 0
                      to table pbbn10.evpn-mac.0
303184             *[EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 1000
                      to table pbbn10.evpn-mac.0
                    [EVPN/7] 11:00:53, routing-instance pbbn10, route-type
Ingress-IM, ISID 2000
                      to table pbbn10.evpn-mac.0
303264             *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-IM, ISID 1000
                     > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303280             *[EVPN/7] 11:00:53, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-IM, ISID 2000
                     > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303328             *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                     > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303344             *[EVPN/7] 11:00:49, remote-pe 100.100.100.2, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                     > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303360             *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:26:88:5f:67:b0
                     > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303376             *[EVPN/7] 11:00:47, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:51:51:51:51:51
                     > to 100.1.12.2 via xe-2/2/0.0, label-switched-path R0toR1
303392             *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                     > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303408             *[EVPN/7] 11:00:35, remote-pe 100.100.100.3, routing-instance
 pbbn10, route-type Egress-MAC, ISID 0
                     > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303424             *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC a8:d0:e5:5b:01:c8
                     > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
303440             *[EVPN/7] 11:00:33, routing-instance pbbn10, route-type
Egress-MAC, ISID 0, BMAC 00:52:52:52:52:52
                     > to 100.1.13.3 via ge-2/0/8.0, label-switched-path R0toR2
```

Copyright © 2017, Juniper Networks, Inc.

show route table mpls.0 detail (PTX Series)

```
user@host> show route table mpls.0 detail
ge-0/0/2.600 (1 entry, 1 announced)
        *L2VPN  Preference: 7
                Next hop type: Indirect
                Address: 0x9438f34
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 567
                Next hop: 10.0.0.1 via ge-0/0/1.0, selected
                Label operation: Push 299808
                Label TTL action: prop-ttl
                Load balance label: Label 299808:None;
                Session Id: 0x1
                Protocol next hop: 10.255.255.1
                Label operation: Push 299872 Offset: 252
                Label TTL action: no-prop-ttl
                Load balance label: Label 299872:Flow label PUSH;
                Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                State: <Active Int>
                Age: 21         Metric2: 1
                Validation State: unverified
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 2-Common L2 VC
                AS path: I
```

show route table mpls.0 ccc ge-0/0/1.1004 detail

```
user@host>show route table mpls.0 ccc ge-0/0/1.1004 detail
mpls.0: 121 destinations, 121 routes (121 active, 0 holddown, 0 hidden)
ge-0/0/1.1004 (1 entry, 1 announced)
        *EVPN   Preference: 7
                Next hop type: List, Next hop index: 1048577
                Address: 0xdc14770
                Next-hop reference count: 3
                Next hop: ELNH Address 0xd011e30
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0xd011e30
                    Next-hop reference count: 3
                    Protocol next hop: 100.100.100.1
                    Label operation: Push 301952
                    Composite next hop: 0xd011dc0 754 INH Session ID: 0x146
                    Indirect next hop: 0xb69a890 1048615 INH Session ID: 0x146
                        Next hop type: Router, Next hop index: 735
                        Address: 0xd00e530
                        Next-hop reference count: 23
                        Next hop: 100.46.1.2 via ge-0/0/5.0
                        Label-switched-path pe4_to_pe1
                        Label operation: Push 300320
                        Label TTL action: prop-ttl
                        Load balance label: Label 300320: None;
                        Label element ptr: 0xd00e580
                        Label parent element ptr: 0x0
                        Label element references: 18
                        Label element child references: 16
                        Label element lsp id: 5
                Next hop: ELNH Address 0xd012070
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0xd012070
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Next-hop reference count: 3
                    Protocol next hop: 100.100.100.2
                    Label operation: Push 301888
                    Composite next hop: 0xd012000 755 INH Session ID: 0x143
                    Indirect next hop: 0xb69a9a0 1048641 INH Session ID: 0x143
                        Next hop type: Router, Next hop index: 716
                        Address: 0xd00e710
                        Next-hop reference count: 23
                        Next hop: 100.46.1.2 via ge-0/0/5.0
                        Label-switched-path pe4_to_pe2
                        Label operation: Push 300304
                        Label TTL action: prop-ttl
                        Load balance label: Label 300304: None;
                        Label element ptr: 0xd00e760
                        Label parent element ptr: 0x0
                        Label element references: 15
                        Label element child references: 13
                        Label element lsp id: 6
                Next hop: ELNH Address 0xd0121f0, selected
                    Next hop type: Indirect, Next hop index: 0
                    Address: 0xd0121f0
                    Next-hop reference count: 3
                    Protocol next hop: 100.100.100.3
                    Label operation: Push 301984
                    Composite next hop: 0xd012180 756 INH Session ID: 0x145
                    Indirect next hop: 0xb69aab0 1048642 INH Session ID: 0x145
                        Next hop type: Router, Next hop index: 801
                        Address: 0xd010ed0
                        Next-hop reference count: 32
                        Next hop: 100.46.1.2 via ge-0/0/5.0
                        Label-switched-path pe4_to_pe3
                        Label operation: Push 300336
                        Label TTL action: prop-ttl
                        Load balance label: Label 300336: None;
                        Label element ptr: 0xd0108c0
                        Label parent element ptr: 0x0
                        Label element references: 22
                        Label element child references: 20
                        Label element lsp id: 7
                State: < Active Int >
                Age: 2:06:50
                Validation State: unverified
                Task: evpn global task
                Announcement bits (1): 1-KRT
                AS path: I
```

show route table mpls.0 protocol evpn

```
user@host>show route table mpls.0 protocol evpn
mpls.0: 121 destinations, 121 routes (121 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299872             *[EVPN/7] 02:30:58, routing-instance mhevpn, route-type
Ingress-IM, vlan-id 10
                      to table mhevpn.evpn-mac.0
300016             *[EVPN/7] 02:30:38, routing-instance VS-1, route-type
Ingress-IM, vlan-id 110
                      to table VS-1.evpn-mac.0
300032             *[EVPN/7] 02:30:38, routing-instance VS-1, route-type
Ingress-IM, vlan-id 120
                      to table VS-1.evpn-mac.0
```

Copyright © 2017, Juniper Networks, Inc.

```
300048             *[EVPN/7] 02:30:38, routing-instance VS-1, route-type
Ingress-IM, vlan-id 130
                      to table VS-1.evpn-mac.0
300064             *[EVPN/7] 02:30:38, routing-instance VS-2, route-type
Ingress-IM, vlan-id 210
                      to table VS-2.evpn-mac.0
300080             *[EVPN/7] 02:30:38, routing-instance VS-2, route-type
Ingress-IM, vlan-id 220
                      to table VS-2.evpn-mac.0
300096             *[EVPN/7] 02:30:38, routing-instance VS-2, route-type
Ingress-IM, vlan-id 230
                      to table VS-2.evpn-mac.0
300112             *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:44:44:44:44:44:44:44:44:44
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300128             *[EVPN/7] 02:29:22, routing-instance mhevpn, route-type
Ingress-Aliasing
                      to table mhevpn.evpn-mac.0
300144             *[EVPN/7] 02:27:06, routing-instance VS-1, route-type
Egress-MAC, ESI 00:44:44:44:44:44:44:44:44:44
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300160             *[EVPN/7] 02:29:22, routing-instance VS-1, route-type
Ingress-Aliasing
                      to table VS-1.evpn-mac.0
300176             *[EVPN/7] 02:27:07, routing-instance VS-2, route-type
Egress-MAC, ESI 00:44:44:44:44:44:44:44:44:44
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300192             *[EVPN/7] 02:29:22, routing-instance VS-2, route-type
Ingress-Aliasing
                      to table VS-2.evpn-mac.0
300208             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-IM, vlan-id 120
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300224             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 mhevpn, route-type Egress-IM, vlan-id 10
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300240             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-IM, vlan-id 110
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300256             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-IM, vlan-id 130
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300272             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-IM, vlan-id 210
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300288             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-IM, vlan-id 220
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300304             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-IM, vlan-id 230
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300320             *[EVPN/7] 02:27:06, routing-instance VS-1, route-type
Egress-MAC, ESI 00:11:11:11:11:11:11:11:11:11
                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300336             *[EVPN/7] 02:27:06, routing-instance VS-1, route-type
Egress-MAC, ESI 00:33:33:33:33:33:33:33:33:33
                      to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
```


Copyright © 2017, Juniper Networks, Inc.

```
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300368             *[EVPN/7] 02:27:07, routing-instance VS-2, route-type
Egress-MAC, ESI 00:33:33:33:33:33:33:33:33:33
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300384             *[EVPN/7] 02:27:07, routing-instance VS-2, route-type
Egress-MAC, ESI 00:11:11:11:11:11:11:11:11:11
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300416             *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:33:33:33:33:33:33:33:33:33
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300432             *[EVPN/7] 02:27:06, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:11:11:11:11:11:11:11:11:11
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
300480             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-MAC
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300496             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-MAC
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300560             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-1, route-type Egress-MAC
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300592             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 VS-2, route-type Egress-MAC
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
300608             *[EVPN/7] 02:29:23
                    > via ge-0/0/1.1001, Pop
300624             *[EVPN/7] 02:29:23
                    > via ge-0/0/1.2001, Pop
301232             *[EVPN/7] 02:29:17
                    > via ge-0/0/1.1002, Pop
301296             *[EVPN/7] 02:29:10
                    > via ge-0/0/1.1003, Pop
301312             *[EVPN/7] 02:27:06
                    > via ae10.2003, Pop
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301360             *[EVPN/7] 02:29:01
                    > via ge-0/0/1.1004, Pop
301408             *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 vpws1004, route-type Egress, vlan-id 2004
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301456             *[EVPN/7] 02:27:06
                    > via ae10.1010, Pop
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301552             *[EVPN/7] 02:27:07, routing-instance VS-1, route-type
Egress-MAC, ESI 00:22:22:22:22:22:22:22:22:22
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301568             *[EVPN/7] 02:27:07, routing-instance VS-2, route-type
```

```
Egress-MAC, ESI 00:22:22:22:22:22:22:22:22:22
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301648              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 vpws1010, route-type Egress, vlan-id 2010
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301664              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 mhevpn, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301680              *[EVPN/7] 02:27:07, remote-pe 100.100.100.2, routing-instance
 mhevpn, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
301696              *[EVPN/7] 02:27:07, routing-instance mhevpn, route-type
Egress-MAC, ESI 00:22:22:22:22:22:22:22:22:22
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301712              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301728              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301744              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-IM, vlan-id 230
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301760              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 vpws1010, route-type Egress, vlan-id 2010
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301776              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 mhevpn, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301792              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-IM, vlan-id 130
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301808              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 vpws1004, route-type Egress, vlan-id 2004
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301824              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 mhevpn, route-type Egress-IM, vlan-id 10
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301840              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 vpws1002, route-type Egress, vlan-id 2002
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301856              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 vpws1003, route-type Egress, vlan-id 2003
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301872              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 vpws1003, route-type Egress Protection, vlan-id 2003
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301888              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 vpws1010, route-type Egress Protection, vlan-id 1010
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301904              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-IM, vlan-id 220
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301920              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-IM, vlan-id 210
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
301936              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-IM, vlan-id 230
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301952              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-SH, vlan-id 230
```


Copyright © 2017, Juniper Networks, Inc.

```
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301968              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-IM, vlan-id 220
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
301984              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-SH, vlan-id 220
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302000              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-IM, vlan-id 210
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302016              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-SH, vlan-id 210
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302032              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302048              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-2, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302064              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302080              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-2, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302096              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302112              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302128              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302144              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-MAC
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302160              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-IM, vlan-id 120
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302176              *[EVPN/7] 02:27:07, remote-pe 100.100.100.1, routing-instance
 VS-1, route-type Egress-IM, vlan-id 110
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302192              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-IM, vlan-id 130
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302208              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-SH, vlan-id 130
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302224              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-IM, vlan-id 120
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302240              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-SH, vlan-id 120
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302256              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-IM, vlan-id 110
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302272              *[EVPN/7] 02:27:07, remote-pe 100.100.100.3, routing-instance
 VS-1, route-type Egress-SH, vlan-id 110
                    > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
```

Copyright © 2017, Juniper Networks, Inc.

```
302288             *[EVPN/7] 02:27:06, remote-pe 100.100.100.1, routing-instance
 mhevpn, route-type Egress-MAC
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302304             *[EVPN/7] 02:27:06, remote-pe 100.100.100.1, routing-instance
 mhevpn, route-type Egress-MAC
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302320             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 mhevpn, route-type Egress-MAC
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302336             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 mhevpn, route-type Egress-MAC
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302352             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 vpws1004, route-type Egress, vlan-id 2004
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302368             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 mhevpn, route-type Egress-IM, vlan-id 10
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302384             *[EVPN/7] 02:27:06, remote-pe 100.100.100.3, routing-instance
 mhevpn, route-type Egress-SH, vlan-id 10
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302400             *[EVPN/7] 02:26:21
                    > via ge-0/0/1.3001, Pop
302432             *[EVPN/7] 02:26:21, remote-pe 100.100.100.3, routing-instance
 vpws3001, route-type Egress, vlan-id 40000
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302448             *[EVPN/7] 02:26:21, remote-pe 100.100.100.1, routing-instance
 vpws3001, route-type Egress, vlan-id 40000
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302464             *[EVPN/7] 02:26:20, remote-pe 100.100.100.2, routing-instance
 vpws3001, route-type Egress, vlan-id 40000
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302480             *[EVPN/7] 02:26:14
                    > via ge-0/0/1.3016, Pop
302512             *[EVPN/7] 02:26:14, remote-pe 100.100.100.1, routing-instance
 vpws3016, route-type Egress, vlan-id 40016
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302528             *[EVPN/7] 02:26:14, remote-pe 100.100.100.2, routing-instance
 vpws3016, route-type Egress, vlan-id 40016
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302560             *[EVPN/7] 02:26:06
                    > via ae10.3011, Pop
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302592             *[EVPN/7] 02:26:07, remote-pe 100.100.100.1, routing-instance
 vpws3011, route-type Egress, vlan-id 401100
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
302608             *[EVPN/7] 02:26:07, remote-pe 100.100.100.2, routing-instance
 vpws3011, route-type Egress, vlan-id 401100
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302624             *[EVPN/7] 02:26:07, remote-pe 100.100.100.3, routing-instance
 vpws3011, route-type Egress Protection, vlan-id 301100
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302656             *[EVPN/7] 02:25:59
                    > via ae10.3006, Pop
                    to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302688             *[EVPN/7] 02:26:00, remote-pe 100.100.100.2, routing-instance
 vpws3006, route-type Egress, vlan-id 400600
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2
302704             *[EVPN/7] 02:26:00, remote-pe 100.100.100.1, routing-instance
 vpws3006, route-type Egress, vlan-id 400600
                  > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
```


Copyright © 2017, Juniper Networks, Inc.

```
302720             *[EVPN/7] 02:25:59, remote-pe 100.100.100.3, routing-instance
 vpws3006, route-type Egress, vlan-id 400600
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
302736             *[EVPN/7] 02:25:59, remote-pe 100.100.100.3, routing-instance
 vpws3006, route-type Egress Protection, vlan-id 300600
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ge-0/0/1.1001      *[EVPN/7] 02:29:23
                    > via ge-0/0/1.2001
ge-0/0/1.2001      *[EVPN/7] 02:29:23
                    > via ge-0/0/1.1001
ge-0/0/1.1002      *[EVPN/7] 02:27:06
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ae10.2003          *[EVPN/7] 02:29:10
                    > via ge-0/0/1.1003
ge-0/0/1.1003      *[EVPN/7] 02:27:06
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3

                     > via ae10.2003
ge-0/0/1.1004      *[EVPN/7] 02:27:06
                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ae10.1010          *[EVPN/7] 02:27:06
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ge-0/0/1.3001      *[EVPN/7] 02:26:20
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
ge-0/0/1.3016      *[EVPN/7] 02:26:13
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ae10.3011          *[EVPN/7] 02:26:06
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1
ae10.3006          *[EVPN/7] 02:25:59
                   > to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe1

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe2

                     to 100.46.1.2 via ge-0/0/5.0, label-switched-path pe4_to_pe3
```

show route table mpls.0 protocol ospf

```
user@host> show route table mpls.0 protocol ospf
mpls.0: 29 destinations, 29 routes (29 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

299952             *[L-OSPF/10] 23:59:42, metric 0
                    > to 10.0.10.70 via lt-1/2/0.14, Pop
                      to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299952(S=0)        *[L-OSPF/10] 23:59:42, metric 0
                    > to 10.0.10.70 via lt-1/2/0.14, Pop
                      to 10.0.6.60 via lt-1/2/0.12, Swap 800070, Push 800030(top)
299968             *[L-OSPF/10] 23:59:48, metric 0
                    > to 10.0.6.60 via lt-1/2/0.12, Pop
```

Copyright © 2017, Juniper Networks, Inc.

show route table mpls.0 extensive (PTX Series)

```
user@host> show route table mpls.0 extensive
ge-0/0/2.600 (1 entry, 1 announced)
TSI:
KRT in-kernel ge-0/0/2.600.0      /32 -> {composite(570)}
        *L2VPN  Preference: 7
                Next hop type: Indirect
                Address: 0x9438f34
                Next-hop reference count: 2
                Next hop type: Router, Next hop index: 567
                Next hop: 10.0.0.1 via ge-0/0/1.0, selected
                Label operation: Push 299808
                Label TTL action: prop-ttl
                Load balance label: Label 299808:None;
                Session Id: 0x1
                Protocol next hop: 10.255.255.1
                Label operation: Push 299872 Offset: 252
                Label TTL action: no-prop-ttl
                Load balance label: Label 299872:Flow label PUSH;
                Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                State: <Active Int>
                Age: 47         Metric2: 1
                Validation State: unverified
                Task: Common L2 VC
                Announcement bits (2): 0-KRT 2-Common L2 VC
                AS path: I
                Composite next hops: 1
                        Protocol next hop: 10.255.255.1 Metric: 1
                        Label operation: Push 299872 Offset: 252
                        Label TTL action: no-prop-ttl
                        Load balance label: Label 299872:Flow label PUSH;
                        Composite next hop: 0x9438ed8 570 INH Session ID: 0x2
                        Indirect next hop: 0x9448208 262142 INH Session ID: 0x2
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.0.0.1 via ge-0/0/1.0
                                Session Id: 0x1
                        10.255.255.1/32 Originating RIB: inet.3
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 1
                                Nexthop: 10.0.0.1 via ge-0/0/1.0
```

show route table mpls.0 (RSVP Route—Transit LSP)

```
user@host> show route table mpls.0

mpls.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

0                  *[MPLS/0] 00:37:31, metric 1
                      Receive
1                  *[MPLS/0] 00:37:31, metric 1
                      Receive
2                  *[MPLS/0] 00:37:31, metric 1
                      Receive
13                 *[MPLS/0] 00:37:31, metric 1
                      Receive
300352             *[RSVP/7/1] 00:08:00, metric 1
```


Copyright © 2017, Juniper Networks, Inc.

```
                    > to 10.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300352(S=0)         *[RSVP/7/1] 00:08:00, metric 1
                    > to 10.64.0.106 via ge-1/0/1.0, label-switched-path lsp1_p2p
300384              *[RSVP/7/2] 00:05:20, metric 1
                     > to 10.64.1.106 via ge-1/0/0.0, Pop
300384(S=0)         *[RSVP/7/2] 00:05:20, metric 1
                     > to 10.64.1.106 via ge-1/0/0.0, Pop
```

show route table vpls_1 detail

```
user@host> show route table vpls_1 detail
vpls_1.l2vpn.0: 1 destinations, 1 routes (1 active, 0 holddown, 0 hidden)
Restart Complete

172.16.1.11:1000:1:1/96 (1 entry, 1 announced)
*L2VPN Preference: 170/-1
Receive table: vpls_1.l2vpn.0
Next-hop reference count: 2
State: <Active Int Ext>
Age: 4:29:47 Metric2: 1
Task: vpls_1-l2vpn
Announcement bits (1): 1-BGP.0.0.0.0+179
AS path: I
Communities: Layer2-info: encaps:VPLS, control flags:Site-Down
Label-base: 800000, range: 8, status-vector: 0xFF
```

show route table vpn-a

```
user@host> show route table vpn-a
vpn-a.l2vpn.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)

+ = Active Route, - = Last Active, * = Both
192.168.16.1:1:1:1/96
                    *[VPN/7] 05:48:27
                      Discard
192.168.24.1:1:2:1/96
                    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                       AS path: I
                     > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
192.168.24.1:1:3:1/96
                    *[BGP/170] 00:02:53, localpref 100, from 192.168.24.1
                       AS path: I
                     > to 10.0.16.2 via fe-0/0/1.0, label-switched-path am
```

show route table vpn-a.mdt.0

```
user@host> show route table vpn-a.mdt.0
vpn-a.mdt.0: 3 destinations, 3 routes (3 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:1:0:10.255.14.216:232.1.1.1/144
                    *[MVPN/70] 01:23:05, metric2 1
                       Indirect
1:1:1:10.255.14.218:232.1.1.1/144
                    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.218
                       AS path: I
                     > via so-0/0/0.0, label-switched-path r0e-to-r1
1:1:2:10.255.14.217:232.1.1.1/144
                    *[BGP/170] 00:57:49, localpref 100, from 10.255.14.217
```

Copyright © 2017, Juniper Networks, Inc.

```
                        AS path: I
> via so-0/0/1.0, label-switched-path r0-to-r2
```

show route table VPN-A detail

```
user@host> show route table VPN-A detail
VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
10.255.179.9/32 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 10.255.179.13:200
                Next hop type: Indirect
                Next-hop reference count: 5
                Source: 10.255.179.13
                Next hop type: Router, Next hop index: 732
                Next hop: 10.39.1.14 via fe-0/3/0.0, selected
                Label operation: Push 299824, Push 299824(top)
                Protocol next hop: 10.255.179.13
                Push 299824
                Indirect next hop: 8f275a0 1048574
                State: (Secondary Active Int Ext)
                Local AS: 1 Peer AS: 1
                Age: 3:41:06 Metric: 1 Metric2: 1
                Task: BGP_1.10.255.179.13+64309
                Announcement bits (2): 0-KRT 1-BGP RT Background
                AS path: I
                Communities: target:1:200 rte-type:0.0.0.0:1:0
                Import Accepted
                VPN Label: 299824 TTL Action: vrf-ttl-propagate
                Localpref: 100
                Router ID: 10.255.179.13
                Primary Routing Table bgp.l3vpn.0
```

show route table VPN-AB.inet.0

```
user@host> show route table VPN-AB.inet.0
      VPN-AB.inet.0: 8 destinations, 8 routes (8 active, 0 holddown, 0 hidden)
      + = Active Route, - = Last Active, * = Both

      10.39.1.0/30       *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
      10.39.1.4/30       *[Direct/0] 00:08:42
                          > via so-5/1/0.0
      10.39.1.6/32       *[Local/0] 00:08:46
                            Local
      10.255.71.16/32    *[Static/5] 00:07:24
                          > via so-2/0/0.0
      10.255.71.17/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                             AS path: I
                          > via so-2/1/0.0, Push 100020, Push 100011(top)
      10.255.71.18/32    *[BGP/170] 00:07:24, MED 1, localpref 100, from
      10.255.71.15
                             AS path: I
                          > via so-2/1/0.0, Push 100021, Push 100011(top)
      10.255.245.245/32  *[BGP/170] 00:08:35, localpref 100
                             AS path: 2 I
                          > to 10.39.1.5 via so-5/1/0.0
      10.255.245.246/32  *[OSPF/10] 00:07:24, metric 1
                          > via so-7/3/1.0
```


Copyright © 2017, Juniper Networks, Inc.

show route table VPN_blue.mvpn-inet6.0

```
user@host> show route table VPN_blue.mvpn-inet6.0
vpn_blue.mvpn-inet6.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

1:10.255.2.202:65536:10.255.2.202/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                    > via so-0/1/3.0
1:10.255.2.203:65536:10.255.2.203/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.203
                      AS path: I
                    > via so-0/1/0.0
1:10.255.2.204:65536:10.255.2.204/432
                   *[MVPN/70] 00:57:23, metric2 1
                      Indirect
5:10.255.2.202:65536:128:::192.168.90.2:128:ffff::1/432
                   *[BGP/170] 00:02:37, localpref 100, from 10.255.2.202
                      AS path: I
                    > via so-0/1/3.0
6:10.255.2.203:65536:64500:128:::10.12.53.12:128:ffff::1/432
                   *[PIM/105] 00:02:37
                      Multicast (IPv6)
7:10.255.2.202:65536:64500:128:::192.168.90.2:128:ffff::1/432
                   *[MVPN/70] 00:02:37, metric2 1
                      Indirect
```

show route table vrf1.mvpn.0 extensive

```
user@host> show route table vrf1.mvpn.0 extensive
1:10.255.50.77:1:10.255.50.77/240 (1 entry, 1 announced)
        *MVPN   Preference: 70
                PMSI: Flags 0x0: Label 0: RSVP-TE:
Session_13[10.255.50.77:0:25624:10.255.50.77]
                Next hop type: Indirect
                Address: 0xbb2c944
                Next-hop reference count: 360
                Protocol next hop: 10.255.50.77
                Indirect next hop: 0x0 - INH Session ID: 0x0
                State: <Active Int Ext>
                Age: 53:03      Metric2: 1
                Validation State: unverified
                Task: mvpn global task
                Announcement bits (3): 0-PIM.vrf1 1-mvpn global task 2-rt-export

                AS path: I
```

show route table inetflow detail

```
user@host> show route table inetflow detail
inetflow.0: 2 destinations, 2 routes (2 active, 0 holddown, 0 hidden)
10.12.44.1,*/48 (1 entry, 1 announced)
        *BGP    Preference: 170/-101
                Next-hop reference count: 2
                State: <Active Ext>
                Local AS: 64502 Peer AS: 64500
                Age: 4
                Task: BGP_64500.10.12.99.5+3792
                Announcement bits (1): 0-Flow
```

Copyright © 2017, Juniper Networks, Inc.

```
                AS path: 64500 I
                Communities: traffic-rate:0:0
                Validation state: Accept, Originator: 10.12.99.5
                Via: 10.12.44.0/24, Active
                Localpref: 100
                Router ID: 10.255.71.161

10.12.56.1,*/48 (1 entry, 1 announced)
        *Flow   Preference: 5
                Next-hop reference count: 2
                State: <Active>
                Local AS: 64502
                Age: 6:30
                Task: RT Flow
                Announcement bits (2): 0-Flow 1-BGP.0.0.0.0+179
                AS path: I
                Communities: 1:1


user@host> show route table green.l2vpn.0 (VPLS Multihoming with FEC 129)
green.l2vpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

10.1.1.2:100:10.1.1.2/96 AD
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
10.1.1.4:100:10.1.1.4/96 AD
                   *[BGP/170] 1d 03:11:02, localpref 100, from 10.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
10.1.1.2:100:1:0/96 MH
                   *[VPLS/170] 1d 03:11:03, metric2 1
                         Indirect
10.1.1.4:100:1:0/96 MH
                   *[BGP/170] 1d 03:11:02, localpref 100, from 10.1.1.4
                      AS path: I, validation-state: unverified
                    >    via ge-1/2/1.5
10.1.1.4:NoCtrlWord:5:100:100:10.1.1.2:10.1.1.4/176
                   *[VPLS/7] 1d 03:11:02, metric2 1
                    >    via ge-1/2/1.5
10.1.1.4:NoCtrlWord:5:100:100:10.1.1.4:10.1.1.2/176
                   *[LDP/9] 1d 03:11:02
                         Discard


user@host> show route table red extensive
red.inet.0: 364481 destinations, 714087 routes (364480 active, 48448 holddown, 1
 hidden)
10.0.0.0/32 (3 entries, 1 announced)
        State: <OnList CalcForwarding>
TSI:
KRT in-kernel 10.0.0.0/32 -> {composite(1048575)} Page 0 idx 1 Type 1 val 0x934342c

    Nexthop: Self
    AS path: [2] I
    Communities: target:2:1
Path 10.0.0.0 from 10.3.0.0 Vector len 4.  Val: 1
        @BGP    Preference: 170/-1
                Route Distinguisher: 2:1
                Next hop type: Indirect
                Address: 0x258059e4
                Next-hop reference count: 2
```


Copyright © 2017, Juniper Networks, Inc.

```
                    Source: 2.2.0.0
                    Next hop type: Router
                    Next hop: 10.1.1.1 via ge-1/1/9.0, selected
                    Label operation: Push 707633
                    Label TTL action: prop-ttl
                    Session Id: 0x17d8
                    Protocol next hop: 10.2.0.0
                    Push 16
                    Composite next hop: 0x25805988 - INH Session ID: 0x193c
                    Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                    State: <Secondary Active Int Ext ProtectionPath ProtectionCand>
                    Local AS:     2 Peer AS:     2
                    Age: 23         Metric2: 35
                    Validation State: unverified
                    Task: BGP_172.16.2.0.0+34549
                    AS path: I
                    Communities: target:2:1
                    Import Accepted
                    VPN Label: 16
                    Localpref: 0
                    Router ID: 10.2.0.0
                    Primary Routing Table bgp.l3vpn.0
                    Composite next hops: 1
                            Protocol next hop: 10.2.0.0 Metric: 35
                            Push 16
                            Composite next hop: 0x25805988 - INH Session ID: 0x193c
                            Indirect next hop: 0x23eea900 - INH Session ID: 0x193c
                            Indirect path forwarding next hops: 1
                                    Next hop type: Router
                                    Next hop: 10.1.1.1 via ge-1/1/9.0
                                    Session Id: 0x17d8
                            2.2.0.0/32 Originating RIB: inet.3
                              Metric: 35                      Node path count: 1
                              Forwarding nexthops: 1
                                    Nexthop: 10.1.1.1 via ge-1/1/9.0
            BGP     Preference: 170/-1
                    Route Distinguisher: 2:1
                    Next hop type: Indirect
                    Address: 0x9347028
                    Next-hop reference count: 3
                    Source: 10.3.0.0
                    Next hop type: Router, Next hop index: 702
                    Next hop: 10.1.4.2 via ge-1/0/0.0, selected
                    Label operation: Push 634278
                    Label TTL action: prop-ttl
                    Session Id: 0x17d9
                    Protocol next hop: 10.3.0.0
                    Push 16
                    Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
                    Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da
                    State: <Secondary NotBest Int Ext ProtectionPath ProtectionCand>

                    Inactive reason: Not Best in its group - IGP metric
                    Local AS:     2 Peer AS:     2
                    Age: 3:34       Metric2: 70
                    Validation State: unverified
                    Task: BGP_172.16.3.0.0+32805
                    Announcement bits (2): 0-KRT 1-BGP_RT_Background
                    AS path: I
                    Communities: target:2:1
                    Import Accepted
```

```
                    VPN Label: 16
                    Localpref: 0
                    Router ID: 10.3.0.0
                    Primary Routing Table bgp.l3vpn.0
                    Composite next hops: 1
                            Protocol next hop: 10.3.0.0 Metric: 70
                            Push 16
                            Composite next hop: 0x93463a0 1048575 INH Session ID:
0x17da
                            Indirect next hop: 0x91e8800 1048574 INH Session ID:
0x17da
                            Indirect path forwarding next hops: 1
                                    Next hop type: Router
                                    Next hop: 10.1.4.2 via ge-1/0/0.0
                                    Session Id: 0x17d9
                            10.3.0.0/32 Originating RIB: inet.3
                              Metric: 70                    Node path count: 1
                              Forwarding nexthops: 1
                                    Nexthop: 10.1.4.2 via ge-1/0/0.0
                #Multipath Preference: 255
                    Next hop type: Indirect
                    Address: 0x24afca30
                    Next-hop reference count: 1
                    Next hop type: Router
                    Next hop: 10.1.1.1 via ge-1/1/9.0, selected
                    Label operation: Push 707633
                    Label TTL action: prop-ttl
                    Session Id: 0x17d8
                    Next hop type: Router, Next hop index: 702
                    Next hop: 10.1.4.2 via ge-1/0/0.0
                    Label operation: Push 634278
                    Label TTL action: prop-ttl
                    Session Id: 0x17d9
                    Protocol next hop: 10.2.0.0
                    Push 16
                    Composite next hop: 0x25805988 - INH Session ID: 0x193c
                   Indirect next hop: 0x23eea900 - INH Session ID: 0x193c Weight 0x1

                    Protocol next hop: 10.3.0.0
                    Push 16
                    Composite next hop: 0x93463a0 1048575 INH Session ID: 0x17da
                  Indirect next hop: 0x91e8800 1048574 INH Session ID: 0x17da Weight
 0x4000
                    State: <ForwardingOnly Int Ext>
                    Inactive reason: Forwarding use only
                    Age: 23         Metric2: 35
                    Validation State: unverified
                    Task: RT
                    AS path: I
                    Communities: target:2:1
```

show route table bgp.evpn.0 extensive |no-more (EVPN)

```
show route table bgp.evpn.0 extensive | no-more
bgp.evpn.0: 6 destinations, 6 routes (6 active, 0 holddown, 0 hidden)
2:1000:10::100::00:aa:aa:aa:aa:aa/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
```


Copyright © 2017, Juniper Networks, Inc.

```
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
         State: Local AS: 17  Peer AS:17  Age:21:12  Metric2:1  Validation State:
 unverified
                Task: BGP_17.1.2.3.4+50756
                AS path: I
                Communities: target:1111:8388708 encapsulation0:0:0:0:3
                Import Accepted
                Route Label: 100
                ESI: 00:00:00:00:00:00:00:00:00:00
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        1.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

2:1000:10::200::00:bb:bb:bb:bb:bb/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS:17  Peer AS:17  Age:19:43  Metric2:1  Validation
State:unverified
                Task: BGP_17.1.2.3.4+50756
                AS path: I
                Communities: target:2222:22 encapsulation0:0:0:0:3
                Import Accepted
                Route Label: 200
                ESI: 00:00:00:00:00:00:00:00:00:00
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

2:1000:10::300::00:cc:cc:cc:cc:cc/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
```

```
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS:17  Peer AS:17  Age:17:21  Metric2:1  Validation State:
 unverified  Task: BGP 17,1,2,3,4+50756
           AS path: I
                Communities: target:3333:33 encapsulation0:0:0:0:3
                Import Accepted
                Route Label: 300
                ESI: 00:00:00:00:00:00:00:00:00:00
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::100::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                PMSI: Flags 0x0: Label 100: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS:17  Peer AS:17 Age:37:01  Metric2:1  Validation State:
 unverified  Task: BGP 17.1.2.3.4+50756
                AS path: I
                Communities: target:1111:8388708 encapsulation0:0:0:0:3
                Import Accepted
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::200::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
```


Copyright © 2017, Juniper Networks, Inc.

```
                PMSI: Flags 0x0: Label 200: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
          State: Local AS: 17 Peer AS: 17   Age:35:22  Metric2:1  Validation
State:unverified  Task: BGP 17.1.2.3.4+50756
           AS path:I  Communities: target:2222:22 encapsulation):0:0:0:0:3

Import Accepted
                Localpref: 100
                Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0

3:1000:10::300::1.2.3.4/304 (1 entry, 0 announced)
        *BGP    Preference: 170/-101
                Route Distinguisher: 1000:10
                PMSI: Flags 0x0: Label 300: Type INGRESS-REPLICATION 1.2.3.4
                Next hop type: Indirect
                Address: 0x9420fd0
                Next-hop reference count: 12
                Source: 10.2.3.4
                Protocol next hop: 10.2.3.4
                Indirect next hop: 0x2 no-forward INH Session ID: 0x0
         State: Local AS: 17 Peer AS:  17  Age 35:22 Metric2:1  Validation State:
 unverified   Task: BGP 17.1.2.3.4+5075
          6  AS path: I   Communities: target:3333:33  encapsulation0:0:0:0:3
Import Accepted   Localpref:100
          Router ID: 10.2.3.4
                Secondary Tables: default-switch.evpn.0
                Indirect next hops: 1
                        Protocol next hop: 10.2.3.4 Metric: 1
                        Indirect next hop: 0x2 no-forward INH Session ID: 0x0
                        Indirect path forwarding next hops: 1
                                Next hop type: Router
                                Next hop: 10.10.10.1 via xe-0/0/1.0
                                Session Id: 0x2
                        10.2.3.4/32 Originating RIB: inet.0
                          Metric: 1                     Node path count: 1
                          Forwarding nexthops: 2
                                Nexthop: 10.92.78.102 via em0.0
```

Copyright © 2017, Juniper Networks, Inc.

show route terse

List of Syntax Syntax on page 666
Syntax (EX Series Switches) on page 666

Syntax
show route terse
<logical-system (all | *logical-system-name*)>

Syntax (EX Series Switches)
show route terse

Release Information
Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Display a high-level summary of the routes in the routing table.

NOTE: For BGP routes, the **show route terse** command displays the local preference attribute and MED instead of the metric1 and metric2 values. This is mostly due to historical reasons.

To display the metric1 and metric2 value of a BGP route, use the show route extensive command.

Options **none**—Display a high-level summary of the routes in the routing table.

logical-system (all | *logical-system-name*)—(Optional) Perform this operation on all logical systems or on a particular logical system.

Required Privilege Level
view

List of Sample Output show route terse on page 668

Output Fields Table 30 on page 447 describes the output fields for the **show route terse** command. Output fields are listed in the approximate order in which they appear.

Table 51: show route terse Output Fields

Field Name	Field Description
routing-table-name	Name of the routing table (for example, inet.0).
number destinations	Number of destinations for which there are routes in the routing table.

Copyright © 2017, Juniper Networks, Inc.

Table 51: show route terse Output Fields *(continued)*

Field Name	Field Description
number routes	Number of routes in the routing table and total number of routes in the following states: • **active** (routes that are active) • **holddown** (routes that are in the pending state before being declared inactive) • **hidden** (routes that are not used because of a routing policy)
route key	Key for the state of the route: • **+**—A plus sign indicates the active route, which is the route installed from the routing table into the forwarding table. • **-** —A hyphen indicates the last active route. • *****—An asterisk indicates that the route is both the active and the last active route. An asterisk before a **to** line indicates the best subpath to the route.
A	Active route. An asterisk (*) indicates this is the active route.
V	Validation status of the route: • **?**—Not evaluated. Indicates that the route was not learned through BGP. • **I**—Invalid. Indicates that the prefix is found, but either the corresponding AS received from the EBGP peer is not the AS that appears in the database, or the prefix length in the BGP update message is longer than the maximum length permitted in the database. • **N**—Unknown. Indicates that the prefix is not among the prefixes or prefix ranges in the database. • **V**—Valid. Indicates that the prefix and autonomous system pair are found in the database.
Destination	Destination of the route.
P	Protocol through which the route was learned: • **A**—Aggregate • **B**—BGP • **C**—CCC • **D**—Direct • **G**—GMPLS • **I**—IS-IS • **L**—L2CKT, L2VPN, LDP, Local • **K**—Kernel • **M**—MPLS, MSDP • **O**—OSPF • **P**—PIM • **R**—RIP, RIPng • **S**—Static • **T**—Tunnel
Prf	Preference value of the route. In every routing metric except for the BGP **LocalPref** attribute, a lesser value is preferred. In order to use common comparison routines, Junos OS stores the 1's complement of the **LocalPref** value in the **Preference2** field. For example, if the **LocalPref** value for Route 1 is 100, the **Preference2** value is -101. If the **LocalPref** value for Route 2 is 155, the **Preference2** value is -156. Route 2 is preferred because it has a higher **LocalPref** value and a lower **Preference2** value.

Copyright © 2017, Juniper Networks, Inc.

Table 51: show route terse Output Fields *(continued)*

Field Name	Field Description
Metric 1	First metric value in the route. For routes learned from BGP, this is the MED metric.
Metric 2	Second metric value in the route. For routes learned from BGP, this is the IGP metric.
Next hop	Next hop to the destination. An angle bracket (>) indicates that the route is the selected route.
AS path	AS path through which the route was learned. The letters at the end of the AS path indicate the path origin, providing an indication of the state of the route at the point at which the AS path originated: • **I**—IGP. • **E**—EGP. • **?**—Incomplete; typically, the AS path was aggregated.

Sample Output

show route terse

```
user@host> show route terse
inet.0: 10 destinations, 12 routes (10 active, 0 holddown, 0 hidden)
+ = Active Route, - = Last Active, * = Both

A V Destination        P Prf   Metric 1   Metric 2  Next hop        AS path
* ? 172.16.1.1/32        O  10          1            >10.0.0.2
  ?                     B 170        100                          I
  unverified                                      >10.0.0.2
* ? 172.16.1.1/32        D   0                       >lo0.2
* V 2.2.0.2/32          B 170        110                          200 I
  valid                                           >10.0.0.2
* ? 10.0.0.0/30         D   0                       >lt-1/2/0.1
  ?                     B 170        100                          I
  unverified                                      >10.0.0.2
* ? 10.0.0.1/32         L   0                        Local
* ? 10.0.0.4/30         B 170        100                          I
  unverified                                      >10.0.0.2
* ? 10.0.0.8/30         B 170        100                          I
  unverified                                      >10.0.0.2
* I 172.16.1.1/32       B 170         90                          200 I
  invalid                                         >10.0.0.2
* N 192.168.2.3/32      B 170        100                          200 I
  unknown                                         >10.0.0.2
* ? 172.16.233.5/32       O  10          1             MultiRecv
```


Copyright © 2017, Juniper Networks, Inc.

test policy

Syntax test policy *policy-name prefix*

Release Information Command introduced before Junos OS Release 7.4.
Command introduced in Junos OS Release 9.0 for EX Series switches.

Description Test a policy configuration to determine which prefixes match routes in the routing table.

NOTE: If you are using the **test policy** command on a logical system, you must first set the CLI to the logical system context. For example, if you want to test a routing policy that is configured on logical system R2, first run the **set cli logical-system R2** command.

Options ***policy-name***—Name of a policy.

prefix—Destination prefix to match.

Additional Information All prefixes in the default unicast routing table (inet.0) that match prefixes that are the same as or longer than the specific prefix are processed by the **from** clause in the specified policy. All prefixes accepted by the policy are displayed. The **test policy** command evaluates a policy differently from the BGP import process. When testing a policy that contains an **interface** match condition in the **from** clause, the **test policy** command uses the match condition. In contrast, BGP does not use the **interface** match condition when evaluating the policy against routes learned from internal BGP (IBGP) or external BGP (EGBP) multihop peers.

When testing a policy, you can see the length of time (in microseconds) required to evaluate the policy and the number of times it has been executed by running the show policy ***policy-name*** **statistics** command.

Required Privilege Level view

Related Documentation

- *Understanding Routing Policy Tests*
- Example: Testing a Routing Policy with Complex Regular Expressions on page 129
- show policy on page 248

List of Sample Output test policy on page 670

Copyright © 2017, Juniper Networks, Inc.

Output Fields For information about output fields, see the output field tables for the show route command, the show route detail command, the show route extensive command, or the show route terse command.

Sample Output

test policy

```
user@host> test policy test-statics 172.16.0.1/8
inet.0: 44 destinations, 44 routes (44 active, 0 holddown, 0 hidden)
Prefixes passing policy:

172.16.3.0/8          *[BGP/170] 16:22:46, localpref 100, from 10.255.255.41
                    AS Path: 50888 I
                     > to 10.11.4.32 via en0.2, label-switched-path l2
172.16.3.1/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
172.16.3.2/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
172.16.3.3/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
172.16.3.4/32         *[IS-IS/18] 2d 00:21:46, metric 0, tag 2
                     > to 10.0.4.7 via fxp0.0
Policy test-statics: 5 prefixes accepted, 0 prefixes rejected
```


Copyright © 2017, Juniper Networks, Inc.